Accurate Cosmology with Galaxy and Quasar Surveys

Boris Leistedt

Submitted for the degree of Doctor of Philosophy Department of Physics and Astronomy University College London September 2014 I, *Boris Leistedt*, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Observations of the cosmic microwave background have led to a golden age of cosmology, where precise measurements can be confronted with predictions from cosmological models. Ongoing and future surveys of the distribution of galaxies will continue this revolution: they will enable us to test the laws of gravity, uncover the properties of dark energy and dark matter, and reinforce the connection to high-energy physics. However, current galaxy survey analyses are already limited by our ability to identify and treat observational systematics, and this problem will be even more pronounced in future experiments. Therefore, it is essential to develop novel methods to deal with these complications when testing cosmological models and searching for new physics. This is the focus of this thesis.

Firstly, I will present measurements of primordial non-Gaussianity obtained from the clustering of quasars from the Sloan Digital Sky Survey. Primordial non-Gaussianity is a powerful probe of inflation, the leading theory of the initial conditions of the universe, but its effects on the distribution of quasars are mimicked by observational systematics. I will describe a framework to deal with these systematics and robustly measure primordial non-Gaussianity from the clustering of quasars.

Secondly, I will present a new set of wavelet transforms on the sphere and the ball. These approaches are highly promising for analysing cosmological and geophysical data and dealing with their systematics in novel ways.

Finally, I will examine the recent claims that extra massive neutrinos can resolve the tensions between cosmic microwave background, galaxy survey and supernova observations. I will demonstrate that this conclusion is premature since it is driven by the least robust data sets. Given the growing number of cosmological observables and their varied levels of robustness, combining data sets and dealing with such tensions will become critical in the near future.

Acknowledgements

First and foremost, I would like to express my deepest gratitude and appreciation to my advisor, Hiranya Peiris. During the past three years, you have been an attentive, supportive, and inspiring mentor, and I am extremely grateful for the infinite amount of time and energy you have invested in me. My time in UCL has been a life changing experience, and I could never hope for a better supervisor.

I would also like to thank the exceptional individuals I have had the chance to collaborate with: Jason McEwen, Daniel Mortlock, Andrew Pontzen, Aurélien Benoit-Lévy, Licia Verde, Nina Roth, and Franz Elsner (in order of appearance). I have learned so much from you, both personally and scientifically, and you made me discover how fun, fascinating, and humanly enriching research could be.

Thank you to Ofer Lahav and Filipe Adballa for the useful discussions and advice during the course of this PhD. I am also grateful to Alan Heavens and Benjamin Joachimi, who kindly accepted to examine this thesis and provide many useful comments and suggestions.

My experience in UCL would not have been complete without the UCL astro-p(i)gs. For all the love, laughs and tears, I thank all the students and post-docs of the Astrophysics Group. Special thoughts go to the inhabitants of G16, this parallel universe where I worked during the three years of my PhD. Thanks to you I have enjoyed every single day in the office.

I am indebted to my parents and my brothers for their unconditional backing over the years. Thank you for encouraging me and providing me everything I needed to follow my dreams.

And finally, thank you Olivier, for standing by my side from the very first day, and making my life so colourful. Without your love and support, this document would simply be blank.

Contents

1	Intr	oduction	15
	1.1	The Universe	15
		1.1.1 Composition and structure	15
		1.1.2 Brief history	19
	1.2	The early Universe	20
		1.2.1 The cosmic microwave background	20
		1.2.2 Inflation and primordial non-Gaussianity	22
	1.3	The evolved Universe	25
		1.3.1 Matter perturbations and structure formation	26
		1.3.2 Galaxy surveys	27
		1.3.3 Scale-dependent halo and galaxy bias	29
		1.3.4 Quasars	31
		1.3.5 Cosmological neutrinos	32
	1.4	Wavelets in cosmology	34
	1.5	Thesis outline	35
2	Esti	nating the large-scale angular power spectrum in the presence of systematics: a	
	case	study of Sloan Digital Sky Survey quasars	37
	2.1	Abstract	37
	2.2	Introduction	38
	2.3	Theory and methods	40
		2.3.1 Background	40

		2.3.2	Partial sky coverage and quadratic estimators	41
		2.3.3	Likelihood analysis and band-powers	46
		2.3.4	Galaxy surveys, shot noise and systematics	50
		2.3.5	Illustration: recovering the power spectrum of the CMASS sample	51
	2.4	Applic	ation to SDSS photometric quasars	54
		2.4.1	Data and subsamples	54
		2.4.2	Theory predictions	56
		2.4.3	Redshift distributions estimates	58
		2.4.4	Masks and systematics	62
		2.4.5	Power spectrum results	64
	2.5	Discus	sion	68
	2.6	Aknov	vledgements	72
	2.7	Apper	ndices	72
		2.7.1	Pixelisation, band-limit and smoothing issues	72
		2.7.2	Karhunen-Loève compression demystified	74
		2.7.3	χ^2 for band-power estimates	76
3	Exp	loiting	the full potential of photometric quasar surveys: optimal power spectra	
3	Expl thro	loiting ough bl	the full potential of photometric quasar surveys: optimal power spectra	י 78
3	Expl thro 3.1	loiting ough bl Abstra	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics	78 78
3	Exp ¹ thro 3.1 3.2	loiting ough bl Abstra Introd	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act	78 78 78 79
3	Exp2 thro 3.1 3.2 3.3	loiting ough bl Abstra Introd Theor	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act uction uction y and methods	78 78 78 79 81
3	Exp) thro 3.1 3.2 3.3	loiting ough bl Abstra Introd Theor 3.3.1	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act	78 78 79 81 81
3	Expl thro 3.1 3.2 3.3	loiting ough bl Abstra Introd Theor 3.3.1 3.3.2	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act	78 78 79 81 81 83
3	Exp) thro 3.1 3.2 3.3	loiting ough bl Abstra Introd 3.3.1 3.3.2 3.3.3	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act	78 78 79 81 81 83 83
3	Exp2 thro 3.1 3.2 3.3	loiting ough bl Abstra Introd 3.3.1 3.3.2 3.3.3 Applic	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act	78 78 79 81 81 83 87 88
3	Expl thro 3.1 3.2 3.3 3.3	loiting ough bl Abstra Introd 3.3.1 3.3.2 3.3.3 Applic 3.4.1	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act uction uction y and methods Power spectrum estimation Extended mode projection Fast cross-power spectra and null-tests eation to SDSS photometric quasars Sample selection and redshift distributions	78 78 79 81 83 83 87 88 89
3	Expl thro 3.1 3.2 3.3 3.4	loiting ough bl Abstra Introd Theor 3.3.1 3.3.2 3.3.3 Applic 3.4.1 3.4.2	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act	78 78 79 81 83 83 87 88 89 92
3	Expl thro 3.1 3.2 3.3 3.4	loiting ough bl Abstra Introd Theor 3.3.1 3.3.2 3.3.3 Applic 3.4.1 3.4.2 3.4.3	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act uction y and methods Power spectrum estimation Extended mode projection Fast cross-power spectra and null-tests scation to SDSS photometric quasars Sample selection and redshift distributions Masks and systematic uncertainties Estimation settings and blind mitigation of systematics	78 78 79 81 83 83 87 88 89 92 95
3	Expl thro 3.1 3.2 3.3 3.4	loiting Jugh bl Abstra Introd Theor 3.3.1 3.3.2 3.3.3 Applic 3.4.1 3.4.2 3.4.3 3.4.4	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act	78 78 79 81 83 83 87 88 89 92 92 95 97
3	Expl thro 3.1 3.2 3.3 3.4	loiting Abstra Abstra Introd Theor 3.3.1 3.3.2 3.3.3 Applic 3.4.1 3.4.2 3.4.3 3.4.4 Concle	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act uction uuction y and methods Power spectrum estimation Extended mode projection Fast cross-power spectra and null-tests sample selection and redshift distributions Masks and systematic uncertainties Estimation settings and blind mitigation of systematics Power spectrum measurements	78 78 79 81 83 83 87 88 89 92 95 97 101
3	Exp) thro 3.1 3.2 3.3 3.4 3.4	loiting Abstra Abstra Introd Theor 3.3.1 3.3.2 3.3.3 Applic 3.4.1 3.4.2 3.4.3 3.4.3 3.4.4 Conch Ackno	the full potential of photometric quasar surveys: optimal power spectra ind mitigation of systematics act uction y and methods Power spectrum estimation Extended mode projection Fast cross-power spectra and null-tests ration to SDSS photometric quasars Sample selection and redshift distributions Masks and systematic uncertainties Estimation settings and blind mitigation of systematics Power spectrum measurements usion	78 78 79 81 81 83 87 88 89 92 95 97 101 103

Con	straints on primordial non-Gaussianity from 800,000 photometric quasars	105
4.1	Abstract	105
4.2	Introduction	106
4.3	PNG with photometric quasars	107
4.4	XDQSOz power spectrum measurements	108
4.5	Halo bias from PNG	110
4.6	Monte Carlo Markov Chain (MCMC) analysis	111
4.7	Results	112
4.8	Future prospects	115
4.9	Conclusion	115
4.10	Acknowledgments	116
S2LE	ET: A code to perform fast wavelet analysis on the sphere	117
5.1	Abstract	117
5.2	Introduction	118
5.3	Wavelets on the sphere	119
	5.3.1 Harmonic analysis on the sphere	120
	5.3.2 Scale-discretised wavelets on the sphere	121
5.4	The S2LET code	125
	5.4.1 Multiresolution algorithm	125
	5.4.2 Implementation	127
	5.4.3 Numerical validation	128
	5.4.4 Future extensions	131
5.5	Examples	132
	5.5.1 Wavelet transform from the command line	132
	5.5.2 Wavelet transform in Matlab and IDL	133
	5.5.3 Wavelet denoising in C	133
5.6	Summary	136
5.7	Acknowledgements	136
Exa	ct Wavelets on the Ball	139
6.1	Abstract	139
6.2	Introduction	140
6.3	Harmonic Analysis on the Ball	141
	Con 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9 4.10 521 5.1 5.2 5.3 5.4 5.4 5.5 5.6 5.7 5.6 5.7 Exac 6.1 6.2 6.3	Constraints on primordial non-Gaussianity from 800,000 photometric quasars 4.1 Abstract 4.2 Introduction 4.3 PNG with photometric quasars 4.4 XDQSOz power spectrum measurements 4.4 XDQSOz power spectrum measurements 4.5 Halo bias from PNG 4.6 Monte Carlo Markov Chain (MCMC) analysis 4.7 Results 4.8 Future prospects 4.9 Conclusion 4.10 Acknowledgments S2LET: A code to perform fast wavelet analysis on the sphere 5.1 Abstract 5.2 Introduction 5.3 Wavelets on the sphere 5.3.1 Harmonic analysis on the sphere 5.3.2 Scale-discretised wavelets on the sphere 5.3.3 Numerical validation 5.4.4 Future extensions 5.5.5 Examples 5.5.1 Wavelet transform from the command line 5.5.2 Wavelet denoising in C 5.5.3 Wavelet enoising in C 5.5.3 Wavelet son the Ball 6.1 Abstract 6.2

		6.3.1 The spherical Laguerre transform	 142
		6.3.2 Relation to the spherical Bessel transform	 144
		6.3.3 The spherical harmonic transform	 147
		6.3.4 The Fourier-Laguerre transform	 148
	6.4	Wavelets on the Ball	 149
		6.4.1 Convolutions	 150
		6.4.2 Exact flaglet transform	 151
		6.4.3 Flaglets and scaling functions	 153
	6.5	Multiresolution Algorithm	 155
		6.5.1 Algorithm	 158
		6.5.2 Fast implementation	 158
		6.5.3 Numerical validation	 160
	6.6	Denoising Illustration	 161
		6.6.1 Wavelet denoising	 161
		6.6.2 Examples	 165
	6.7	Conclusions	 168
7	No 1	ew cosmological concordance with massive sterile neutrinos	171
	7.1	Abstract	 171
	7.2	Introduction	 171
	7.3	Data and methods	 173
	7.4	No new concordance with sterile neutrinos.	 174
	7.5	Constraints on neutrino masses from robust datasets	 176
	7.6	Conclusions	 178
	7.7	Acknowledgments	 180
Bi	bliog	aphy	180

List of Figures

1.1	Current picture of cosmic history, as suggested by cosmological observations and	
	explained by the concordance model.	20
1.2	The angular power spectra of the CMB temperature and polarization fluctuations	
	measured by <i>Planck</i> (Planck Collaboration, 2013i)	21
1.3	Example of an inflaton potential, from Baumann (2009)	23
1.4	Power spectrum and bispectrum in Fourier space	24
1.5	Matter power spectrum from a range of probes, reproduced from Hlozek et al. (2012),	
	where the data sets are described in detail	26
1.6	Left: measurements of the 3D power spectrum of galaxies from the Baryon Acoustic	
	Oscillation Survey (Anderson et al., 2012). Right: tomographic angular power	
	spectrum of luminous red galaxies from SDSS DR7 (Thomas et al., 2010)	29
1.7	Kaiser effect and scale-dependent halo bias due to primordial non-Gaussianity	30
1.8	Typical spectral energy distribution of a quasar, from Vanden Berk et al. (2001)	31
1.9	Effect of massive neutrinos on the matter and CMB power spectra, reproduced from	
	Park et al. (2012).	33
1.10	Example of "mexican hat" wavelets constructed to be localised in time and frequency	
	domains (Couderc and Zareba, 1998). These wavelets are redundant, <i>i.e.</i> , a given	
	scale or frequency is probed by more than one wavelet	34
2.1	Fiducial CMB and galaxy survey masks used to calculate the covariance matrices of	
	Fig. 2.3.	45

2.2	Fiducial CMB and CMASS angular power spectra used to calculate the covariance	
	matrices of Fig. 2.3.	46
2.3	Covariance matrices $ V_{\ell\ell'} $ of the PCL and QML estimates calculated for the CMB and	
	LSS spectra and masks in Figs. 2.1 and 2.2 in the absence of noise and systematics.	47
2.4	Mask of the CMASS DR9 sample in the Northern Galactic Cap	52
2.5	Average PCL and QML estimates of the 600 CMASS mock catalogues	53
2.6	Distributions of the photometric and spectroscopic redshift estimates of the RQCat	
	UVX-LOWZ sample, cross-matched with the SDSS-DR7, BOSS and 2SLAQ spectro-	
	scopic quasar catalogues.	56
2.7	Histograms of the redshifts and apparent magnitudes of UVX-LOWZ objects in RQCat,	
	and of objects with good spectra found in the SDSS-DR7, BOSS and 2SLAQ spectro-	
	scopic quasar catalogues.	57
2.8	Low-resolution redshift- and magnitude-dependent completeness corrections for	
	estimating the redshift distributions of the RQCat subsamples through cross-matching	
	with the SDSS-DR7, BOSS and 2SLAQ spectroscopic quasar catalogues	58
2.9	Final estimates of the redshift distributions of the four RQCat photometric subsamples,	
	calculated by cross-matching with the SDSS-DR7 spectroscopic quasar catalogue.	59
2.10	Theory predictions for the four RQCat subsamples computed with CAMB_SOURCES	63
2.11	Systematics templates used in this analysis, and the (dimensionless) angular power	
	spectra $ ilde{\mathcal{C}}_\ell$ of their overdensity maps.	63
2.12	Masks used for the power spectrum analysis of RQCat, in Equatorial coordinates.	64
2.13	Illustration of the suboptimality of the PCL estimator in the case of the Mid+High- z	
	subsample.	65
2.14	QML estimates of the (dimensionless) auto-power spectra of the overdensity maps of	
	the four RQCat samples presented in Fig. 2.9.	69
2.15	QML estimates of the (dimensionless) cross-power spectra of the RQCat overdensity	
	maps using the same conventions as Fig. 2.14	70
2.16	QML estimates for the (dimensionless) cross-power spectra of the four RQCat over-	
	density maps with the systematics templates using the same conventions as Fig. 2.14.	70
2.17	Beams that apply to the power spectrum estimates of pixelised, smoothed maps at	
	HEALPIX resolution N _{side}	74
2.18	Histograms of the χ^2 values of the PCL (left) and QML (right) estimates of the 600	
	CMASS mock catalogues.	76

3.1	Redshift distributions of the photometric quasar samples selected with $\mathrm{P}_{\mathrm{QSO}} > 0.8$	
	and top hat windows on the photometric redshift estimates \hat{z}_p	88
3.2	Templates of the main systematics selected as potential contaminants in the photo-	
	metric quasar catalogues, and used within the extended mode projection approach.	90
3.3	A subset of systematics templates originating from the decorrelation (using Mask 1)	
	of the basic set of systematics, a subset of which is presented in Fig. 3.2	91
3.4	Masks constructed for the power spectrum analysis of the XDQSOz samples	93
3.5	Comparison of the reduced χ^2 obtained with the optimal and approximate QML	
	cross-correlation estimators	96
3.6	Distributions of the reduced χ^2 arising from null tests performed with the cross-power	
	spectra of the four quasar samples with the ~ 3700 orthogonal modes of the extended	
	set of systematics templates.	96
3.7	Auto- and cross-angular power spectra of the four XDQSOz quasar samples measured	
	with the quadratic maximum likelihood (QML) estimator.	98
3.8	Same as Fig. 3.7, but using extended mode projection.	99
41	Redshift distributions of the four quasar samples used in this analysis parametrised	
7.1	as superpositions of normal distributions. The shaded regions show the $1/2$ and 3σ	
	regions explored when adding 5% uncertainty to the parameters of these fits, which	
	are included in the MCMC analysis	100
12	Constraints on the guasar bias model described in Eq. 4.4 . The solid line shows the	109
4.2	Fiducial model with $h_{-} = 1$, $\theta_{-} = 5$, and the shaded hands show the 1π constraints	
	inductal model with $b_0 = 1, \beta = 5$, and the shaded bands show the 1 σ constraints (b = 0.06 \pm 0.15, 68% CL β and f marginalized) from the VDOSOg power	
	$(b_0 = 0.90 \pm 0.15, 00\%$ CL, β and $f_{\rm NL}$ marginalised) from the ADQSOZ power	
	distributions. The coloured hands show the results when fiving θ_{1} , 5 and ellowing a	
	distributions. The coloured bands show the results when fixing $\beta = 5$ and anowing a different bias emplitude in each redshift bin to demonstrate the chility of the overall	
	model to simultaneously describe the four semples	110
4.0	Construction on local type (110
4.3	Constraints on local-type $f_{\rm NL}$ (in the ACDM+ $f_{\rm NL}$ model, with $n_{f_{\rm NL}} = g_{\rm NL} = 0$)	
	using the power spectrum analysis of XDQSOz quasars, for different bias models and	
	incorporating uncertainties in the redshift distributions and cosmological parameters.	
	The error bars show the 1 and 2σ constraints, the dashed line shows $f_{\rm NL} = 0$, and	
	the shaded bands show the constraints from <i>Planck</i> (Planck Collaboration, 2013h).	111

4.4	1σ and 2σ joint constraints on $f_{\rm NL}$ and $g_{\rm NL}$ for the $(b_0,\beta)+n(z){\rm +cosmo}$ case,	
	<i>i.e.</i> , marginalising over the uncertainties in the cosmological parameters, redshift	
	distributions, and bias model	113
4.5	1σ and 2σ joint constraints on $\tilde{f}_{\rm NL}$ and $n_{f_{\rm NL}}$ for the extended model of Eq. 4.3,	
	marginalising over the same parameters as Fig. 4.4.	113
F 1	We vale to and easing function competence d with the coals discussion d (CD) models and	
5.1	Parling construction constructed with the scale-discretised (SD), needlet and	
	B-spline generating functions with parameters $\lambda = 3$ and $J_0 = 2$ and for band-innit	104
- 0	$L = 128. \dots \dots$	124
5.2	Wavelets for scales $j \in \{2, 3, 4, 5\}$ and scaling function constructed through a tiling	105
	of the harmonic line using scale-discretised functions.	125
5.3	Scale-discretised wavelet transform of a band-limited topography map of the Earth	
	for $\lambda = 3$, $J_0 = 2$ and $L = 128$, i.e. with the scale-discretised wavelets shown in	
	Figure 5.2	126
5.4	Numerical accuracy and computation time of the scale-discretised wavelet transform	
	computed with S2LET	129
5.5	Wavelet denoising by hard-thresholding, using parameters $\lambda = 2$ and $J_0 = 0$ and	
	scale-discretised generating functions.	137
6.1	Spherical Laguerre sampling scheme on $r \in [0,1]$ for increasing band-limit	144
6.2	First six spherical Laguerre basis functions $K_p(r)$ constructed on $r \in [0,1]$ and the	
	associated sample positions	145
6.3	Slices of an axisymmetric flaglet wavelet kernel constructed on the ball of radius	
	R = 1, translated along the radial half-line.	152
6.4	Tiling of Fourier-Laguerre space at resolution $L = N = 64$ for flaglet parameters	
	$\lambda = \nu = 2$, giving $J = J' = 7$	156
6.5	Slices of four successive axisymmetric flaglet wavelet kernels, probing different scales	
	in angular and radial space.	157
6.6	Numerical accuracy and computation time of the Fourier-Laguerre transform com-	
	puted with FLAG.	162
6.7	Numerical accuracy and computation time of the flaglet transform computed with	
	FLAGLET	163
6.8	Flaglet decomposition of the N-body simulation dataset considered for the first	
	denoising example	166

6.9	Denoising of an N-body simulation.	167
6.10	Denoising of a seismological Earth model	169
7.1	Constraints on the $\Lambda {\rm CDM} + {\rm N}_{\rm eff} + m_{\nu,{\rm sterile}}^{\rm eff}$ model, showing that non-zero sterile	
	neutrino mass is only favoured as a result of a tension between the CMB and cluster	
	data (PlaSZ, Xray) in the $\sigma_8\text{-}\Omega_m$ plane, and the degeneracy between σ_8 and neutrino	
	mass	173
7.2	Persistence of the tension as the minimal Λ CDM model is extended in the neutrino	
	sector, <i>i.e.</i> , as N_{eff} and massive active or sterile neutrinos are added	178
7.3	Constraints on the Λ CDM+ $r_{0.002}$ + N _{eff} + $m_{\nu, \rm sterile}^{\rm eff}$ model, illustrating the persisting	
	tension between X-ray clusters and CMB+BAO in the $\sigma_8\text{-}\Omega_m$ plane, despite an	
	apparent reconciliation of the BICEP and <i>Planck</i> results on $r_{0.002}$.	179

List of Tables

2.1	Number of objects in the four RQCat redshift subsamples, and for which good spectra	
	(and thus good spectroscopic redshifts) were found in the SDSS-DR7 quasar catalogue.	62
2.2	Systematics thresholds used to restrict the power spectrum analysis of RQCat to the	
	most reliable regions of the sky to minimise contamination from calibration errors.	64
2.3	The chi square values for the auto- and cross-power spectra of the four RQCat samples	
	presented in Figs. 2.14 and 2.15, with and without mode projection (mp) \ldots	68
2.4	The chi square values for the cross-power spectra of the four RQCat samples with the	
	systematics templates, presented in Fig. 2.16, using the same conventions as Table 2.3.	69
3.1	The chi-square values for the power spectra measured using extended mode projec- tion with Mask 2, presented in Fig. 3.8	102
4.1	PNG constraints for different models. The base model is $(b_0, \beta) + \cos(\theta + n(z))$. The	
	first super/subscript corresponds to the 68% CL, and the 95% CL can be obtained by	
	adding the second number.	114
5.1	Order of magnitude of the accuracy of the <code>HEALPix</code> spherical harmonic transform	131
7.1	Evidence ratios between the minimal Λ CDM model and the extended neutrino models.	175
7.2	Marginalised 95% CL constraints on the $\Lambda {\rm CDM} + {\rm N}_{\rm eff} + \sum m_{\nu}$ model from a variety	
	of robust LSS datasets with the <i>Planck</i> CMB temperature and lensing measurements.	177
7.3	Same as Table 7.2, but for the $\Lambda \text{CDM} + N_{\text{eff}} + m_{\nu, \text{sterile}}^{\text{eff}}$ model, showing tight con-	
	straints on the mass of sterile neutrinos.	177

7.4	Evidence ratios $\ln[E_{\Lambda CDM}/E_{ext.}]$ between the minimal ΛCDM model and the $\Lambda CDM + r_{0.}$	002+
	${\rm N}_{\rm eff}+m_{\nu,{\rm sterile}}^{\rm eff}$ model, showing that sterile neutrinos are not favoured by the data,	
	even when adding the BICEP results.	179

Introduction

"The purpose of models is not to fit the data but to sharpen the questions.". S. Karlin

1.1. The Universe

Cosmology is the quantitative study of the Universe as a whole: its content, origin, evolution, and fate. Its recent success as a precision science mostly rests on two pillars: large data sets mapping the Universe on a wide range of scales and frequencies, and a remarkable model describing these observations at unprecedented fidelity. Yet, many questions remain unsolved, and cosmology is an active field of research, requiring a vital balance of theory and observation, and overlapping with numerous other disciplines including astronomy, chemistry, particle physics, statistics, signal processing, and machine learning.

1.1.1. Composition and structure

The fundamental assumption of the current concordance model— Λ CDM, summarised in the following sections—is the homogeneity and isotropy of the Universe on large scales. Supported by a range of observations, this idea has been promoted to a guiding axiom: the Cosmological Principle, stating the Universe looks the same for all observers, whoever and wherever they are. A natural consequence of this axiom is the assumption that the *observable Universe* is a fair sample, *i.e.*, representative of the whole Universe, and that the same physical laws apply throughout. Violations of these assumptions are under investigation, but they are not considered in this thesis.

One of the pillars of the ACDM model is the theory of General Relativity (GR), which describes the interplay between mass, gravity, and space time. A good review GR from a cosmological perspective is provided in Dodelson (2003), and we use the same conventions here. Following the principles of GR, the Universe can be viewed as an evolving object, the structure and dynamics of which are influenced by its content, and vice versa. More precisely, in the current concordance model, the Universe is mainly composed of four ingredients:

- Baryonic matter: "visible" matter interacting through gravitational, electromagnetic, strong and weak forces, composing stars, planets, and living organisms. Making up only 4.9% of the total energy budget today, it mostly contains hydrogen and light elements formed in the early Universe.
- Relativistic species: electromagnetic radiation and neutrinos, mainly produced in the early Universe in the form of the cosmic microwave and neutrino backgrounds (CMB and CνB, respectively). They are currently subdominant in energy density compared to the other components, but played an essential role in earlier epochs (see details below).
- Cold dark matter (CDM): non-baryonic, pressureless, non-relativistic matter, that interacts very weakly (or not at all) through the electromagnetic force. Although it hasn't been directly observed or produced, its existence is well motivated by a range of observations including the CMB, galaxy rotation curves, galaxy clustering (Davis et al., 1982; Oort, 1932; Rubin and Ford, 1970; Zwicky, 1933) and gravitational lensing of galaxies and clusters (Clowe et al., 2004, 2006; Dietrich et al., 2012; Massey et al., 2007; Taylor et al., 1998). It makes up 26.8% of the energy budget of the Universe, and played a crucial role in the formation of structures in the Universe, such as galaxies and clusters of galaxies. CDM remains by far the simplest and most successful joint explanation of a plethora of phenomena. The detection of its associated particles as well as understanding baryonic physics in the CDM context (e.g., in galaxy formation) are active fields of research (see *e.g.*, Aalseth et al. 2011a,b; Angle et al. 2008; Bernabei et al. 2008).
- Dark energy (Λ): it dominates the energy budget of the Universe today (68.3%) and drives a phase of accelerated expansion (Perlmutter et al., 1999; Riess et al., 1998, 2004). It is currently indistinguishable from a cosmological constant, which poses a deep theoretical challenge. In particular, when considered as vacuum energy, the theoretical prediction under estimates the observed value by 120 orders of magnitude. This discrepancy is likely due to our

inability to unify quantum physics and GR and correctly compute all the contributions to the energy of the vacuum. Modern reviews of these theoretical challenges can be found in Carroll (2001), Clifton et al. (2012), and Martin (2012). Dark energy is a pressing question, and the search for deviations from a cosmological constant are high priorities of modern observational cosmology.

Cosmological observables such as the CMB and galaxy surveys, detailed in subsequent sections, inform us about the proportions of these four constituents, which have dominated the energy budget in various epochs, and influenced the formation of structure and the expansion history of the Universe.

The idea of an expanding Universe was introduced to resolve the apparent violation of homogeneity in Hubble's observations of the recession of galaxies. The expansion is encoded in the scale factor *a*, which is isotropic and a function of physical time *t* only, with a = 1 today. The space-time metric $g_{\mu\nu}$ of the expanding, homogenous, isotropic Universe takes a simple form,

$$g_{\mu\nu}dx^{\mu}dx^{\nu} = -c^2dt^2 + a^2(t)\left(\frac{dr^2}{1-kr^2} + r^2(d\theta^2 + \sin^2 d\phi^2)\right).$$
 (1.1)

In this metric, often called the Friedmann-Lemaitre-Robertson-Walker (FLRW) metric, (r, θ, ϕ) are the 3D polar coordinates labelling the points of the 3-dimensional constant-time slice, thereby defining comoving observers. k is a measure of the spatial curvature, reduced to canonical values k = +1, 0, -1 corresponding to a constant-positive-curvature, a flat, and a negative-curvature 3D spaces, respectively. Another important quantity is the conformal time η ,

$$\eta \equiv \int \frac{dt}{a(t)},\tag{1.2}$$

which is used to defined the comoving particle horizon in the time interval $t - t_i$, denoted by $\chi(t) = \eta - \eta_i$. The latter is the maximum comoving distance travelled by a photon traveled between t_i and t, and therefore defines a causal horizon beyond which particles have not been causally connected since t_i . In addition, solving for the geodesics equations $g_{\mu\nu}dx^{\mu}dx^{\nu}$ yields that the physical momentum and energy of photons scale as a^{-1} as the Universe expands, the result leading to the definition of redshift z,

$$\frac{\lambda_{\rm obs}}{\lambda_{\rm em}} \equiv 1 + z = \frac{a(t_{\rm obs})}{a(t_{\rm em})},\tag{1.3}$$

connecting the wavelengths and scale factors at emission and observation times (t_{em} and t_{obs} , respectively). For observations made today, we have $t_{obs} = today$ and $a(t_{obs}) = 1$ by convention.

For small redshifts, expansion around t_{obs} yields the famous Hubble law $cz \approx H_0 d$ with the Hubble constant defined as

$$H_0 = \left. \frac{1}{a} \frac{da}{dt} \right|_{t_{\rm obs}},\tag{1.4}$$

sometimes expressed in terms of the dimensionless parameter h as $H_0 = 100 \ h \text{kms}^{-1} \text{Mpc}^{-1}$. While early measurements of cepheids yielded $h \sim 0.7$ (Sandage, 1958), current CMB, supernovae and gravitational lensing measurements are able to pin down H_0 to percent level. However, some of the most recent measurements are discrepant (Freedman et al., 2001; Hinshaw et al., 2013; Planck Collaboration, 2013k; Riess et al., 2011a), and it is unclear whether this points to new physics or to residual systematic uncertainties in the measurements. Extra relativistic species in the early universe, such as neutrinos, could alleviate these tensions if they are confirmed. This issue is briefly considered below and in the last chapter of this thesis.

Assuming GR as a theory of gravity, the dynamics of the Universe are beautifully encapsulated in terms of the Friedmann and Lemaître solution of the Einstein equations,

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G\rho}{3} - \frac{k}{a^2} + \frac{\Lambda}{3}$$
(1.5)

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(3P+\rho) + \frac{\Lambda}{3}.$$
 (1.6)

These equations—often called the Friedmann equations—connect the content and the (background) expansion of the Universe, entirely captured in the time-dependent scale factor a(t). The components are encoded in the equation of state $P = P(\rho)$, which can be written as $P = w\rho$ for perfect fluids. Relativistic particles have w = 1/3, matter w = 0, and cosmological constant w = -1, although the search for deviations from w = -1 is an active research topic.

Using the equation of state—a byproduct of the Einstein and Friedmann equations—,

$$\dot{\rho} + 3\frac{\dot{a}}{a}(\rho + P) = 0,$$
 (1.7)

one finds the solution $\rho(a) = \rho_0 a^{-3(1+w)}$. This result, in conjunction with the Friedmann equations, gives

$$a(t) = \begin{cases} t^{2/3(1+w)}, & w \neq -1 \\ e^{\bar{H}t}, & w = -1. \end{cases}$$
(1.8)

for a flat universe. In other words, a *radiation-dominated* Universe has $a \propto \sqrt{t}$, whereas *matter-domination* yields $a \propto t^{2/3}$. A cosmological constant drives an exponential solution (with a constant

Hubble parameter $H = \overline{H}$).

Finally, the Friedmann equations can be rewritten in terms of the critical density $\rho_{cr} = 3H^2/8\pi G$, to pull out an explicit curvature term

$$\Omega(a) - 1 = \frac{k}{H^2 a^2},$$
(1.9)

where $\Omega(a)$ is the total energy density in units, $\Omega(a) = \rho/\rho_{cr}$. Observationally, the Universe appears to be spatially flat (*i.e.*, $\Omega = 1$), but this raises an important fine-tuning problem — the *flatness problem*—, for the initial conditions of the cosmological model, as detailed below.

1.1.2. Brief history

The observable Universe was smaller, denser and hotter in the past. In the standard cosmological paradigm, described in the next section, an initial phase of inflation is assumed to have caused a 60 *e*-fold ($=e^{60} \approx 10^{26}$) accelerated expansion and imprinted a distinctive, nearly scale-invariant spectrum of perturbations on a smooth background. Inflation is discussed in more details below. After this initial phase, the Universe was filled with an ionised baryon-electron-photon plasma. The mean free path of photons—subject to Thomson scattering—was extremely short due to the high number density of free electrons. During this radiation-dominated epoch, the Universe expanded as $a \propto \sqrt{t}$ and cooled, steadily decreasing the mean photon energy, and eventually becoming dominated by matter rather than radiation. Recombination occurred at a temperature of about 3000 K, when the Universe was about 380,000 years old. Electrons and protons of the primordial plasma combined into neutral atoms, while photons began to free-stream, their mean free path becoming larger than the Hubble scale. This is when the CMB was emitted. The Universe then remained matter dominated for a long period, expanding as $a \propto t^{2/3}$, and allowing matter perturbations to grow as the scale factor. Cosmic structures started to form: stars, quasars, galaxies, clusters of galaxies. The period between recombination and the formation of light sources is known as the dark ages. Reionisation occurred when the first objects were energetic enough to ionise neutral hydrogen, slowly filling the Universe with an ionised plasma. This process is believed to have occurred between 150 million and one billion years after inflation. These estimates will soon be refined by precise observations of the polarisation of the CMB, and by 21-cm experiments, which will map the neutral hydrogen in the epoch of reionisation. At the age of 10 billion years, the



Figure 1.1: Our picture of cosmic history, as suggested by cosmological observations and explained by the concordance model.

Universe started to undergo a new period of accelerated expansion, presumably driven by dark energy, which stopped the growth of matter perturbations on all scales. Figure 1.1 gives a cartoon overview of the various stages of the history of the Universe, starting from a phase of inflation to the recent dark energy domination epoch. If the accelerated expansion persists, it will prevent any more structures from entering the cosmological horizon, and new gravitationally bound structures from forming.

1.2. The early Universe

1.2.1. The cosmic microwave background

As previously discussed, recombination of electrons and protons occurred when the Universe cooled below \sim 3000 K, leading to the emission of relic radiation known as the cosmic microwave background (CMB). The CMB is the cornerstone of modern cosmology and one of the pillars of the Λ CDM model. First detected by Penzias and Wilson (1965), then identified by Dicke et al. (1965), its spectrum is the most perfect black body spectrum ever measured. It was first measured by COBE (Smoot et al., 1992), which also detected its temperature fluctuations—of the order of 1 in

100,000— which have since then been studied by a great number of experiments. In particular, the WMAP satellite mission (Spergel et al., 2003) revealed the extraordinary informative power of these fluctuations, and firmly established cosmology as a precision science. More recently, the *Planck* satellite (Planck Collaboration, 2013g) refined our understanding of these fluctuations. *Planck* also bridged the gap between WMAP, mapping ≥ 0.2 degree scales, and experiments like the South Pole Telescope (Keisler et al., 2011) and the Atacama Cosmology Telescope (Dunkley et al., 2011), which focused on very small ($\leq 0.1 \text{ deg}^2$) angular scales. The angular power spectra of the temperature and polarisation fluctuations from the first *Planck* release are shown in Figure 1.2. Their amplitude and features are intimately related to the composition and physical laws of the early Universe. For instance, the peaks of the power spectra constrain the curvature, the amount of dark matter and dark energy, whereas the damping tail informs us about the number of relativistic species at the time of recombination.



Figure 1.2: The angular power spectra of the CMB temperature and polarization fluctuations measured by *Planck* (Planck Collaboration, 2013i).

Not only is the physics of the CMB itself well understood and tested, but CMB experiments also provide a number of secondary observables. In particular, the lensing of the CMB by the intervening matter can now be mapped at great accuracy (Das et al., 2011; Planck Collaboration, 2013d; van Engelen et al., 2012), and is becoming established as a powerful cosmological probe.

Finally, the CMB temperature maps can be used to search for physics that is not captured by angular power spectra. Indeed, the latter only contain all the information if the CMB is Gaussian. This has been tested to excellent precision, but current measurements leaves room for small deviations, which will be discussed in the next section. Beyond (non-)Gaussianity of the CMB, mapbased techniques are used to search for specific physics or features in the CMB, such as signatures of collisions with bubble Universes, cosmic strings, and deviations from homogeneity or isotropy. These are beyond the scope of this thesis, but some of the most recent investigations can be found in Feeney et al. (2013); Planck Collaboration (2013a,b,c) and references therein.

1.2.2. Inflation and primordial non-Gaussianity

The CMB offers a unique picture of the early Universe, and Figure 1.2 shows that it is remarkably well described by the Λ CDM model. However, a number of phenomena are unexplained without appropriate initial conditions, set by an initial phase of *inflation*. A pedagogic review of these problems and how to resolve them is presented in Baumann (2009). Here I only summarise how the inflationary paradigm addresses the *horizon* and *flatness* problems.

The *horizon problem* results from the observation that the CMB is uniform to 10^{-5} , and that this uniformity cannot be explained by causality alone. Indeed, in Λ CDM Fourier modes can only evolve from outside to inside the horizon, since the radiation- and then matter-dominated Universe has been in decelerating expansion from recombination until recent times, before dark energy started to dominate. In fact, the particle horizon in the Λ CDM model is approximately 1 degree on across the CMB sky, which is at odds with the measured full-sky uniformity.

The *flatness problem* results from analysing Equation 1.9 and noticing that spatial flatness is an unstable point; in other words any deviation from flatness grows with time. Since CMB observations indicate that the Universe is spatially flat to percent level (*i.e.*, the total density today is close to the critical density), it must have been even flatter in the past. One can extrapolate these constraints to find *e.g.*, $|\Omega - 1| < O(10^{-16})$ at the time of nucleosynthesis, yielding a striking fine-tuning problem.

Inflation is an elegant solution to these problems, and is currently the leading theory for the initial conditions of the Universe. It consists of a period of quasi-de Sitter expansion with $H^2 \gg |\dot{H}|$, where comoving separations grow as $a(t) = e^{Ht}$ and can cross the horizon. But one of the features of inflation is that it must end, *i.e.*, it must only last for a finite period, and then return to a standard picture of matter and radiation domination. Therefore, a cosmological constant alone is not a viable solution since it would give constant H. This corresponds to the "old inflation" scenario (Guth, 1981; Guth and Weinberg, 1981): a scalar field trapped in a false vacuum. Therefore, a viable model of inflation requires a physical *clock* (Linde, 1982). One simple implementation of

this idea is a rolling scalar field with Lagrangian $\mathcal{L} = -\frac{1}{2}\partial_{\mu}\phi\partial^{\mu}\phi - V(\phi)$. The potential energy produces an exponential expansion rate, solving both the horizon and flatness problems provided the scale factor increases by at least a factor of e^{60} . The condition $H^2 \gg |\dot{H}|$ therefore translates into $V(\phi) \gg \dot{\phi}^2$, and inflation ends when reaching the bottom of the potential, to undergo oscillations and reheat the Universe. This process is illustrated in Figure 1.3. This *reheating* converts the potential energy of the inflaton into particles of the Standard Model. (see *e.g.*, Allahverdi et al. 2010 for a review). Quantum density and metric fluctuations are generated during inflation, yielding powerful predictions, such as the near scale-invariance of its scalar perturbations and the existence of tensor perturbations producing gravitational waves. The latter leave a distinctive large-scale angular power in the B-mode polarisation of the CMB, the detection of which is one of the main goals of ongoing CMB experiments.

Figure 1.3: Example of an inflaton potential, from Baumann (2009). CMB fluctuations are created by quantum fluctuations $\delta\phi$, and acceleration occurs when the potential energy of the field dominates over its kinetic energy. Inflation ends at ϕ_{end} when the kinetic energy has grown to become comparable to the potential energy. At reheating, the energy density of the inflaton is converted into radiation.

Inflation is a successful phenomenological model, but its exact physical implementation is currently poorly understood. Dozens of classes of models exist, with hundreds of different variants, all solving the previous horizon and flatness problems (see Planck Collaboration et al. 2013 for a set of models tested with *Planck* data). In addition, there is currently a growing ambition to implement inflation in the context of a more fundamental theory, such as string theory. Therefore, pinning down the detailed physics of inflation and reducing the space of models with observational constraints is a pressing question. One of the most important observables that can help discriminate between classes of inflationary models is primordial non-Gaussianity (PNG), denoting deviation from Gaussianity in the primordial fluctuations. In the CMB, since angular power spectra only measure the Gaussian information, PNG is measured by higher-order correlations, such as bispectra and trispectra, *i.e.*, correlations between triplets and quadruplets of modes on the sky, as illustrated in Figure 1.4.



Figure 1.4: Power spectrum and bispectrum in Fourier space. Conditions $P(|k_1-k_2|)$ and $B(k_1, k_2, k_3) = B(\frac{k_3}{k_1}, \frac{k_2}{k_1})$ ensure that they are equivalent to the 2- and 3-point correlations in real space. The power spectrum is a one-dimensional curve; the bispectrum is a two-dimensional surface which is a function of the triangle formed by (k_1, k_2, k_3) . In the context of inflation, limiting configurations (squeezed, equilateral, folded) correspond to different physical mechanisms.

Specific physical mechanisms during inflation will create correlations in different configurations and shapes. For this reason, PNG is a powerful window on inflation. In particular, canonical singlefield slow-roll inflation predicts negligible level of squeezed-type PNG (see *e.g.*, Allen et al. 1987; Bartolo et al. 2004; Maldacena 2003; Salopek and Bond 1990). Any measurement of deviations from this prediction can thus provide evidence for non-standard inflationary physics.

This specific form of PNG arising from squeezed configurations is the so-called *local* model, where the primordial potential ϕ is modified by including higher order terms,

$$\Phi = \phi + f_{\rm NL}[\phi^2 - \langle \phi^2 \rangle] + g_{\rm NL}[\phi^3 - 3\phi \langle \phi^2 \rangle], \qquad (1.10)$$

where all fields are evaluated at the same spatial coordinate, and $f_{\rm NL}$ and $g_{\rm NL}$ are real-valued

constants (often called the skewness and kurtosis parameters).

The most stringent constraints on local PNG currently come from higher-order statistics of the CMB. Most recently, the *Planck* Collaboration reported $-8.9 < f_{\rm NL} < 14.3$ (95% CL) (Planck Collaboration, 2013h), while constraints on the kurtosis have been obtained from the WMAP satellite, also consistent with $g_{\rm NL} = 0$. While these results are compatible with Gaussian initial conditions, their uncertainties still leave room for non-standard inflation models. Importantly, PNG can be detected in surveys of the large-scale distribution of galaxies, as detailed below. Performing this measurement in real data is one of the objectives of this thesis.

1.3. The evolved Universe

The CMB is a powerful window on the early Universe, but is not sensitive to all predictions of the Λ CDM model. For example, the effects the dark energy are only detectable on the largest angular scales through the Integrated Sachs-Wolfe (ISW) effect, the blueshifting of CMB photons due to evolving potential wells at late times. Relativistic species like neutrinos are hard to distinguish from radiation, and theories of gravity can only be tested indirectly through their geometrical effects. Finally, the CMB is a surface by nature, and CMB temperature experiments like *Planck* are already on the verge of being limited by the number of modes ($\sim 10^6$ modes). Even though CMB polarisation and lensing have not yet reached their full maturation, they will also be ultimately limited by the statistical power of the last scattering surface.

On the contrary, the large-scale structure of the Universe is a three-dimensional structure giving access to a much larger number of modes, and a plethora of additional observables. Ongoing and future experiments like the Dark Energy Survey¹ (DES), Euclid², the Square Kilometer Array³ (SKA), the Large Synoptic Survey Telescope⁴ (LSST), will probe $> 10^9$ modes and an immense range of physical processes, scales and epochs, which will allow us to test detailed scale- and redshift-dependent models of dark matter, dark energy, gravity, and baryonic physics.

¹http://www.darkenergysurvey.org/

²http://sci.esa.int/euclid/

³https://www.skatelescope.org/

⁴http://www.lsst.org/lsst/

1.3.1. Matter perturbations and structure formation

Figure 1.5 shows the power spectrum of the matter perturbations today (z = 0), *i.e.*, fluctuations in the dark matter and baryon fluid, superimposed with a variety of observables. This again highlights the remarkable success of the Λ CDM model to describe cosmological observations in a range of scales and frequencies. The form of the matter power spectrum can be derived with linear perturbation theory and evolving the primordial scale-invariance power spectrum $P(k) \propto k$, as summarised below.

Figure 1.5: Matter power spectrum from a range of probes, reproduced from Hlozek et al. (2012), where the data sets are described in detail.

The turnover scale around $k \sim 0.02 \ h {\rm Mpc}^{-1}$ is due to the very different physics that affects the modes entering the horizon before or after matter-radiation equality. The largest scale modes are always bigger than the horizon ("super-horizon") and don't evolve through causal processes. During radiation- and matter-dominated epochs, they grow linearly with conformal time, and become constant once dark energy starts to dominate. At intermediate scales, modes enter the horizon after matter-radiation equality and evolve in a dark matter-dominated Universe. Since cold dark matter is pressureless, in this case the perturbations also grow linearly with conformal time. The smallest-scale modes, however, enter the horizon before matter-radiation equality and evolve in a radiation-dominated Universe. Therefore, the power at these scales is suppressed by a factor of

 $\sim k^{-4}$ due to oscillations in the photon plasma, yielding the $\sim k^{-3}$ scaling of the matter power spectrum at $k > 0.1 h \text{Mpc}^{-1}$.

The matter power spectrum is a rich source of cosmological information, since it is sensitive to the composition, initial conditions, and evolution of the Universe.

1.3.2. Galaxy surveys

Unlike the CMB, the distribution of matter is not directly observable. It is sampled by tracers of the *light* distribution, such as galaxies and quasars, which probe the large-scale structure through their positions, and the magnification and lensing of their light by gravity. This thesis focuses on the study of the clustering of galaxies and quasars found by wide area surveys. Obviously, other tracers can be used, such as the distribution of neutral hydrogen, measured in the Ly α forest of quasar spectra or with future 21cm radio surveys. Their study can bring a great deal of information about baryonic physics and the z > 1 Universe.

There are two distinct types of surveys of the large-scale structure (with a spectrum of intermediate cases, not considered here): spectroscopic and photometric surveys. In the former, one uses a high-resolution spectrograph to stare at bright sources and acquire their spectral energy distribution to estimate their type, characteristics and redshift. Hence, such surveys provide the 3D positions of galaxies and quasars, and can be used to map the large-scale structure in great detail. However, they are limited by the sensitivity of the spectrographs, and by the time and financial constraints needed to acquire good spectra of large numbers of sources. Photometric surveys, on the contrary, make use of broad band filters and wide-field cameras to overcome these constraints. They can detect faint sources (typically hundreds of millions of galaxies) over large areas of the sky (from hundreds to tens of thousands of square degrees), giving access to larger volumes and higher statistical power. Obviously, other complications arise, since the light spectrum is now only observed through a handful of filters. Topical issues include detecting sources and classifying sources (e.g., distinguishing galaxies from stars), estimating their fluxes and shapes in noisy images, and estimating their photometric redshift (photo-z). Most of these issues remain open problems and intense areas of research. Interestingly, some of them benefit from machine learning techniques (e.g., Collister and Lahav 2004; Graff et al. 2014), such as neural networks, which efficiently find classification and estimation rules in existing data and prove useful when physical models are not available or perform poorly (e.g., for photo-z estimation and star-galaxy separation).

Photometric surveys also rely different types of analysis compared to spectroscopic surveys. In general, photo-z estimates contain significant uncertainties, causing the radial distribution of detected galaxies to be poorly known. Hence one cannot probe the 3D distribution of matter, for example through 3D statistics such as the Fourier power spectrum P(k). The most common alternative is to measure the *projected* clustering of galaxies, involving 2D statistics such as angular power spectra or 2D correlation functions, in redshift slices of size greater than the photometric redshift uncertainties. In the best cases (with good photo-z), using sufficiently small bins and crosscorrelations between overlapping redshift samples enable the recovery of most of the information compared to the 3D distribution. Some of the most recent measurements of the 3D and 2D clustering of galaxies are presented in Figure 1.6. A notable advantage of exploiting angular power spectra in redshift bins is to keep the angular and radial dimensions *decoupled* in the analysis. This can facilitate the treatment of redshift dependent effects (sometimes difficult to incorporate in P(k)approaches⁵), and observational complications — or systematics —, which usually affect the angular and radial information differently. The identification and robust treatment of systematics is a central point of this thesis. Indeed, future surveys will not be limited by statistical uncertainties, but rather by our ability to deal with systematics.

The balance between depth, volume, photo-*z* uncertainties, and systematics can advantage spectroscopic and photometric surveys in different applications. The former are currently used for accurate measurements of the 3D power spectrum and bispectrum, mainly aiming at sub-percent level measurements of baryon acoustic oscillation feature. But the largest ongoing and future surveys —the Kilo Degree Survey⁶ (KiDS), DES, Euclid, LSST— are all photometric. Beyond the gain in volume and statistical power, this is also because such surveys are optimal for measuring and combining information from the clustering and gravitational lensing of galaxies. These quantities do not only probe different scales and types of physics, but also suffer from different systematics. In particular, cosmic shear — the magnification of light and distortion of shapes of galaxies due to gravity— is a powerful window on the matter distribution, but also involves complications such as intrinsic alignment and shape measurement noise. The agreement but also complementarity between various probes of the large-scale structure is illustrated in Figure 1.5. The Λ CDM model remarkably fits all datasets, but future experiments will decrease the errors on these measurements by orders of magnitude. Combining data sets and testing models will then be challenging.

⁵Most spectroscopic P(k) analyses assume no redshift evolution within the redshift range of the sample. It can be shown that sample spanning a large redshift range requires the a full 3D approach, for example involving the Fourier-Bessel transform, to exploit all the information.

⁶http://kids.strw.leidenuniv.nl/

Figure 1.6: Left: measurements of the 3D power spectrum of galaxies from the Baryon Acoustic Oscillation Survey (Anderson et al., 2012). Right: tomographic angular power spectrum of luminous red galaxies from SDSS DR7 (Thomas et al., 2010).

1.3.3. Scale-dependent halo and galaxy bias

Connecting the distribution of light tracers to the underlying matter is a difficult task, due our limited knowledge of the physics controlling the formation and evolution of galaxies in dark matter halos. These complications are usually encapsulated in a *galaxy (or quasar) bias* term, which is added on top of the *halo bias*, to relate the galaxy and the dark matter overdensity fields. In its simplest incarnations —well verified observationally for certain types of galaxies— the bias is linear and scale-independent on large-scales, and becomes non-linear, and scale-dependent on small scales. The development of bias models is an active field of research, involving theoretical investigations and extensive comparisons to cosmological simulations. A major challenge of ongoing and future surveys will be to connect phenomenological bias models (fitted to measurements of the clustering of galaxies as a function of scale, luminosity, galaxy type) to the underlying physics of dark matter halo and galaxy formation.

Interestingly, the halo bias itself can be affected by primordial physics. The example considered in this thesis is primordial non-Gaussianity, which couples large and small scale modes in the initial conditions. This coupling can propagate and modify the abundance of halos in overdense regions (by tilting the halo mass function), therefore also affecting the number of galaxies in a scale-dependent manner. The effect can be visualised in Figure 1.7: non-zero local PNG will



Figure 1.7: The Kaiser effect (Kaiser 1984) describes the fact that galaxies form at the peaks of the primordial density fluctuations. Squeezed non-Gaussianity increases the large scale modulation of the density peaks and thus impacts the clustering of tracers. Compared to the Gaussian case, not only are tracers more massive but they are more correlated over large distances, which modifies the bias of the power spectrum at large-scales.

increase the abundance of massive halos on either large or small scales (for positive and negative $f_{\rm NL}$, respectively). Therefore, the net signature of primordial non-Gaussianity on the large-scale structure is an enhancement of the bias of tracers on large-scales. Quantitatively, non-zero $f_{\rm NL}$ and $g_{\rm NL}$ will modify the "Gaussian" bias $b^{\rm G}(z)$ into (Dalal et al., 2008; Desjacques and Seljak, 2010b; Matarrese and Verde, 2008; Slosar et al., 2008; Smith et al., 2012)

$$b^{\rm NG}(k,z) = b^{\rm G}(z) + \frac{\beta_f(z)f_{\rm NL} + \beta_g(z)g_{\rm NL}}{\alpha(k,z)}$$
 (1.11)

$$\beta_f = 2\delta_c \left(b^{\mathcal{G}}(z) - 1 \right) \tag{1.12}$$

$$\beta_g = 3 \frac{\partial \log n}{\partial f_{\rm NL}} \tag{1.13}$$

$$\alpha(k,z) = \frac{2k^2 T(k) D(z)}{3\Omega_m H_0^2}$$
(1.14)

where δ_c is the spherical collapse threshold, *n* the halo number density, and $\alpha(k, z)$ relates the linear density field and the primordial potential through $\delta_{\text{lin}}(k, z) = \alpha(k, z)\Phi(k)$. These expressions neglect an additional small contribution induced by the effect of PNG on the halo mass function (which can absorbed in b^G). For simplicity, the implicit mass dependence of b^G , β_f , and β_g is also neglected, and the bias b^G is considered to be only redshift dependent. Note that for β_g we use the fitting formula from Smith et al. (2012), which is valid for the range of masses and scales considered in this thesis.

Primordial non-Gaussianity is powerful window on the early Universe, and it is remarkable that

its effect can be measured in the bias of large-scale structure tracers. However, measurements of the bias on large scales are largely limited by cosmic variance and the presence of spurious large-scale correlations, often produced by contamination from observational systematics. This thesis deals with this specific problem, *i.e.*, measuring the bias and constraining primordial non-Gaussianity in the presence of these two complications.

1.3.4. Quasars

Quasars are the compact, energetic regions surrounding the central supermassive black hole of massive galaxies. The violent accretion of material by black holes makes quasars the most luminous objects in the known Universe. They have a very distinctive spectral energy distribution, show in Figure 1.8, and radiate in a range of frequencies, from X-rays to the far-infrared, with a peak in the ultraviolet-optical bands. Some quasars are also strong sources of radio emission and of gamma-rays.

Among all tracers of the large-scale structure, quasars are the most biased, and their bias strongly increases with redshift ($b \sim 2, 3, 5$ at z = 1.5, 2, 2.5 respectively), mostly because they form and reside in average mass halos across cosmic times (Fanidakis et al., 2013; Tinker et al., 2010).



Figure 1.8: Typical spectral energy distribution of a quasar, from Vanden Berk et al. (2001).

Thanks to the distinctive quasar spectrum, spectroscopic surveys of quasars don't suffer from

stellar contamination. However, they are based on small samples of bright sources pre-selected in specific redshift ranges. This is due to the limited sensitivity of multi-fiber spectrographs and the resources needed to verify large numbers of sources.

This thesis focuses on surveys of *photometric* quasars, *i.e.*, quasar candidates identified with an imaging survey, here the Sloan Digital Sky Survey⁷ (SDSS). The main advantage of this type of survey is to detect a large number of quasars up to redshift $z \sim 4$, and probe much larger volumes than spectroscopic surveys, therefore significantly decreasing the variance of large-scale clustering. This makes quasars ideal populations for PNG studies. However, a significant fraction of these photometric candidates are likely to be stars, mostly because it is difficult to distinguish quasars from stars and estimate their redshift using only photometric information. The performances of the star-quasar classification are also heavily redshift-dependent as the emission lines fall in or out of the photometric filters as they get redshifted. Observational systematics couple with these effects to create a complex spatially varying stellar contamination and depth in photometric quasar samples. These effects lead to spurious large-scale correlations and prevent from harnessing the statistical power of photometric quasars to constrain PNG. The first part of this thesis (chapters 1-3) largely deals with the resolution of this problem.

1.3.5. Cosmological neutrinos

Neutrinos are electrically neutral fermions playing an important role in sub-atomic interactions. Solar and atmospheric experiments have revealed that they are massive, in contrast with the predictions from the standard model of particle physics. Yet, their precise masses have not been measured. This is an active research area in cosmology because the number and masses of neutrinos affect both the CMB and matter anisotropies. In particular, ongoing and future CMB and galaxy surveys should provide precise measurements of these quantities, and guide the extensions of the standard models of cosmology and particle physics.

In the early universe, neutrinos are relativistic and behave like radiation, therefore affecting the acoustic peaks and damping tail of the CMB power spectrum. The latter is not very sensitive to the mass of neutrinos, but rather to their number, usually parametrised by N_{eff} , the number of extra relativistic species (or *dark radiation*). The fiducial value $N_{\text{eff}} = 3.046$ corresponds to three *active* neutrinos (of left-handed chirality), the difference from three being due to the small

⁷http://www.sdss.org/



Figure 1.9: Effect of massive neutrinos on the matter and CMB power spectra, reproduced from Park et al. (2012).

amount of entropy from electron/positron annihilation that gets transferred to the neutrino. N_{eff} is by design equal to three in the idealised case that all of this entropy is transferred to photons. Current estimates of N_{eff} have a slight $(1 - 2\sigma)$ preference for $N_{\text{eff}} > 3.046$ (Dunkley et al., 2011; Keisler et al., 2011; Planck Collaboration, 2013k), especially when combining the CMB with local measurements of H_0 (Freedman et al., 2001; Hinshaw et al., 2013; Planck Collaboration, 2013k; Riess et al., 2011a).

In the late universe, massive neutrinos free-stream and suppress the growth of structure on small scales, in proportion to their total mass, as shown in Figure 1.9. Therefore, the latter can be constrained by any observable of the matter power spectrum, including the clustering of galaxies and quasars, cosmic shear, abundances of clusters of galaxies, and redshift space distortions. Some recent measurements of the σ_8 parameter, the RMS of matter fluctuations on a 8 h^{-1} Mpc scale, are in significant tension with the CMB estimate. It was also argued that including very massive neutrinos in the Λ CDM model could resolve these discrepancies.

It is currently unclear whether the discrepancies on H_0 and σ_8 measurements are hints of new physics or due to residual systematics in the measurements. The last chapter of this thesis investigates these tensions, and whether extending Λ CDM with extra massive neutrinos (active or sterile) can resolve them. It is focused on tracking down the origin of the tensions and testing the extended models with robust combinations of data sets.

1.4. Wavelets in cosmology

The uncertainty principle of harmonic analysis states that functions cannot be perfectly localised in both real and frequency spaces simultaneously. In other words, a Dirac delta function (perfectly localised) in one space is necessarily unbounded (*i.e.*, has infinite support) in the other space. Wavelets are the compromise solution: they are *bounded* (or *band-limited*) in both spaces, as shown in Figure 1.10, and can be constructed to probe specific scales and frequency ranges. The prove particularly useful to represent or analyse data with localised features. In fact, they are a key ingredient of the notion of *sparsity*, which aims at finding and exploiting sets of functions or *dictionaries* to optimally represent data with the smallest number of components. Many natural datasets turn out to be *sparse* in wavelet-type bases, which has lead the development of groundbreaking algorithms in applications such as data compression and reconstruction. Comprehensive overviews of the fundamentals of wavelets and sparsity can be found in *e.g.*, Daubechies 1992a; Mallat 1999.



Figure 1.10: Example of "mexican hat" wavelets constructed to be localised in time and frequency domains (Couderc and Zareba, 1998). These wavelets are redundant, *i.e.*, a given scale or frequency is probed by more than one wavelet.

In cosmology, wavelets have been commonly used to search for anomalous features in the CMB (*e.g.*, Feeney et al. 2011a), as well as constraining dark energy (McEwen et al., 2007a, 2008) and primordial non-Gaussianity (Cayón et al., 2001; Vielva et al., 2004). The latter modifies the average amplitude and size of hot and cold spots, which are efficiently represented in wavelet bases. They are also used in several raw data analysis operations, for example to decompose and process CMB polarisation beams and galaxy shapes (*e.g.*, Kuijken 2006; Refregier and Bacon 2003). Most recently, the perspectives of using wavelets and sparsity for galaxy survey data analysis have

become apparent. Indeed, most features of the large-scale structure —filaments, clusters, voids can be represented with a small number of localised components. The chapters 4 and 5 of this thesis deal with developing wavelets on the 2D sphere and on the 3D ball with interesting localisation and reconstruction properties, and lay down the perspectives of applying them to CMB and galaxy survey data.

1.5. Thesis outline

Each chapter of this thesis consists of a paper written by myself and collaborators. I was critically involved in every aspect of these papers, including the development of the theory and methods, the implementation and application to data, and the writing of the manuscript, where I was the main author. The chapters of this thesis are reproductions of the following papers as they appear in print. Minor formatting changes have been made to made to match thesis requirements. The titles, co-authors and publication details of the papers comprising each chapter are listed below. All were carried out in collaboration with the named co-authors.

Chapter 1: Estimating the large-scale angular power spectrum in the presence of systematics: a case study of Sloan Digital Sky Survey quasars.

This work was published as Boris Leistedt, Hiranya V. Peiris, Daniel Mortlock, Aurélien Benoit-Lévy, Andrew Pontzen, 2013, Monthly Notices of the Royal Astronomical Society, 435 (3): 1857-1873, and was carried out in collaboration with the named co-authors.

Chapter 2: Exploiting the full potential of photometric quasar surveys: Optimal power spectra through blind mitigation of systematics.

This work was published as Boris Leistedt, Hiranya V. Peiris, 2014, Monthly Notices of the Royal Astronomical Society, 444(1): 2-14, and was carried out in collaboration with the named co-authors.

Chapter 3: Constraints on primordial non-Gaussianity from 800,000 photometric quasars.This work was carried out in collaboration with Hiranya V. Peiris and Nina Roth, and was submitted to Physical Review Letters in August 2014. It is currently in review.
Chapter 4: *S2LET*: a code to perform fast wavelet analysis on the sphere.

This work was published as Boris Leistedt, Jason D. McEwen, Pierre Vandergheynst, Yves Wiaux, 2013, Astronomy & Astrophysics, 558, A128, and was carried out in collaboration with the named co-authors.

Chapter 5: Exact Wavelets on the Ball.

This work was published as Boris Leistedt, Jason D. McEwen, 2012, IEEE Transactions on Signal Processing, 60, 6257-6269, and was carried out in collaboration with the named co-authors.

Chapter 6: No new cosmological concordance with massive sterile neutrinos.

This work was published as Boris Leistedt, Hiranya V. Peiris, Licia Verde, 2014, Physical Review Letters, 113, 041301, and was carried out in collaboration with the named co-authors.

Estimating the large-scale angular power spectrum in the presence of systematics: a case study of Sloan Digital Sky Survey quasars

"The first principle is that you must not fool yourself and you are the easiest person to fool". R. Feynman

2.1. Abstract

The angular power spectrum is a powerful statistic for analysing cosmological signals imprinted in the clustering of matter. However, current galaxy and quasar surveys cover limited portions of the sky, and are contaminated by systematics that can mimic cosmological signatures and jeopardise the interpretation of the measured power spectra. We provide a framework for obtaining unbiased estimates of the angular power spectra of large-scale structure surveys at the largest scales using quadratic estimators. The method is tested by analysing the 600 CMASS mock catalogues constructed by Manera et al. (2013) for the Baryon Oscillation Spectroscopic Survey (BOSS). We then consider the Richards et al. (2009) catalogue of photometric quasars from the Sixth Data Release (DR6) of the Sloan Digital Sky Survey (SDSS), which is known to include significant stellar contamination and systematic uncertainties. Focusing on the sample of ultraviolet-excess (UVX) sources, we show that the excess clustering power present on the largest-scales can be largely mitigated by making use of improved sky masks and projecting out the modes corresponding to the principal systematics. In particular, we find that the sample of objects with photometric redshift $1.3 < \tilde{z}_p < 2.2$ exhibits no evidence of contamination when using our most conservative mask and mode projection. This indicates that any residual systematics are well within the statistical uncertainties. We conclude that, using our approach, this sample can be used for cosmological studies.

2.2. Introduction

The cosmic microwave background (CMB) and the large-scale structure (LSS) of galaxies contain a wealth of physical information that can be used to test models of the origin and evolution of the Universe. Both are well-described by correlated Gaussian random fields, and therefore can be characterised by two-point statistics (see, e.g., Bond et al. 1998, 2000; Tegmark et al. 2002a). In particular, the angular power spectrum is a natural tool for CMB data analysis, and has also proved useful for the study of the clustering properties of galaxy surveys in redshift bins. Such tomographic approaches will be essential for exploiting next generation surveys such as the Dark Energy Survey¹ (DES), which will provide large photometric catalogues where the uncertainties on the redshift estimates complicate a full three-dimensional analysis.

However, data unavoidably contain non-cosmological contributions, for example due to instrumental errors and systematic uncertainties. These contaminants result in additional correlations in the measured power spectra, and can compromise our interpretation of the observables if not correctly treated. In the context of galaxy surveys, observational systematics and calibration errors can result in extra clustering power over a wide range of scales (see e.g., Huterer et al. 2013; Ross et al. 2011, 2012a; Thomas et al. 2010, 2011). This proves especially problematic at the largest scales, since the corresponding modes need to be constrained from partial sky data. These modes are nonetheless crucial for testing early universe theories such as cosmological inflation (e.g., Albrecht and Steinhardt 1982; Guth 1981; Linde 1982), the standard paradigm for describing the origin of structure in the universe. Future galaxy surveys will be able to test this paradigm very precisely, particularly through the search for signatures of primordial non-Gaussianity (PNG). PNG creates a scale-dependent galaxy bias affecting the 2-point clustering properties of LSS tracers at the largest scales (Dalal et al., 2008; Matarrese and Verde, 2008). Hence, these scales, which can be strongly affected by systematics, require particularly careful treatment.

 $^{^1}$ www.darkenergysurvey.org

Quasars, being highly biased tracers of LSS, are excellent candidates to study the scale- and redshift dependence of galaxy bias, for example to constrain PNG. However, since current spectroscopic samples have low number densities and cannot compete with PNG constraints derived from the CMB, one must resort to photometric quasar catalogues. The largest catalogues currently available are extracted from the Sloan Digital Sky Survey (SDSS), and were used to study PNG and the integrated Sachs-Wolfe effect (ISW) (Giannantonio et al., 2006, 2008; Slosar et al., 2008; Xia et al., 2010, 2011). However, these studies demonstrated the high sensitivity of the correlation functions to the sky masks under consideration, indicating the presence of significant levels of contamination by stars and calibration-related systematics. In particular, recent work by Pullen and Hirata (2013), corroborated by Giannantonio et al. (2014), confirmed the high levels of contamination in the Richards et al. (2009) catalogue of SDSS photometric quasars, leading to concerns about the use of this sample for clustering measurements. In this work, we use sample reduction, masking and mode projection to identify a subset of objects in this catalogue that can be used for cosmological analyses. We concentrate on the main systematics found by previous studies, and analyse their impact on the clustering measurements through auto- and cross-correlation of redshift subsamples with each other, and with templates of the systematics. Note that cross-correlating with external data can also prove useful in identifying and mitigating the systematics in quasar samples.

When analysing photometric catalogues in redshift bins, the theory power spectrum predictions require precise estimates of the redshift distributions, which are compromised by the large uncertainties of the photometric redshifts. This issue is critical for photometric quasars, since their redshift estimates are significantly more uncertain than for other types of galaxies and include a significant fraction of catastrophic failures. We investigate the use of spectroscopic catalogues for calculating robust and unbiased redshift distribution estimates for the photometrically-selected quasar subsamples.

In addition to data quality and modelling issues, various methodological issues arise when estimating the power spectrum of a galaxy survey for comparison with theory. The pseudo-spectrum (Hivon et al., 2002; Wandelt et al., 2001) and quadratic maximum likelihood (Bond et al., 2000; Tegmark, 1997) estimators were developed to measure the power spectrum in the presence of sky cuts. However, numerous technical subtleties and constraints due to pixelisation or limited computer resources are implicit in these estimators, and can create significant biases if not handled carefully, as concluded by several studies on the CMB (see e.g., Copi et al. 2011; Efstathiou 2004a; Eriksen et al. 2007; Pontzen and Peiris 2010). This paper aims to clarify these technicalities in the context of galaxy surveys.

This article is organised as follows. In Sec. 2 we define and illustrate the properties of quadratic power spectrum estimators, and demonstrate their validity by applying them to a set of mock catalogues. In Sec. 3 we turn to the Richards et al. (2009) catalogue of SDSS photometric quasars. We present our data samples, redshift distribution estimates, masks, and power spectrum measurements, and discuss the impact of the main systematics on these measurements. Our conclusions are presented in Sec. 4. Further technical details on smoothing and masking rules, Karhunen-Loève compression and χ^2 measures are contained in appendices.

2.3. Theory and methods

2.3.1. Background

We consider a real signal $x(\mathbf{n})$ on the unit sphere S^2 , equivalently described in terms of its spherical harmonic coefficients $\{x_{\ell m}\}$ with $\ell \in \mathbb{N}$ and $m \in \{-\ell, \ldots, \ell\}$. The angular power spectrum $\{C_\ell\}$ of x is defined as

$$C_{\ell} = \sum_{m=-\ell}^{\ell} \frac{|x_{\ell m}|^2}{2\ell + 1},$$
(2.1)

and corresponds to the average power in fluctuations on scales of order $180/\ell$ degrees on the sphere. Assuming that x is a realisation of an underlying random field denoted by X, the power spectrum of x can be viewed as a compression technique, and used to perform statistical inference on physical models of X. In particular, this compression is lossless if X is an isotropic Gaussian random field, and the power spectrum is then a sufficient statistic containing all the relevant information in the realisation x. Moreover, the 'observed' power spectrum C_{ℓ} is a realisation of a 'theory' power spectrum C_{ℓ} that fully characterises the field of interest X. The variance of the former, known as cosmic variance, depends on the number of modes on the sky and is given by

$$\operatorname{Var}(\mathcal{C}_{\ell}) = \frac{2C_{\ell}^2}{2\ell + 1}.$$
(2.2)

In practice, real data contain a finite amount of information, and the continuous signal x is observed at finite resolution on the sphere. In the context of LSS surveys, galaxy catalogues are usually constructed from raw imaging data and then reduced into pixelised overdensity maps $\mathbf{x} = (x_0, \dots, x_{N_{pix}-1})$ where $x_i = x(\mathbf{n}_i)$ and \mathbf{n}_i is the centre of the *i*th pixel on the sphere. More details on the construction of such maps from the source number counts will be given in Sec. 2.3.4. The average correlation between the pixels, namely the pixel-pixel covariance matrix, depends on the theory power spectrum through a Legendre expansion, i.e.,

$$\mathbf{S} = \langle \mathbf{x} \mathbf{x}^t \rangle = \sum_{\ell} C_{\ell} \mathbf{P}^{\ell}, \qquad (2.3)$$

where $(\mathbf{P}^{\ell})_{ij} = (2\ell + 1)/4\pi P_{\ell}(\mathbf{n}_i \cdot \mathbf{n}_j)$ is a useful matrix notation (Tegmark, 1997). A quadratic estimator for the power spectrum of full-sky pixelised data is given by the projection of the data onto the Legendre matrices, i.e.,

$$C_{\ell} = \mathbf{x}^t \mathbf{P}^{\ell} \mathbf{x},\tag{2.4}$$

which is the pixel-space equivalent of Eq. (3.1) (see discussion and references in Pontzen and Peiris 2010).

In this section we have used an arbitrary equal-area pixelisation scheme, but henceforth we will adopt the HEALPIX conventions (Górski et al., 2005). In defining Eq. (2.4), we only considered full sky coverage. This assumption will be relaxed in the next section. We also implicitly assumed that the power spectrum of the pixelised map x was equal to that of the continuous signal x. This approximation is only true at high resolution when the pixel size is small compared with $180/\ell$, and the integrals in the spherical harmonics and Legendre transforms are correctly approximated by matrix multiplications through quadrature, as in Eq. (2.4). The bias induced by pixelisation as a function of ℓ is critical for low-resolution power spectrum estimation, and needs to be corrected. This issue is investigated in Appendix 2.7.1, and the following sections will assume that the relevant corrections have been applied.

2.3.2. Partial sky coverage and quadratic estimators

Due to contamination or inaccessibility of certain regions of the sky, most cosmological applications involve signals that only cover a portion of the sphere. The power spectrum must then be calculated from a cut-sky map $\tilde{\mathbf{x}}$. From a theoretical perspective, the latter can be viewed as the restriction of the full sky map \mathbf{x} using a binary mask $\mathbf{m} = (m_0, \dots, m_{N_{pix}-1})$, such that $m_i = 0$ for masked pixels and $m_i = 1$ elsewhere. Masked/unmasked vectors or matrices are related to each other through an operator **K**, a diagonal matrix such that $(\mathbf{K})_{ij} = m_i \delta_{ij}$ (with δ_{ij} the Kronecker delta) removing pixels that lie inside the mask². In what follows, the addition of a tilde will represent cut-sky quantities.

In the presence of partial sky coverage, applying Eq. (2.4) on the cut-sky map $\tilde{\mathbf{x}}$ leads to a cut-sky power spectrum $\{\tilde{C}_{\ell}\}$ that considers the zones inside the mask as *data*, i.e., as pixels with $x(\mathbf{n}) = 0$. Consequently \tilde{C}_{ℓ} differs from the quantity of interest C_{ℓ} and is not a realisation of the underlying theory spectrum C_{ℓ} . Inverting this effect involves deconvolving the effect of the mask from the observed power spectrum $\{\tilde{C}_{\ell}\}$, leading to the definition of the 'pseudo-spectrum' (PCL) estimator (Brown et al. 2005; Efstathiou 2004b; Hivon et al. 2002; Wandelt et al. 2001),

$$\hat{C}_{\ell}^{\text{PCL}} = \sum_{\ell'} (\mathbf{M}^{-1})_{\ell\ell'} \tilde{\mathcal{C}}_{\ell'}, \qquad (2.5)$$

where $\tilde{\mathcal{C}}_{\ell'} = \tilde{\mathbf{x}}^t \tilde{\mathbf{P}}^\ell \tilde{\mathbf{x}}$ are the cut-sky estimates. The coupling matrix is defined as

$$(\mathbf{M})_{\ell\ell'} = \operatorname{Tr} \tilde{\mathbf{P}}^{\ell} \tilde{\mathbf{P}}^{\ell'}, \qquad (2.6)$$

and is a function of the mask only. Since the variance of each ℓ -mode depends on number of times it is observed, the minimum variance, namely the cosmic variance presented in Eq. (3.2), is achieved on the full sky only. Partial sky coverage decreases the number of observed modes, and the variance of the PCL estimates in the absence of noise is approximately

$$\operatorname{Var}(\hat{C}_{\ell}^{\operatorname{PCL}}) \approx \frac{1}{f_{\operatorname{sky}}} \operatorname{Var}(\mathcal{C}_{\ell}).$$
 (2.7)

Here, $f_{sky} = \sum_i m_i / N_{pix}$ is the fraction of the sky covered by the mask, with $f_{sky} = 1$ corresponding to full sky coverage. Equation (2.7) is a good approximation for small scale modes, which remain numerous after masking. The exact expression for the variance in the Gaussian framework is given in Eq. (2.11), and must be used for low- ℓ modes since they are sensitive to the shape of the mask.

The PCL approach is simply an inversion of the mask and does not attempt to minimise the loss of information caused by the decrease in the number of observed modes. In fact, it is well known that the PCL estimates are only optimal (i.e., unbiased, minimum variance estimates) for a flat power spectrum (see e.g., Efstathiou 2004b and Pontzen and Peiris 2010). This equivalence will prove useful in the context of galaxy surveys, as we shall see in the next sections. To recall the

²Hence $\tilde{\mathbf{b}} = \mathbf{K}\mathbf{b}$ for any data vector *b*, while for any matrix **B** we write $\tilde{\mathbf{B}} = \mathbf{K}\mathbf{B}\mathbf{K}$ implicitly taking advantage of the property $\mathbf{K}^t = \mathbf{K}$.

definition of the optimal estimator, we consider the generic class of quadratic estimators of the form

$$\hat{C}_{\ell} = \tilde{\mathbf{x}}^t \mathbf{E}^{\ell} \tilde{\mathbf{x}}.$$
(2.8)

In this formalism, the PCL estimator reads

$$\mathbf{E}_{\mathrm{PCL}}^{\ell} = \sum_{\ell'} \left(\mathbf{M}^{-1} \right)_{\ell\ell'}^{\mathrm{PCL}} \tilde{\mathbf{P}}^{\ell'}, \tag{2.9}$$

with the coupling matrix $(\mathbf{M})_{\ell\ell'}^{PCL} = \text{Tr } \tilde{\mathbf{P}}^{\ell} \tilde{\mathbf{P}}^{\ell'}$. In the Gaussian case, the expected value of the generic quadratic estimator is given by

$$\langle \hat{C}_{\ell} \rangle = \text{Tr} \; \tilde{\mathbf{C}} \mathbf{E}^{\ell},$$
 (2.10)

and its variance by

$$\mathbf{V}_{\ell\ell'} = \langle \hat{C}_{\ell} \hat{C}_{\ell'} \rangle - \langle \hat{C}_{\ell} \rangle \langle \hat{C}_{\ell'} \rangle = 2 \operatorname{Tr} \tilde{\mathbf{C}} \mathbf{E}^{\ell} \tilde{\mathbf{C}} \mathbf{E}^{\ell'}.$$
(2.11)

As a result, in the presence of sky cuts the uncertainties on the power spectrum estimates are typically significantly correlated. Uncorrelated error bars can be obtained by diagonalising the covariance matrix and using the resulting rotation matrix to transform the power spectrum estimates and the theory predictions (Tegmark, 1997; Tegmark et al., 2002a).

In the previous equations, $\tilde{\mathbf{C}}$ denotes the cut-sky pixel-pixel covariance matrix, which can be modelled as the superposition of a signal part $\tilde{\mathbf{S}}$ calculated with a theory prior $\{C_{\ell}\}$ and noise, i.e.,

$$\tilde{\mathbf{C}} = \langle \tilde{\mathbf{x}} \tilde{\mathbf{x}}^t \rangle = \tilde{\mathbf{S}} + \tilde{\mathbf{N}}.$$
 (2.12)

The pixel-pixel covariance matrix must also incorporate any additional signal present in the data, such as the systematics, as detailed in Sec. 2.3.4.

The minimum variance estimator in the Gaussian framework, first introduced in Tegmark (1997), is the so-called quadratic maximum likelihood (QML) estimator. The latter reads

$$\mathbf{E}_{\text{QML}}^{\ell} = \sum_{\ell'} \left(\mathbf{M}^{-1} \right)_{\ell\ell'}^{\text{QML}} \frac{1}{2} \tilde{\mathbf{C}}^{-1} \tilde{\mathbf{P}}^{\ell'} \tilde{\mathbf{C}}^{-1},$$
(2.13)

and uses the deconvolution matrix

$$(\mathbf{M})_{\ell\ell'}^{\text{QML}} = \frac{1}{2} \text{Tr} \, \tilde{\mathbf{C}}^{-1} \tilde{\mathbf{P}}^{\ell} \tilde{\mathbf{C}}^{-1} \tilde{\mathbf{P}}^{\ell'}.$$
(2.14)

In the Gaussian, isotropic case, QML is a lossless estimator that recovers all the relevant information contained in the data. The deconvolution and covariance matrices of the estimates are then equal to the inverse of the Fisher information matrix. While being no longer theoretically optimal for anisotropic theories, Pontzen and Peiris (2010) showed that QML remained superior to PCL.

PCL and QML can be contrasted in terms of the computational complexity and quality of the final power spectrum estimates. They both depend on the mask, but QML additionally requires an accurate model of the pixel-pixel covariance matrix. This model requires priors on the fiducial theory $\{C_\ell\}$ and on the additional correlations present in the data, such as noise and systematics. Although the pseudo-spectrum estimator does not explicitly use such priors, it is equivalent to a maximum likelihood analysis when a flat spectrum is assumed in place of a more motivated choice for the pixel-pixel covariance matrix (Efstathiou, 2004b; Pontzen and Peiris, 2010). As a result, PCL yields nearly optimal estimates when the power spectrum is close to flat and with no anisotropic contributions. Moreover, the sensitivity to the shape of the spectrum decreases at small scales as the number of observed modes increases. A simple inversion of the mask then maximises the likelihood function and the variance of PCL reaches its minimum, namely the inverse of the Fisher matrix (Efstathiou, 2004b, 2006), i.e.,

$$\mathbf{V}_{\ell\ell'}^{\text{PCL}} \approx \frac{2\mathcal{C}_{\ell}\mathcal{C}_{\ell'}}{2\ell+1} (\mathbf{M}^{-1})_{\ell\ell'}^{\text{PCL}} \approx (\mathbf{M}^{-1})_{\ell\ell'}^{\text{QML}} = \mathbf{V}_{\ell\ell'}^{\text{QML}}.$$
(2.15)

Note that this result is only valid in the regime where the signal dominates in the covariance matrix.

To illustrate the contrast between the PCL and QML estimates, we calculated and compared their covariance matrices in three realistic settings of interest, involving typical masks and spectra of CMB and galaxy catalogues. The masks are shown in Fig. 2.1, the theory power spectra in Fig. 2.2, and the resulting covariance matrices in Fig. 2.3. For the CMB, we considered the KQ85 mask (smoothed, galaxy part only) and the best-fit theory angular power spectrum of the Wilkinson Microwave Anisotropy Probe 9-year data release (WMAP, see e.g., Bennett et al. 2012; Hinshaw et al. 2012). As a realistic setting for a galaxy survey, we considered a mask created from the sky coverage of the SDSS DR6, enlarged and smoothed for stability of the estimates, shown in Fig. 2.1 (the variance rapidly becomes unstable for complex mask shapes and requires binning, as detailed in Sec. 2.3.3). We used CAMB_SOURCES³ (Challinor and Lewis, 2011) to project the matter power spectrum P(k) (corresponding to the WMAP 9-year cosmology) into an angular power spectrum using the redshift distribution of the CMASS sample, using a fixed astrophysical bias $b_g = 2$. The

³http://camb.info/sources/



Figure 2.1: Fiducial CMB and galaxy survey masks used to calculate the covariance matrices of Fig. 2.3. These masks are in Galactic coordinates and approximate the WMAP KQ85 mask and the SDSS DR6 sky coverage respectively.

CMASS sample will be described in further detail in Sec. 2.3.5.

The CMB spectrum varies across three orders of magnitude in the range $0 < \ell < 50$ and, as expected, QML performs better than PCL in this range. In particular, the variance of the largest scale modes is up to 20% smaller compared with PCL. For this mask, PCL is a good estimator for $\ell > 50$; its variance is typically within ~ 10% of that expected for a maximum likelihood estimator, although the degree of suboptimality will depend on the noise and the geometry of sky cut under consideration (Efstathiou, 2004a,b, 2006; Hamimeche and Lewis, 2009). The variance of the estimates is much larger for the galaxy survey mask, but since the CMASS spectrum varies by only one order of magnitude in $0 < \ell < 200$, the variance of the PCL and QML estimates in fact only differ by a few percent for $\ell > 15$. In conclusion, since LSS power spectra are close to flat, PCL yields nearly optimal estimates, when no motivated prior for \tilde{C} is available. However, the degree of suboptimality increases when considering masks with complex geometries, and in the presence of non-isotropic contributions (e.g., spatial fluctuations due to calibration errors). In this case, the implicit assumptions in the PCL estimator are poor priors for a maximum likelihood analysis, and a better model of the pixel-pixel covariance matrix must be used in the QML estimator to obtain optimal estimates.

The complexity and resources involved in the PCL and QML algorithms also differ considerably. PCL benefits from fast, low-memory algorithms, and therefore can be applied to resolutions and multipole ranges which are amply sufficient for galaxy survey analyses. By contrast, QML involves the inversion of large covariance matrices and the execution of non-symmetric matrix multiplications. Although our optimised QML algorithm advantageously balances the work-load across processors and minimises memory use, typical galaxy surveys such as the SDSS ($f_{sky} \sim 1/6$) can only be analysed at HEALPIX $N_{side} = 64$ on a personal computer. To ensure that the estimates are minimally



Figure 2.2: Fiducial CMB and CMASS angular power spectra used to calculate the covariance matrices of Fig. 2.3. The grey bands show the cosmic variance.

affected by the resolution, accurate smoothing and masking rules are detailed in Appendix 2.7.1. Neglecting these considerations can lead to significant biases if the estimation is performed at low-resolution. Since high- ℓ maximum likelihood estimates might be desirable at high-resolution in some cases, Gruetjen and Shellard (2012) recently proposed a iterative algorithm (similar to the Newton-Raphson iterative scheme presented in the next section) to converge to the maximum likelihood solution using an augmented PCL basis. This approach will prove useful for obtaining optimal estimates of the damping tail of the CMB spectrum, where a full QML analysis is intractable, and PCL estimates are sub-optimal due to the exponential decay and anisotropic contributions to these modes (Hamimeche and Lewis, 2009). However, this improved estimator is unnecessary for the study of galaxy surveys, since PCL is nearly optimal at the smallest scales where shot noise dominates. Alternatively, the complexity of the QML estimator can be reduced by adopting Karhunen-Loève compression (Tegmark et al., 1997, 1998, 2002b; Vogeley and Szalay, 1996); for further details see Appendix 2.7.2. The results of this paper were obtained without this technique since we were able to run the estimator at N_{side} = 64 without any other approximation, thus covering the scales which are not dominated by shot noise.

2.3.3. Likelihood analysis and band-powers

The size and the shape of the mask strongly influence the variance of power spectrum estimates, and thus, the quality of any subsequent analysis. Typical CMB masks only cause a modest increase in the variance (e.g., $\sim 30\%$ for a WMAP-style mask with $f_{\rm sky} = 75\%$) compared with the full sky case. By contrast, galaxy survey masks typically have $f_{\rm sky} < 20\%$ and can increase the variance of



Figure 2.3: Covariance matrices $|V_{\ell\ell'}|$ of the PCL and QML estimates calculated for the CMB and LSS spectra and masks in Figs. 2.1 and 2.2 in the absence of noise and systematics. For a CMB spectrum, the QML variances (top and middle right) are smaller and less correlated than PCL variances (top and middle left), and the resulting PCL estimates are significantly suboptimal compared to QML. For the LSS spectrum, which is considerably flatter, this distinction is much less pronounced (bottom left and right panels), as expected by the reduction of QML to PCL in the case of a flat spectrum.

the individual multipoles by factors of 5-10.

This problem is usually addressed by estimating the power spectrum in multipole bands, which smoothes the power spectrum, Gaussianises the likelihood function, reduces the sensitivity to the input prior, and decreases the variance of the estimates. In particular, this approach proves useful for PCL since its prior can be significantly suboptimal when using complex masks and in the presence of anisotropic contributions. The PCL and QML estimators in Eqs. (2.9) and (2.13) are straightforwardly transformed into band-power estimators by considering binned Legendre matrices $\tilde{\mathbf{P}}^b = \sum_{\ell \in L_b} \tilde{\mathbf{P}}^{\ell}$, where L_b denotes the range of multipoles included in the *b*-th bin. The pixel-pixel covariance matrix $\tilde{\mathbf{C}}$ is unchanged and calculated as before, but the coupling matrices in the binned formulation are given by

$$(\mathbf{M})_{bb'} = \sum_{\ell \in L_b} \sum_{\ell' \in L_{b'}} (\mathbf{M})_{\ell\ell'}.$$
 (2.16)

The expectation value of each band-power estimate reads

$$\langle \hat{C}_b \rangle = \sum_{\ell} W_{b\ell} C_{\ell}, \qquad (2.17)$$

where the window function $W_{b\ell}$ is constructed with the coupling matrix through

$$W_{b\ell} = \sum_{b'} (\mathbf{M})_{bb'}^{-1} (\mathbf{M})_{b'\ell}.$$
 (2.18)

Hence, to be confronted with band-power estimates, the theory power spectrum must be transformed into band-powers using these window functions.

Using large band-powers guarantees that the likelihood function is sufficiently close to Gaussian for QML to deliver maximum-likelihood estimates. However, intermediate situations with smaller bins and complex masks might not fulfil this condition, leading to sub-optimal estimates. This issue can be addressed by iteratively converging to the maximum-likelihood solution using a Newton-Raphson scheme (Bond et al., 1998, 2000; Knox et al., 1998). With the so-called Newton-Raphson maximum likelihood (NRML) estimator, each iteration improves the previous one using

$$\delta C_{\ell} = \sum_{\ell'} (\mathbf{M}^{-1})^{\text{QML}}_{\ell\ell'} \frac{1}{2} \text{Tr} \left[\left(\mathbf{x} \mathbf{x}^{t} - \tilde{\mathbf{C}} \right) \left(\tilde{\mathbf{C}}^{-1} \tilde{\mathbf{P}}^{\ell} \tilde{\mathbf{C}}^{-1} \right) \right].$$
(2.19)

This approach is in fact equivalent to calculating the *i*-th estimate by feeding QML with the previous

iteration. The formulation of the NRML estimator then simplifies to

$$\hat{C}_{\ell}^{i+1} = \tilde{\mathbf{x}}^t (\mathbf{E}^{\ell})^i \tilde{\mathbf{x}}, \tag{2.20}$$

where $(\mathbf{E}^{\ell})^i$ explicitly makes use of the covariance matrix calculated from \hat{C}^i_{ℓ} , namely $\tilde{\mathbf{C}} = \sum_{\ell} \tilde{\mathbf{P}}^{\ell'} \hat{C}^i_{\ell} + \tilde{\mathbf{N}}$. Hence, a critical issue is to construct a reliable covariance matrix at each iteration.

Both PCL and QML estimates can be extended to estimate the angular cross-power spectra (and hence cross band-powers) between two maps denoted by \tilde{x}_1 and \tilde{x}_2 . The PCL estimator reads

$$\hat{C}_{\ell}^{\text{PCL,cross}} = \tilde{\mathbf{x}}_{1}^{t} \mathbf{E}_{\text{PCL}}^{\ell} \tilde{\mathbf{x}}_{2}, \qquad (2.21)$$

and implicitly assumes flat auto-power spectra and zero cross-power spectrum as priors. On the contrary, since QML derives from a likelihood function, the estimator must be adapted (following e.g., Padmanabhan et al. 2005; Tegmark and de Oliveira-Costa 2001) by considering the input data vector as a concatenation of the two maps,

$$\tilde{\mathbf{x}} = \begin{pmatrix} \tilde{\mathbf{x}}_1 \\ \tilde{\mathbf{x}}_2 \end{pmatrix}, \qquad (2.22)$$

and by using a pixel-pixel covariance matrix which incorporates all the information about the maps and their cross-correlation,

$$\tilde{\mathbf{C}} = \begin{pmatrix} \tilde{\mathbf{C}}_{11} & \tilde{\mathbf{C}}_{12} \\ \tilde{\mathbf{C}}_{12}^{\dagger} & \tilde{\mathbf{C}}_{22} \end{pmatrix}.$$
(2.23)

This formulation makes use of priors for the three power spectra, and must include models of the additional correlations and noise present in the data. The auto- and cross-spectra can be simultaneously estimated from $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{C}}$ using the usual formulation of QML, i.e., Eqs. (2.13) and (2.14), In particular, the matrices to be used in the estimator with the covariance matrix of Eq. (2.23) in order to calculate \hat{C}_{ℓ}^{11} , \hat{C}_{ℓ}^{22} and \hat{C}_{ℓ}^{12} read

$$\tilde{\mathbf{P}}_{11}^{\ell} = \begin{pmatrix} \tilde{\mathbf{P}}^{\ell} & 0\\ 0 & 0 \end{pmatrix}, \tilde{\mathbf{P}}_{22}^{\ell} = \begin{pmatrix} 0 & 0\\ 0 & \tilde{\mathbf{P}}^{\ell} \end{pmatrix}, \tilde{\mathbf{P}}_{12}^{\ell} = \begin{pmatrix} 0 & \tilde{\mathbf{P}}^{\ell}\\ \tilde{\mathbf{P}}^{\ell} & 0 \end{pmatrix}.$$
(2.24)

2.3.4. Galaxy surveys, shot noise and systematics

Galaxy catalogues are usually provided as lists of objects whose positions and properties, such as photometric colours, were measured by an instrument in the context of a sky survey. To relate to the dark matter distribution, a catalogue must be pixelised into a number count map $\tilde{\mathbf{G}}$, and then transformed into an overdensity map $\tilde{\mathbf{x}}$. Given a pixelisation scheme, if G_i denotes the number of objects in the *i*-th pixel (described by its size Ω_i and the position of its centre \mathbf{n}_i), the overdensities are constructed as

$$\tilde{x}_i = \frac{\tilde{G}_i}{\Omega_i \bar{G}} - 1, \tag{2.25}$$

where $\bar{G} = N_{obj}/\Delta\Omega$ is the average number of objects per steradian. N_{obj} is the total number of objects in the catalogue, and $\Delta\Omega$ is the total surface outside the mask. The Poisson sampling of the observed tracers naturally gives rise to shot noise, characterised by a diagonal noise matrix \tilde{N} (Huterer et al., 2001) such that $(\tilde{N})_{ij} = \bar{G}^{-1}\delta_{ij}$, where δ is the Kronecker delta. Moreover, the power spectrum estimates of galaxy survey overdensity maps always include a constant bias term due to shot noise,

$$\langle \hat{C}_{\ell} \rangle = C_{\ell} + \frac{1}{\bar{G}}.$$
(2.26)

Since the cosmological contribution C_{ℓ} steadily decreases with ℓ in the linear regime, shot noise usually dominates on small scales for catalogues with small number densities (i.e., for ℓ higher than a limit determined by \bar{G}).

So far, we assumed that the map x was the result of cosmological clustering encapsulated by a theory power spectrum $\{C_\ell\}$. In other words, the relevant correlations in the map were due to $\{C_\ell\}$ and the shot noise $1/\overline{G}$. However, observations $\tilde{\mathbf{x}}^{\text{obs}}$ are often contaminated by various signals introducing spurious correlations and requiring appropriate modelling to avoid suboptimal estimates. In particular, $\tilde{\mathbf{x}}^{\text{obs}}$ can always be described as the superposition of the true cosmological signal $\tilde{\mathbf{x}}^{\text{true}}$ and a contamination part due to systematics. We will assume that templates of the systematics are available, namely n_{sys} maps denoted by \mathbf{c}^k with $k = 1, \ldots, n_{\text{sys}}$. In this case, one can adopt a model for the contamination signal, and estimate its parameters from the data. The best-fit contamination model is then subtracted from the measured power spectra. This approach was recently used to correct the angular power spectra and 2-pt correlation functions of the CMASS sample with a best-fit linear contamination model (Ho et al., 2012; Ross et al., 2011).

A more robust approach to mitigate the influence of the systematics is to incorporate their

contribution in the pixel-pixel covariance matrix with large coefficients ξ_k , i.e.,

$$\tilde{\mathbf{C}}_{ij} = \tilde{\mathbf{S}}_{ij} + \tilde{\mathbf{N}}_{ij} + \sum_{k} \xi_k \tilde{c}_i^k \tilde{c}_j^k.$$
(2.27)

This technique, known as mode projection (Ho et al., 2008; Slosar et al., 2004), assigns a very large variance to the modes corresponding to the systematics in pixel space, such that they do not influence the power spectrum estimates⁴. As a result, the QML estimates are unbiased in the power spectrum of $\tilde{\mathbf{x}}^{\text{true}}$ as $\xi_i \to \infty$, provided that the contamination signal can be described as a linear combination of the templates. In the Bayesian perspective, mode projection is equivalent to marginalising over the parameters of a linear model of the contamination. Hence, non-linear contamination or neglected systematics leave residual biases that must be eliminated by other means, for example through masking or modelling of the contamination signal.

2.3.5. Illustration: recovering the power spectrum of the CMASS sample

The ongoing Baryon Oscillation Spectroscopic Survey (BOSS) is part of SDSS-III (Eisenstein et al., 2011) and aims to measure the spectroscopic redshifts of 1.5 million galaxies, 160,000 quasars and various ancillary targets from SDSS photometry (Gunn et al., 1998, 2006). The CMASS spectroscopic sample includes extended sources selected using colour magnitude cuts to produce a roughly volume limited sample in the redshift range 0.4 < z < 0.7. In the DR9 release of BOSS, CMASS uses data taken up to the end of July 2011 and covers 3344 deg² in the Northern and Southern Galactic caps.

Manera et al. (2013) presented a set of 600 mock catalogues for the CMASS DR9 sample, constructed based on a Λ CDM cosmology defined by { $\Omega_m = 0.274, \Omega_b h^2 = 0.0224, h = 0.70, n_s = 0.95, \sigma_8 = 0.8$ }, and evolved using 2nd-order Lagrangian perturbation theory (Crocce et al., 2006; Scoccimarro, 1997). These mock catalogues were used to compute accurate covariance matrices for CMASS and constrain cosmological parameters (Sánchez et al., 2012), test deviations from General Relativity (Samushia et al., 2013) and measure the scale-dependent halo bias (Ross et al., 2013).

To evaluate the performance of the previously described angular power spectrum estimators, we compared the mean and standard deviation of the power spectrum estimates of the mock catalogues with the theoretical expectations. We only considered the CMASS DR9 mock catalogues in the

⁴This approach does not prevent the extraction of cosmological signals that happen to have the same power spectra as the systematics. In particular, projecting out a systematic is equivalent to ignoring one mode defined in pixel space.



Figure 2.4: Mask of the CMASS DR9 sample in the Northern Galactic Cap, in equatorial coordinates.

Northern Galactic cap (NGC), covering 2635 deg² as shown in Fig. 2.4. We computed a theory prediction for the angular power spectrum of the underlying dark matter with CAMB_SOURCES (Challinor and Lewis, 2011) using a redshift distribution parametrised by an Edgeworth expansion of the redshift histogram of CMASS objects.

We calculated the PCL and QML estimates of the individual mock catalogues at HEALPIX resolutions N_{side}^{PCL} =128 and N_{side}^{QML} =64. The mean and variance of the estimates compared with theoretical expectations are shown in Fig. 2.5. Note that the error bars of the power spectrum estimates are correlated; see Eq. (2.11). We subtracted the shot noise and also multiplied the theory prediction by a scale-independent bias of 1.9, the value used to construct these mock catalogues (Manera et al., 2013). We were able to recover the theory power spectrum and the covariance of the estimates with good precision using the smoothing, masking and band-limit rules defined in Appendix 2.7.1. The estimates and the theory predictions from CAMB_SOURCES were insensitive to small changes in the redshift distribution. However, we observed that inconsistent smoothing, masking and band-limit rules led to biases in the recovered power spectra, resulting from a mismatch between the model pixel-pixel covariance matrices and the information content in the data. As expected, due to the flatness of the power spectrum and the simple geometry of the mask, QML performed only marginally better than PCL, which yielded nearly-optimal estimates in the absence of systematics. The theory spectrum was converted into band-powers using the exact window functions in Eqs. (2.17) and (2.18), which are shown in the bottom panel of Fig. 2.5.



Figure 2.5: Average PCL and QML estimates of the 600 CMASS mock catalogues (Manera et al., 2013). Shot noise was subtracted from the estimates, and the angular power spectrum was calculated with CAMB_SOURCES (black dashed line) using a scale-independent astrophysical bias $b_g = 1.9$. The theory band-powers (black solid line) were obtained by applying the exact PCL and QML binning window functions of Eqs. (2.18) and (2.17) (bottom panel). We were able to recover the theory power spectrum and the covariance of the estimates with good precision using the smoothing, masking and band-limit rules defined in Appendix 2.7.1.

2.4.1. Data and subsamples

We considered the Richards et al. (2009) catalogue of photometric quasars, which is based on the Sixth Data Release (DR6) of the SDSS (Adelman-McCarthy et al., 2008). The objects in this catalogue - which we call RQCat as in Pullen and Hirata (2013) - were photometrically selected by a binary Bayesian classifier trained in 4D colour space using several catalogues of spectroscopicallyconfirmed stars and quasars. The final version of the catalogue includes 1,172,157 objects with several quality and technical flags, which can be exploited to apply further systematics cuts and obtain cleaner samples. The Bayesian classifier was initially applied to all point sources in the DR6 release without restriction, and explicitly assumed that 95% of the input objects were stars (i.e., constant star and quasar priors $p_{\text{star}} = 0.95$ and $p_{\text{quasar}} = 0.05$). It was also applied with stronger priors ($p_{\text{star}} = 0.98$) to objects with photometric redshifts \tilde{z}_p (estimated by the SDSS pipeline, as opposed to the true redshift z) in three redshift ranges $0 < \tilde{z}_p \le 2.2, 2.2 < \tilde{z}_p \le 3.5$ and $\tilde{z}_p > 3.5$ to achieve higher efficiency⁵ and completeness. The efficiency is degraded in the two higher redshift ranges since the Bayesian classifier performs worse due to the overlap between the stellar and quasar loci in the colour space (at $z \sim 2.6$). More sophisticated algorithms such as XDQSO (Bovy et al., 2011a) were developed to specifically address this issue and identify higher-redshift quasars for spectroscopic follow-up in the context of BOSS (Ross et al., 2012b). However, in this work we focused on $z \leq 2.2$ objects, where the performance of the binary Bayesian classifier used by Richards et al. (2009) is satisfactory.

Colour-based selection of quasars is difficult, and RQCat is expected to be significantly contaminated by stars, which are often misclassified as quasars due to similar colours. Hence, the full catalogue cannot be used as a statistical sample for direct power spectrum analysis due to its low efficiency (lower than 80%). In this work, we restricted ourselves to good UV-excess low-redshift objects, defined as u - g < 1.0 and $\tilde{z}_p < 2.2$ (the corresponding flags are GOOD>0, UVX=1 and LOWZ=1). This sample, denoted by UVX-LOWZ, is the least contaminated by stars and achieves 96.3% ± 1.2 efficiency (Richards et al., 2009).

⁵The efficiency, or purity, of a catalogue denotes the fraction of objects which are quasars. It characterises the ability of the classifier to separate quasars from stars.

Previous studies considered the UVX sources in RQCat for cosmological analyses, e.g., studying the environment of quasars (Myers et al., 2006, 2007a), detecting the ISW effect (Corasaniti et al., 2005; Giannantonio et al., 2006, 2008; Giannantonio et al., 2012) and constraining PNG (Slosar et al., 2008; Xia et al., 2010, 2011). Recent work by Pullen and Hirata (2013), corroborated by Giannantonio et al. (2014), found that the UVX objects were significantly contaminated, as indicated by the cross-spectra of redshift bins which exhibited excess power at the largest scales.

The two main sources of systematics in quasar photometric catalogues are contamination and calibration errors. The origin of contamination lies in the classification stage: selecting quasars based on photometric data is a complex task. Various objects can be misclassified as quasars, and are thus present in the final catalogues. Since the clustering properties of these contaminants differ from those of quasars, they affect the measured power spectra and can jeopardise the interpretation of the data. Calibration errors, on the other hand, are present in the catalogue regardless of the ability of the classifier to separate stars and quasars. In the ideal case, a perfect classifier applied to all point sources detected by a given instrument will lead to a sample with clustering properties purely due to cosmological physics. However, real instruments are not perfectly calibrated, and observing conditions also change with time, introducing spurious correlations due to variations in the number of detected sources on the sky. In addition, calibration errors impact the apparent magnitude estimates, which propagate through the Bayesian classifier (since the latter does not model or account for them), inducing a spatial dependence in its efficiency.

Contamination and calibration issues can be addressed in different ways. First, it is important to reduce the catalogue of interest by selecting the most reliable objects (here UVX sources) and also restricting the analysis to the most reliable areas of the sky. Secondly, corrections can be applied to the power spectrum estimates themselves to minimise the remaining spurious correlations. Alternatively, one can opt for a Bayesian analysis and marginalise over the systematics in the cosmological analysis. In this study, we focused on the sample reduction approach. We separated the UVX-LOWZ sample into four subsamples by selecting objects with photometric redshifts \tilde{z}_p in bins with ranges [0.5, 1.3], [1.3, 1.8], [1.8, 2.2] and [1.3, 2.2]. These samples are called Low-z, Mid-z, High-z and Mid+High-z respectively, as referred to as the RQCat subsamples. We rejected low redshift quasars ($\tilde{z}_p < 0.5$) because their power spectra were severely contaminated: this can be attributed to strong stellar contamination, and to the fraction of low-z quasars which are extended sources and were therefore not processed by the Bayesian classifier. The corresponding incompleteness is non-trivial and likely to depend on observational effects such as dust absorption and seeing variations.



Figure 2.6: Distributions of the photometric and spectroscopic redshift estimates (\tilde{z}_p and \tilde{z}_s) of the RQCat UVX-LOWZ sample, cross-matched with the SDSS-DR7, BOSS and 2SLAQ spectroscopic quasar catalogues (blue, red and green dots). The dashed lines indicate the photometric redshift cuts used to assemble the four RQCat subsamples. The photometric redshift estimates are seen to be unreliable and cannot be used to estimate the redshift distributions of the photometrically-selected subsamples. On the contrary, the cross-matched samples have reliable spectroscopic redshifts, and can be used for this purpose. However, one must apply the relevant completeness corrections in order to account for the change of selection function between the photometric and cross-matched subsamples (due to e.g., different magnitude limits).

2.4.2. Theory predictions

In order to calculate theoretical predictions for the angular power spectra of the four RQCat subsamples, we used CAMB_SOURCES (Challinor and Lewis, 2011), a high-precision code which projects the 3D matter power spectrum P(k) into angular auto- and cross-power spectra. Since this study is focused on the impact of the systematics on the observed power spectra, we fixed the cosmological parameters to *Planck* Λ CDM best-fit values⁶ (Planck Collaboration, 2013g). We opted

⁶Fixed to $\Omega_c h^2 = 0.188, \Omega_m = 0.315, \Omega_b h^2 = 0.02205, H_0 = 67.3 \text{ kms}^{-1} \text{Mpc}^{-1}, \ln(10^{10} A_s) = 3.089, n_s = 0.9603$, and $\tau = 0.089$.



Figure 2.7: Histograms of the redshifts and apparent magnitudes of UVX-LOWZ objects in RQCat, and of objects with good spectra found in the SDSS-DR7, BOSS and 2SLAQ spectroscopic quasar catalogues. In addition to the higher number of objects, the SDSS-DR7 cross-matched sample has redshift and magnitude distributions close to that of RQCat, indicating similar selection functions. Hence, the redshift distributions of the RQCat subsamples can be estimated using the cross-matched sample, with only minor completeness corrections to account for the differences in magnitude limits. By contrast, BOSS and 2SLAQ target significantly different redshift and magnitude ranges, and the completeness corrections required to estimate the redshift distributions of RQCat are strongly redshift- and magnitude-dependent. In addition, the latter are limited by sample variance due to the smaller number of objects.

for a scale-independent linear bias to relate the observed galaxy clustering to dark matter.

Although the matter power spectrum P(k) only depends on cosmological parameters, computing angular power spectrum predictions requires additional knowledge about the distributions and properties of the samples under consideration. In particular, the auto-angular power spectrum of a sample of interest reads

$$C_{\ell} = \frac{2}{\pi} \int dk k^2 P(k) [W_{\ell}(k)]^2, \qquad (2.28)$$

where the window function $W_{\ell}(k)$ includes several cosmological effects, as detailed in Challinor and Lewis (2011). In fact, $W_{\ell}(k)$ requires two quantities in addition to the standard cosmological parameters: the unit-normalised redshift distribution of tracers, denoted by n(z), and the logarithmic slope of the number counts s, which accounts for the effect of magnification due to lensing. A simplified formula for $W_{\ell}(k)$, which we use for illustration purposes only (see Challinor and Lewis (2011) for the full formalism), reads

$$W_{\ell}(k) = \int dz \left[b_g n(z) + 2(2.5s - 1)f(z) \right] D(z) j_{\ell}(kr),$$
(2.29)

where b_g is the linear galaxy bias, D(z) is the growth factor, j_ℓ is the spherical Bessel function, r(z)



Figure 2.8: Low-resolution redshift- and magnitude-dependent completeness corrections for estimating the redshift distributions of the RQCat subsamples through cross-matching with the SDSS-DR7, BOSS and 2SLAQ spectroscopic quasar catalogues. As expected from Fig. 2.7, the redshift distribution of the cross-matched sample with SDSS-DR7 only requires weak magnitude corrections to relate to that of RQCat. On the contrary, BOSS and 2SLAQ require significant corrections, since they target different redshift and magnitude ranges than RQCat.

is the comoving distance, and f(z) is the lensing window function, giving rise to magnification. In this work, we assumed that the logarithmic slope of the number counts, defined as

$$s = \frac{d\log N(m)}{dm},\tag{2.30}$$

was constant in each redshift range. We estimated its value for the four RQCat subsamples by calculating the slope of the histogram of number counts in terms of the *g*-band PSF magnitude at g = 21. We found s = 0.18, 0.89, 0.89 and 0.87 respectively for these subsamples, consistent with previous studies in similar redshift ranges (Pullen and Hirata, 2013; Xia et al., 2009). The choice of a scale- and redshift-independent linear bias is motivated by previous studies of RQCat and of low-redshift quasars in general (e.g., Giannantonio et al. 2014; Myers et al. 2007a; Pullen and Hirata 2013; Sherwin et al. 2012). The redshift evolution of the bias in 0.5 < z < 2.2 proves to be smaller than the uncertainty on the completeness corrections, and thus marginally affects the predicted angular power spectra.

2.4.3. Redshift distributions estimates

The quasars in each RQCat subsample are characterised by a normalised redshift distribution n(z), with n(z)dz corresponding to the probability of finding a quasar with redshift between z and z + dz. Consequently, n(z) incorporates the physical distribution (i.e., originating from the quasar



Figure 2.9: Final estimates of the redshift distributions of the four RQCat photometric subsamples, calculated by cross-matching with the SDSS-DR7 spectroscopic quasar catalogue (Schneider et al., 2010). The accuracy of the redshift distribution estimates is essential for obtaining robust angular power spectrum estimates. Maps of the photometric and cross-matched samples are shown on the right subpanels at HEALPIX resolution $N_{side} = 64$. The dashed vertical lines in the main subpanels indicate the photometric redshift cuts used to construct the RQCat subsamples. Since the cross-matched samples have different selection functions, their redshift distributions (red thin histograms) must be corrected in order to accurately estimate the redshift distributions of the RQCat samples. The thick black histograms show the final estimates obtained by applying magnitude- and pixel-dependent completeness corrections, and were fitted with a superposition of Gaussian distributions (solid lines) for use in CAMB_SOURCES.

luminosity function), the survey characteristics (such as the magnitude limits), and the photometric redshift cuts used to construct the subsample under consideration. The simplest estimator for n(z)is a normalised histogram of the photometric redshifts of all objects in each subsample of interest. However, in practice, this approach does not yield good estimates due to the large uncertainties in quasar photometric redshift estimates. To illustrate this issue, we used the SDSS-DR7, BOSS and 2SLAQ spectroscopic quasar catalogues (Croom et al., 2009; Pâris et al., 2012; Schneider et al., 2010) to find objects in the four RQCat subsamples for which reliable spectra, and thus good spectroscopic redshifts, were available. Figure 2.6 shows the photometric and spectroscopic redshifts (\tilde{z}_p and \tilde{z}_s , respectively) of the cross-matched objects. The dispersion of the points (\tilde{z}_p , \tilde{z}_s) around $\tilde{z}_p = \tilde{z}_s$ demonstrates that the photometric redshift estimates of quasars suffer from large uncertainties and catastrophic failures, yielding a large fraction of spectroscopic redshifts outside the photometric windows used to construct the RQCat subsamples, indicated by the dashed lines. The photometric redshift estimates do not accurately follow the underlying redshift distributions n(z), and cannot be used to compute accurate angular power spectrum predictions.

Nevertheless, one can use the redshift distributions of the cross-matched samples, which can be calculated with great accuracy using the spectroscopic redshift estimates, whose uncertainties are negligible compared to the precision required for n(z). However, the redshift, magnitude and spatial distributions of the cross-matched samples may deviate from those of the photometric samples due to differences in their selection functions, caused, e.g., by different magnitude limits, redshift ranges, or sky coverage. In order to avoid biases in the redshift distribution estimates, one must include a completeness correction factor, denoted by f_c , which is a function of redshift z, magnitude g and position on the sky n.

In practice, we divide the redshift and magnitude domains into bins denoted by $[z_i^{\min}, z_i^{\max}]$ (centred at z_i) and $[g_j^{\min}, g_j^{\max}]$ (centred at g_j) respectively, and the sky into pixels denoted by k. For each RQCat subsample, an estimator of n(z) at redshift $z = z_i$ is given by

$$\tilde{n}(z_i) = C \sum_{jk} \frac{N_{\rm obj}({\rm cross}, [z_i^{\rm min}, z_i^{\rm max}], [g_j^{\rm min}, g_j^{\rm max}], k)}{\tilde{f}_c(z_i, g_j, k)},$$
(2.31)

where *C* is a normalisation constant (such that $\sum_{i} \tilde{n}(z_{i}) = 1$) and $N_{obj}(cross, [z_{i}^{min}, z_{i}^{max}], [g_{j}^{min}, g_{j}^{max}], k)$ is the number of objects in the *cross-matched* sample in the *i*th redshift bin, *j*th magnitude bin and *k*th pixel.

With an estimator of n(z) in hand, we now discuss how to construct the completeness correction.

A simple estimator for the correction in the (i, j, k)th volume element reads

$$\tilde{f}_{c}(z_{i}, g_{j}, k) = \frac{N_{\rm obj}({\rm cross}, [z_{i}^{\min}, z_{i}^{\max}], [g_{j}^{\min}, g_{j}^{\max}], k)}{N_{\rm obj}({\rm photo}, [z_{i}^{\min}, z_{i}^{\max}], [g_{j}^{\min}, g_{j}^{\max}], k)},$$
(2.32)

where $N_{obj}(photo, [z_i^{min}, z_i^{max}], [g_j^{min}, g_j^{max}], k)$ is the number of objects in the *photometric* sample (namely, one of the RQCat subsamples), thus constructed using the photometric redshifts.

The resolution of previous multidimensional histograms is limited by sample variance and by the uncertainties on (both photometric and spectroscopic) redshift and magnitude estimates⁷. To estimate the redshift distribution of the RQCat subsamples, we investigated the use of the SDSS-DR7, BOSS and 2SLAQ spectroscopic quasar catalogues. The redshift and magnitude distributions of the UVX-LOWZ objects in RQCat cross-matched with these catalogues are shown in Fig. 2.7. We found that the BOSS- and 2SLAQ-based cross-matched samples were small and had selection functions quite different from RQCat. As a result, the estimator required significant completeness corrections, and the multidimensional histograms had to be constructed with large bins to reduce sample variance and biases due to uncertainties on the redshift and magnitude estimates. On the other hand, cross-matching with SDSS-DR7 led to large samples which had selection functions similar to that of the ROCat subsamples. Fig. 2.8 shows low-resolution estimates of completeness corrections arising when cross-matching the whole UVX-LOWZ sample (union of the RQCat subsamples) with SDSS-DR7, BOSS and 2SLAO, where spatial dependence was neglected for purposes of illustration. Even at such low resolution, the completeness corrections for BOSS and 2SLAQ exhibit strong redshift and magnitude dependences (and are limited by sample variance due to the low number of cross-matched objects). For these reasons, the final redshift distribution estimates of the four RQCat samples were calculated using the SDSS-DR7 catalogue. The right panels of Fig. 2.9 show number count maps of the RQCat subsamples and the SDSS-DR7-based cross-matched samples. The numbers of objects in the respective subsamples are summarised in Table 2.1.

With the high number of cross-matched objects, we were able to test various assumptions for the completeness corrections, such as the weakness of the redshift-dependence, and choose resolutions that yielded the best estimates. Since SDSS-DR7 spans the same redshift range as RQCat, no redshift-dependent corrections were required (Fig. 2.8 shows that the correction is only weakly redshift dependent, apart from a mild transition at z = 2. Neglecting this transition did not significantly impact the final estimates). We used magnitude-dependent completeness corrections at resolution $\Delta g = 1.0$, due to the different magnitude limits in the SDSS photometric and spectroscopic data.

 $^{^{7}}$ We will neglect the uncertainty in the sky position since it is negligible compared to the pixel size used for the completeness correction.

Table 2.1: Number of objects in the four RQCat redshift subsamples (N_{obj} photo), and for which good spectra (and thus good spectroscopic redshifts) were found in the SDSS-DR7 quasar catalogue (N_{obj} cross). These cross-matched samples were used to estimate the redshift distributions shown in Fig 2.9.

	Low-z	Mid-z	High-z	Mid+High-z
N _{obj} photo	95,185	109,713	92,740	202,453
N _{obj} cross	19,328	17,589	10,654	28,243

Finally, the SDSS-DR7 spectroscopic quasar catalogue was assembled from different data releases of SDSS and is known to be non-uniform on the sky (Schneider et al., 2010). Indeed, Fig. 2.9 shows that the cross-matched samples contain regions with a greater number of objects, in some cases exploring fainter magnitudes. To address this spatial dependence, we calculated the magnitude-dependent completeness corrections in individual pixels at HEALPIX resolution $N_{side} = 16$. The final estimates $\tilde{n}(z_i)$ are shown in Fig. 2.9, and were fitted by superpositions of Gaussian distributions for use in CAMB_SOURCES.

Figure 2.10 shows the estimates obtained with CAMB_SOURCES (thus using the full formalism from Challinor and Lewis 2011) with our best estimates for s and n(z), compared with the same estimates where magnification and completeness corrections were neglected. The large differences between the resulting angular power spectra demonstrate that these effects must be accounted for, and carefully estimated from the data in order to avoid significant biases in the theory predictions.

2.4.4. Masks and systematics

We considered five sources of systematics: stellar contamination, dust absorption, seeing, airmass and sky brightness. Following Pullen and Hirata (2013), we constructed the stellar density map from SDSS DR6 point sources with 18.0 < r < 18.5 and i < 21.3. For the extinction, we used the dust maps from Schlegel et al. (1998) with the corrections by Peek and Graves (2010). Templates for seeing, airmass and sky brightness were constructed with the MANGLE software (Hamilton and Tegmark, 2004; Swanson et al., 2008) using data retrieved from the FIELDS table in the SDSS CAS server. All maps were binned onto the HEALPIX grid at resolution $N_{side} = 128$, and are shown in Fig. 2.11.

We designed three sky masks by excluding pixels based on their values in the systematics maps. The thresholds are summarised in Table 2.2, and the resulting masks are presented in Fig. 2.12.



Figure 2.10: Theory predictions for the four RQCat subsamples computed with CAMB_SOURCES. The black lines show the final theory predictions, calculated with the best redshift distribution estimates presented in Sec. 2.4.3, for a *Planck* cosmology and a fixed galaxy bias $b_g = 2.3$. The blue dashed lines show the predictions obtained by neglecting the effect of magnification, and the red dot-dashed lines by neglecting the completeness correction in the estimated redshift distributions (i.e., using the redshift histograms of the cross-matched samples without accounting for the differences in magnitude limits).



Figure 2.11: Systematics templates used in this analysis, and the (dimensionless) angular power spectra \tilde{C}_{ℓ} of their overdensity maps.

Table 2.2: Systematics thresholds used to restrict the power spectrum analysis of RQCat to the most reliable regions of the sky to minimise contamination from calibration errors. The maps of the systematics are shown in Fig. 2.11, and the resulting masks in Fig. 2.12.

Systematic (unit)	Mask 1	Mask 2	Mask 3
Seeing (arcsec)	2.0	1.6	1.55
Reddening (mag)	0.05	0.05	0.045
Stellar density (stars/deg ²)	562	400	350
Airmass (mag)	1.4	1.3	1.25
Sky brightness (nmgy/arcsec ²)	2×10^{-9}	1.8×10^{-9}	1.75×10^{-9}



Figure 2.12: Masks used for the power spectrum analysis of RQCat, in Equatorial coordinates. Retained regions are based on thresholds summarised in Table 2.2 and the systematics templates of Fig. 2.11. Additional excised rectangles follow Pullen & Hirata (2012). The three masks respectively have $f_{sky} = 0.148, 0.121, and 0.101$.

Following Pullen and Hirata (2013), we also excised rectangular regions with missing data⁸. Our first mask is therefore very similar to those used in previous studies of RQCat and constitutes the reference mask. The two other masks use more aggressive systematics cuts.

The cut-sky angular power spectra, \tilde{C}_{ℓ} , of the systematics maps for the three masks are shown in Fig. 2.11. Interestingly, stellar density and dust absorption templates display strong large-scale power ($\ell < 30$), and calibration error templates (seeing, airmass, sky brightness) have notable features at $\ell \sim 70,110$ and 150. The masked template maps have these features reduced, but not eliminated. Since the data are known to be affected by these systematics, the measured spectra are likely to be contaminated at these multipoles.

2.4.5. Power spectrum results

We obtained angular band-power estimates with the QML estimator and multipole bins of size

⁸In equatorial (J2000) coordinates, the discarded angular rectangles are $(\alpha, \delta) = (122^{\circ} - 139^{\circ}, -1.5 - (-0.5)^{\circ}), (121^{\circ} - 126^{\circ}, 0^{\circ} - 4^{\circ}), (119^{\circ} - 128^{\circ}, 4^{\circ} - 6^{\circ}), (111^{\circ} - 119^{\circ}, 6^{\circ} - 25^{\circ}), (111.5^{\circ} - 117.5^{\circ}, 25^{\circ} - 30^{\circ}), (110^{\circ} - 116^{\circ}, 32^{\circ} - 35^{\circ}), (246^{\circ} - 251^{\circ}, 8.5^{\circ}, 13.5^{\circ}), (255^{\circ} - 270^{\circ}, 20^{\circ} - 40^{\circ}), (268^{\circ} - 271^{\circ}, 46^{\circ} - 49^{\circ}), (232^{\circ} - 240^{\circ}, 26^{\circ} - 30^{\circ}).$



Figure 2.13: Illustration of the suboptimality of the PCL estimator in the case of the Mid+High-*z* subsample. The top panels show the covariance matrices of the PCL and QML estimates with $\Delta \ell = 11$ for the three masks of Fig. 2.12. The suboptimality of the PCL prior is measured by the fractional increase of variance compared to QML, shown in the right column. The bottom panel shows the resulting effects on the power spectrum estimates, which are more pronounced for the second and third masks due to their complex geometry.

 $\Delta \ell = 11$, which led to a good balance in terms of multipole resolution and variance of the estimates. We did not use the PCL estimator for the final results because the geometry of the second and third masks, in addition to the presence of systematics, yielded significantly suboptimal estimates. To illustrate this point, Fig. 2.13 shows a comparison of the PCL and QML covariance matrices and the band-power estimates of the Mid+High-z subsample for the three masks. Any significant increase of the PCL variance compared to that of QML, especially on diagonal- and nearly-diagonal elements which contain the most significant contributions, demonstrates the suboptimality of the PCL prior. For the first mask, the PCL variance of these elements is at most $\sim 20\%$ greater than the QML variance, indicating that the resulting estimates are nearly optimal. However, for the second and third masks, these elements have a PCL variance up to $\sim 50\%$ greater than that of QML, and the resulting PCL estimates significantly differ from the optimal QML estimates, as shown in the bottom panel of Fig. 2.13. This effect is less pronounced for larger multipole bins (e.g., $\Delta \ell = 31$), as the likelihood becomes less sensitive to the priors on the pixel-pixel covariance matrix. However, the resulting loss of resolution prevents the study of localised multipole ranges affected by systematics. For these reasons we opted for the QML estimator with $\Delta \ell = 11$ in the final analysis. We systematically marginalised over the values the monopole and the dipole by projecting them out. They are poorly constrained from cut-sky data, and may affect the power spectrum estimates over a wide range of multipoles when deconvoling the cut-sky power spectrum into full sky estimates⁹. We used the values $\bar{G}^{-1} = 1.95 \cdot 10^{-5}, 1.55 \cdot 10^{-5}, 1.85 \cdot 10^{-5}$ and $8.15 \cdot 10^{-6}$ respectively for the shot noise of the four RQCat subsamples, calculated from the average number count per steradian assuming 5% stellar contamination.

The auto- and cross-spectra of the four RQCat samples are presented in Figs. 2.14 and 2.15, and the χ^2 values of the theory prediction are listed in Table 2.3. We subtracted the shot noise from the auto-spectra, and used a constant bias, $b_g = 2.3$, following previous studies of these data (Giannantonio et al., 2006, 2008; Pullen and Hirata, 2013; Slosar et al., 2008; Xia et al., 2010). The theory predictions are summarised in Fig. 2.10. We also used the exact window functions $W_{b\ell}$ for converting the theory power spectra into band-powers; see Eq. (2.17). Figure 2.16 shows the cross-correlation power spectra of the quasar samples with the systematics templates, and Table 2.4 lists the corresponding χ^2 values. Details of the χ^2 computation are contained in Appendix 2.7.3.

In Figs. 2.14 and 2.15, the top panels show the final band-power estimates, where the pixel space modes corresponding to the five systematics templates were projected out. The effect of mode

⁹In standard P(k) analyses, this issue is resolved by applying an integral constraint to the power spectrum estimates (see, e.g., Tegmark et al. 2002a).

projection on the estimates is illustrated in the bottom panels, showing the differences in the QML estimates. Hence, these values can be added to the estimates in the top panels to recover the results without mode projection. The change in the covariance of the estimates due to mode projection is negligible.

Reference mask

Our first mask, which is similar to that used in previous studies of RQCat (Giannantonio et al., 2006, 2008; Pullen and Hirata, 2013; Slosar et al., 2008; Xia et al., 2010), is mostly based on extinction, stellar density and seeing cuts, and also excises a few pixels with extreme values of airmass and sky brightness. When using this reference mask, the auto-spectrum estimates of the four RQCat subsamples exhibit significant excess power in the first multipole bin. In particular, the cross-correlation of the Low-*z* sample with the other samples confirm the presence of systematics in common. The cross-spectra of the quasar subsamples with the systematics templates, shown in Fig. 2.16, enable us to identify the main sources of contamination responsible for this excess power. In addition to seeing and airmass, which are the main contaminants in the four samples, stellar contamination affects the Low-*z* sample, and dust extinction and sky brightness contaminate the Mid-*z* and High-*z* samples.

The auto- and cross-spectra are marginally improved by projecting out the modes corresponding to the systematics templates, as shown by the small decrease in the χ^2 values, summarised in Tables 2.3 and 2.4. In particular, the large-scale power excess persists, confirming the conclusions by Pullen and Hirata (2013) that the contamination must involve non-linear combinations of systematics, or else systematics which have not been accounted for.

Improved masks

Our second mask is based on more restrictive cuts on the systematics, the most important of which are seeing and stellar density cuts. Using this mask not only improves the overall quality of the estimates, as measured by the χ^2 , but also eliminates the excess power at low ℓ in all subsamples except the Low-*z* one. Interestingly, the cross-spectra of the Low-*z* sample with the others exhibit no excess power. This indicates that the systematics responsible for the excess in the Low-*z* sample are successfully mitigated in the Mid-*z* and High-*z* samples, and thus in the Mid+High-*z* sample. The cross-spectra with the systematics templates are significantly decreased, although dust extinction and seeing still affect the Mid-*z* and High-*z* samples. Mode-projection further improves the quality

Table 2.3: The chi square values for the auto- and cross-power spectra of the four RQCat samples presented in Figs. 2.14 and 2.15, with and without mode projection (mp). The number of degrees of freedom is $\nu - p = 13$, and the probability to exceed (PTE) the observed chi squares are shown in parentheses. The results are shown in bold when PTE < 1% (corresponding to $\chi^2_{13} = 27.7$).

$\chi^2_{\nu-p}$	Mask 1		Mask 2		Mask 3	
	no mp	mp	no mp	mp	no mp	mp
Low-z	56.6 (2e-7)	44.4 (3e-5)	16.4 (0.23)	12.4 (0.49)	25.6 (0.02)	20.2 (0.09)
Mid-z	23.9 (0.03)	22.9 (0.04)	11.2 (0.59)	11.1 (0.60)	8.8 (0.78)	8.7 (0.79)
High-z	16.4 (0.22)	14.4 (0.34)	12.9 (0.45)	11.8 (0.54)	11.0 (0.61)	11.5 (0.57)
Mid+High-z	13.6 (0.40)	13.1 (0.44)	8.2 (0.83)	8.3 (0.82)	4.2 (0.99)	4.5 (0.98)
Low-z / Mid-z	33.2 (2e-4)	23.7 (5e-3)	27.7 (1e-4)	15.8 (0.07)	27.9 (1e-4)	15.3 (0.08)
Low-z / High-z	10.2 (0.68)	11.5 (0.24)	12.1 (0.52)	7.3 (0.60)	6.2 (0.94)	3.5 (0.94)
Low/MidHigh-z	26.3 (0.02)	19.9 (0.02)	17.9 (0.16)	9.7 (0.37)	13.9 (0.38)	8.3 (0.50)
Mid-z / $High-z$	6.1 (0.94)	4.3 (0.89)	11.0 (0.61)	3.9 (0.92)	9.2 (0.75)	2.7 (0.97)

of the estimates, but does not eliminate the excess power in the Low-z sample.

The third mask is based even more stringent cuts on the systematics. The cross-spectra with the systematics templates show that using this mask further decreases the influence of extinction, airmass and seeing on the Mid-*z* and High-*z* samples. However, it fails to remove the excess power in the Low-*z* sample, and dust extinction continues to impact the Mid-*z* sample. Yet, no statistical anomalies are observed in the auto-spectra of Mid-*z*, High-*z* and Mid+High-*z* samples. Mode-projection further improves the χ^2 values of all auto- and cross-spectra. In particular, the χ^2 values for the cross-spectra between the Mid-*z* and the High-*z* samples significantly improve, indicating a successful mitigation of the remaining levels of extinction.

In summary, when using the third mask and mode projection, the auto-spectra of the Mid-z, the High-z and Mid+High-z samples exhibit no evidence of systematics, and are well-fitted by the theory prediction. The cross-spectra of these samples with the Low-z sample are also not anomalous, indicating the absence of systematics in common. Mode-projection eliminates the contributions (to linear order) of the five systematics we have considered, so that any remaining contamination can only involve non-linear combinations or unidentified systematics. Their presence is confirmed in the Low-z subsample, but the other subsamples exhibit no evidence of contamination, indicating that any residual systematics are well within the sample variance.

2.5. Discussion

We have investigated the problem of estimating angular auto- and cross-power spectra on the



Figure 2.14: QML estimates of the (dimensionless) auto-power spectra of the overdensity maps of the four RQCat samples presented in Fig. 2.9. The estimates were calculated for the three masks in Fig. 2.12 (blue diamond, red triangles and black squares) for multipole bins of size $\Delta \ell = 11$, and compared with the theory prediction for a *Planck* cosmology (solid line). The top panel shows the estimates obtaining by mode-projecting the five systematics of Fig 2.11. The bottom panel indicates the difference in the estimates when mode projection is not used, i.e., $\hat{C}_b^{\text{no mp}} - \hat{C}_b^{\text{mp}}$, and can be added to the top panel to recover the estimates without mode projection (the change in the covariance of the estimates due to mode projection is negligible).

Table 2.4: The chi square values for the cross-power spectra of the four RQCat samples with the systematics templates, presented in Fig. 2.16, using the same conventions as Table 2.3.

$\chi^2_{\nu-p}$	Stellar density	Extinction	Airmass	Seeing	Sky brightness
	Masks 1 / 2 / 3	Masks 1 / 2 / 3	Masks 1 / 2 / 3	Masks 1 / 2 / 3	Masks 1 / 2 / 3
Low-z	17/11.2/15.5	10.5 / 9.1 / 14.3	23.9 / 8.2 / 6.7	50 / 25.2 / 12.5	11.3 / 4.1 / 6.9
Mid-z	4.0 / 9.5 / 15.3	12.3 / 17.7 / 13.1	26.9 / 16.1 / 8.2	19.9 / 11.1 / 11.3	8.4 / 6.4 / 7.7
High-z	3.3 / 5.7 / 6.5	28 / 26.9 / 47.1	5.8 / 3.0 / 6.8	32 / 20.1 / 12.1	27.9 / 18.6 / 15
Mid+Hig	h2.8 / 5.3 / 7.7	11.2 / 14.3 / 23	20.2 / 9.1 / 6.1	34.7 / 14.8 / 9.6	19.8 / 10.8 / 6.6



Figure 2.15: QML estimates of the (dimensionless) cross-power spectra of the RQCat overdensity maps using the same conventions as Fig. 2.14.



Figure 2.16: QML estimates for the (dimensionless) cross-power spectra of the four RQCat overdensity maps with the systematics templates using the same conventions as Fig. 2.14.

largest scales in the presence of systematics, and applied this framework to the UVX sources in the RQCat catalogue of SDSS photometric quasars. Previous studies (Giannantonio et al., 2014; Pullen and Hirata, 2013; Xia et al., 2011) indicated that this catalogue was not suitable for clustering analyses due to the high-levels of contamination by the systematics. We examined these conclusions by focusing on $0.5 < \tilde{z}_p \le 2.2$ objects, divided into four redshift bins, and attempted to remove the influence of the main systematics using improved sky masks and mode projection. We also improved the theoretical predictions by making use of refined redshift distributions estimated by cross-matching objects in RQCat with the SDSS-DR7 spectroscopic catalogue and applying robust completeness corrections.

In agreement with previous studies, we found that $\tilde{z}_p < 1.3$ objects exhibited significant levels of contamination by systematics, in particular dust absorption, airmass and seeing, which could not be eliminated by masking and mode projection. The remaining excess power on large scales points to the presence of unknown or non-linear combinations of the systematics in this redshift bin.

The large-scale excess power observed in the auto- and cross-spectra of $1.3 < \tilde{z}_p \le 2.2$ objects was eliminated by using improved sky masks based on templates of five of the main systematics. The agreement with the theory predictions further improved when projecting out the modes corresponding to these systematics. Within the statistical uncertainties, we found no evidence for remaining contamination in this sample. We conclude that photometric quasar samples can be made suitable for cosmological studies.

We did not attempt to model the contamination signal, but rather constructed a sample (through object selection and masking) which exhibited negligible levels of contamination. This approach relies on the ability to measure these levels of contamination, in our case through auto- and cross-spectra of the samples and cross-spectra with systematics templates, which are limited by the variance of the estimates due to the shot noise and sky coverage. Consequently, more contaminated samples (such as the whole RQCat) with higher number densities require a model of the systematics in order to obtain clustering measurements that are not dominated by spurious correlations. In this context, mode projection can be used to marginalise over the parameters of linear contamination models while estimating the power spectrum. However, future surveys may require more sophistic-ated algorithms, or even Bayesian component separation models similar to COMMANDER-RULER (Eriksen et al., 2006, 2008) used in the context of *Planck* (Planck Collaboration, 2013f).

The clustering properties of the large galaxy and quasar catalogues produced by next generation photometric surveys such as DES will put tight constraints on models of galaxy bias and PNG. However, stellar contamination and calibration errors will always be present in galaxy survey
data, and may compromise our understanding of the observables if not correctly treated. Thus, as sample variance becomes steadily smaller with increasing catalogue sizes, efficient strategies for mitigating the systematics, such as those presented here, will become critical for the cosmological interpretation of these surveys.

2.6. Aknowledgements

We thank Stephen Feeney, Filipe Abdalla and Ofer Lahav for useful discussions and comments. BL is supported by the Perren Fund and the IMPACT Fund. HVP is supported by STFC, the Leverhulme Trust, and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007- 2013) / ERC grant agreement no 306478-CosmicDawn. ABL is supported by the Leverhulme Trust and STFC. AP is supported by the Oxford Martin School. We acknowledge use of the following public software packages: HEALPIX (Górski et al., 2005); MANGLE (Hamilton and Tegmark, 2004; Swanson et al., 2008) and CAMB_SOURCES (Challinor and Lewis, 2011). We acknowledge use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science.

2.7. Appendices

2.7.1. Pixelisation, band-limit and smoothing issues

The equations defining the PCL and QML estimators do not impose maximum multipoles for the reconstruction, the coupling or the covariance matrices. Such bounds are called band-limits, and are in practice imposed by the finite information content of x, which depends on the resolution of the map and any additional operations such as smoothing. In particular, pixelising the signal x into a (full or cut-sky) map x induces a distortion in the power spectrum estimates, which we parametrise as

$$\mathcal{C}_{\ell}^{\mathrm{pix}} = b_{\ell}^2 \mathcal{C}_{\ell},\tag{2.33}$$

where C_{ℓ}^{pix} and C_{ℓ} are the power spectra of x and x respectively. The continuous map x is usually not accessible, but can be approximated using a high-resolution pixelisation, which in practice will correspond to the highest resolution at which the data are available (and will only require a small correction of the pixelisation-induced bias as detailed below). x will then refer to a smoothed, lower-resolution map constructed from x, which will be used to estimate the power spectrum on a specific range of scales. This approach is motivated by the complexity of the PCL and QML pixel-space estimators which depends on the number of pixels in the mask. It is usually desirable to use the lowest resolution for which the power spectrum can be accurately estimated in the range of multipoles of interest.

The beam b_ℓ is decomposed into a pixelisation-induced part $b_\ell^{\rm pix}$ and a smoothing part $b_\ell^{\rm pix}$. In the previous sections we have assumed that pixelisation distortions were negligible, i.e., $b_{\ell}^{\text{pix}} \approx 1$ for all ℓ , but this assertion is true for a mode ℓ only if the pixels of x are small compared with $180/\ell$ degrees. We considered the HEALPIX pixelisation of the sphere, where a map at resolution $N_{\rm side}$ has $12N_{\rm side}^2$ equal-area pixels. In terms of its effect on the power spectrum, pixelising a signal at resolution $\mathrm{N}_{\mathrm{side}}$ is well approximated by smoothing with a Gaussian kernel of full width at half maximum (FWHM) $41.7/N_{side}$ degrees. This beam, shown in Fig. 2.17, smoothly decays as ℓ increases, and imposes an effective band-limit of $\ell_{\rm max} = 7 N_{\rm side}$ on the pixelised map. However, it is well-known that the accessible multipoles for a map at HEALPIX resolution N_{side} lie within $\ell \in [0, 2N_{side}]$. Higher multipoles are not accessible because the integrals in the spherical harmonics and Legendre transforms are not accurately approximated by matrix multiplications at this resolution. It is essential to smooth the initial map before degrading it in order to avoid a mismatch between the band-limits, while insuring that the power spectrum can be reconstructed up to $\ell_{\rm max} = 2 N_{\rm side}$. We investigated this issue and found that band-limiting the map at $\ell_{\rm max} = 4 N_{\rm side}$ gives optimal performance. This can be realised by a Gaussian smoothing of FWHM $60/N_{\rm side}$ degrees, as illustrated in Fig. 2.17.

However, the smoothing procedure introduces information from outside the mask into the map, hence biasing the estimates (Aurich and Lustig, 2011; Copi et al., 2011; Feeney et al., 2011b). This effect depends on the resolution of the map and the shape of the mask, and can significantly affect the power spectrum on a wide range of scales. To avoid this smoothing-induced contamination, the mask must be extended. A systematic procedure is to smooth the complement of the mask, i.e., 1 - m, and keep the pixels below a certain threshold (0.01 in this work). But because the Gaussian kernel is band-limited in harmonic space but not compact in pixel space, the smoothing will always leak contamination signal into the data. Also, depending on the threshold, the mask extension might



Figure 2.17: Beams that apply to the power spectrum estimates of pixelised, smoothed maps at HEALPIX resolution $N_{\rm side}$. The beam must be accounted for in the construction of the model covariance matrices, and inverted when comparing the estimates with theory.

be large and thus significantly increase the variance of the estimates. This can be resolved by using a top-hat smoothing of diameter $90/N_{side}$ degrees, which band-limits the data at $\ell_{max} = 4N_{side}$ similar to the Gaussian smoothing. The top-hat kernel is compact and the mask extension is smaller and requires no threshold. However its extent in harmonic space is infinite (in other words it is not band-limited), as shown in Fig. 2.17, which can introduce further contamination in the map due to approximated smoothing. This can be avoided by performing the top-hat smoothing at high-resolution with a large band-limit. Alternatively, the smoothing can easily be performed in pixel space through explicit convolution since the kernel has a simple shape and the procedure can be parallelised.

In conclusion, the power spectrum of a high-resolution signal can be estimated from a map at lower resolution N_{side} provided that the band-limits, the pixelisation and smoothing-induced biases are correctly handled. Before degradation, the initial map is smoothed with a Gaussian kernel of FWHM 60/N_{side} degrees or a top-hat kernel of diameter 90/N_{side} degrees, and the mask must be extended accordingly to minimise the contamination due to smoothing. The estimation can then be performed on the low resolution map in the range $\ell \in [0, 4N_{side}]$ and the signal covariance matrix created with a prior $b_{\ell}^2 C_{\ell}$ up to $\ell_{max} = 4N_{side}$ in order to incorporate all the information in the data. Dividing by the beam b_{ℓ}^2 leads to accurate unbiased band-power estimates of C_{ℓ} in the range $\ell \in [0, 2N_{side}]$.

2.7.2. Karhunen-Loève compression demystified

The Karhunen-Loève (KL) transform (Tegmark et al., 1997, 1998, 2002b; Vogeley and Szalay, 1996) refers to finding a basis in which the transformed data pixels are statistically orthogonal with respect to a prior on the power spectrum. The transformation matrix **B** satisfies

$$\tilde{\mathbf{x}} = \mathbf{B}\tilde{\mathbf{y}}.\tag{2.34}$$

such that

$$\langle \tilde{y}_n \tilde{y}_{n'}^* \rangle = \lambda_n \delta_{nn'}, \tag{2.35}$$

and the columns \mathbf{B}_n are the eigen-vectors of the pixel-pixel covariance matrix, i.e.,

$$\tilde{\mathbf{C}}\mathbf{B}_n = \lambda_n \mathbf{B}_n. \tag{2.36}$$

Hence $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ retain the same information and KL compression is equivalent to a principal component analysis performed through singular value decomposition (SVD) of the pixel-pixel covariance matrix. However, this approach is only optimal if the noise covariance is diagonal. Anisotropic noise requires a so-called prewhitening operation, changing the eigen-problem to be solved into

$$\tilde{\mathbf{S}}\mathbf{B}_n = \lambda_n \tilde{\mathbf{N}}\mathbf{B}_n,\tag{2.37}$$

where the eigen-vectors now diagonalise both the signal and the noise covariance matrices. If the transformation is renormalised such that $\mathbf{B}_{n}^{\dagger}\mathbf{N}\mathbf{B}_{n} = 1$ before transforming the map x, the coefficients λ_{n} can be interpreted as signal-to-noise ratios, i.e.,

$$\langle \tilde{y}_n \tilde{y}_{n'}^* \rangle = \delta_{nn'} (1 + \lambda_n). \tag{2.38}$$

The power spectrum is invariant under rotation in the isotropic case, it can be estimated from \mathbf{y} provided that the matrices $\tilde{\mathbf{C}}$ and \mathbf{P}^{ℓ} are in the same coordinate system. In the KL basis the most informative contributions to the spectrum explicitly come from the first KL modes, as they correspond to high SNR modes. Hence KL compression is often used to remove the noisiest modes and speed up power spectrum estimation, in particular for the matrix inversion in the QML estimator. However, this gain is usually offset by the cost of the preliminary SVD required to compute the KL transformation matrix.

Moreover, although removing the noisiest modes leaves the PCL and QML estimates unbiased, it increases their variance and potentially impacts the estimates on a wide range of scales. In particular,



Figure 2.18: Histograms of the χ^2 values of the PCL (left) and QML (right) estimates of the 600 CMASS mock catalogues. We used multipole bins of size $\Delta \ell = 11$, and the theory prediction from CAMB SOURCES. The boxes show the Poisson errors due to sample variance. The histograms were normalised and compared with the theoretical χ^2 probability distributions (thick lines). Their good agreement demonstrates the validity of the χ^2 given by Eq. 2.39.

assuming a constant diagonal noise, the KL modes calculated from a full sky covariance matrix **S** directly relate to the theory spectrum. The largest modes correspond to the largest C_{ℓ} 's with multiplicity $2\ell + 1$. On the cut sky, this degeneracy is broken and each mode uniquely relates to a linear combination of C_{ℓ} 's with finite support peaking at the previous full sky value (thus conserving the order of the modes). Small masks increase the scope of this combination, and removing small KL modes can thus impact a wide range of multipoles. Logically, considering the CMB spectrum with a typical CMB mask shows that the lowest KL modes, the noisiest, relate to high ℓ multipoles, and removing them leaves the low ℓ power spectrum estimates unchanged. However for a galaxy survey the theory spectrum is flatter and the mask larger, and the noisiest modes then relate to a wide range of multipoles. Therefore, removing them also impacts the largest scales. Fortunately this effect is reduced when performing the estimation in bins, and the KL compression has a negligible impact on the quality and the variance estimates.

We did not use KL compression in the context of our analyses because the gain in computer time is marginal. We were able to run the QML estimator at $N_{\rm side} = 64$ without any approximation.

2.7.3. χ^2 for band-power estimates

To compare power spectrum measurements with theory predictions we used a χ^2 defined as

$$\chi^2 = \sum_{bb'} (\hat{C}_b - C_b) \mathbf{V}_{bb'}^{-1} (\hat{C}_{b'} - C_{b'})$$
(2.39)

where \hat{C}_b denotes the (PCL or QML) band-power estimates and $C_{b'}$ the theory band-powers, constructed from the theory power spectrum C_ℓ (calculated from CAMB_SOURCES) using the window functions $W_{b\ell}$ defined in Eqs. (2.17) and (2.18). The covariance matrix V is calculated using Eq. (2.11) and is equal to the inverse of the Fisher matrix when the QML estimator is used. Note that the shot noise $\frac{1}{G}$ must be subtracted from the auto-spectrum estimates \hat{C}_b .

It is well known that observed power spectra C_{ℓ} as calculated by Eq. (3.1) are only described by Gaussian statistics at high- ℓ when the central limit theorem applies. For low- ℓ estimates, one must resort to alternative likelihood functions (see, e.g., Hamimeche and Lewis 2009; Jaffe et al. 1999). However, when estimating band-powers rather than individual multipoles, these effects can be neglected, and a Gaussian likelihood can be used to compare the estimates with theory band-powers. To illustrate this point, Fig. 2.18 shows the χ^2 values for the PCL and QML estimates of the 600 CMASS mock catalogues with bin size $\Delta \ell = 11$ in the range [2,120], corresponding to 11 degrees of freedom. The normalised histograms follow the theoretical χ^2 distributions, demonstrating the validity of a Gaussian likelihood, as expected when using band-powers.

3

Exploiting the full potential of photometric quasar surveys: optimal power spectra through blind mitigation of systematics

"If tortured sufficiently, data will confess to almost anything". F. Menger

3.1. Abstract

We present optimal measurements of the angular power spectrum of the XDQSOz catalogue of photometric quasars from the Sloan Digital Sky Survey. These measurements rely on a quadratic maximum likelihood estimator that simultaneously measures the auto- and cross-power spectra of four redshift samples, and provides minimum-variance, unbiased estimates even at the largest angular scales. Since photometric quasars are known to be strongly affected by systematics such as spatially-varying depth and stellar contamination, we introduce a new framework of *extended mode projection* to robustly mitigate the impact of systematics on the power spectrum measurements. This technique involves constructing template maps of potential systematics, decorrelating them on the sky, and projecting out modes which are significantly correlated with the data. Our method is able to simultaneously process several thousands of nonlinearly-correlated systematics, and mode projection is performed in a blind fashion. Using our final power spectrum measurements, we find a good agreement with theoretical predictions, and no evidence for further contamination by systematics. Extended mode projection not only obviates the need for aggressive sky and quality cuts, but also provides control over the level of systematics in the measurements, enabling the

search for small signals of new physics while avoiding confirmation bias.

3.2. Introduction

Quasars are bright, highly biased tracers of the large scale structure (LSS) of the universe, and are useful for testing cosmological models in large volumes and over extended redshift ranges. In particular, their bias can inform us about the abundance and mass of dark matter halos in which they form, opening new windows on galaxy formation and the astrophysics of active galactic nuclei (e.g., Fan et al. 2006). In addition, they can be used to constrain primordial non-Gaussianity (PNG) which is predicted to enhance the bias of LSS tracers on large scales (Dalal et al., 2008; Komatsu et al., 2009; LoVerde et al., 2008; Matarrese and Verde, 2008). However, these applications require accurate auto- and cross-correlation power spectrum measurements, which are complicated by the presence of significant systematics in the data, and the difficulty of creating large, deep quasar catalogues. Indeed, although quasar candidates are easily confirmed with spectroscopy, quasars are point sources, and most point sources in the sky are stars. Since acquiring high resolution spectra is a time- and resource-consuming process, the creation of large quasar catalogues is cumbersome and can only be realised by preselecting targets and scheduling them for spectroscopic follow-up. In optical frequencies, catalogues of photometric quasars were constructed from the Sloan Digital Sky Survey (SDSS, Gunn et al. 2006), and promising candidates were followed up using the SDSS spectrograph, yielding large catalogues of confirmed quasars such as the Baryon acoustic OScillations Survey (BOSS, Dawson et al. 2013).

While spectroscopic catalogues don't suffer from stellar contamination, photometric samples are larger and extend to fainter magnitudes, and can therefore yield more precise measurements of the clustering and the bias of quasars. However, photometric data are significantly contaminated by multiple sources of systematics, either intrinsic (*e.g.*, dust extinction), observational (*e.g.*, seeing, airmass), or instrumental (*e.g.*, instrument calibration), and star-quasar separation using only colour information is also nontrivial. These systematics affect the properties of the raw images in complex ways, propagate into the final quasar catalogues, and create spurious spatial correlations (spatially-varying star-quasar separation efficiency will induce a spurious clustering signal, see *e.g.*, Huterer et al. 2013). Finally, some of these correlations may also be imprinted in spectroscopic catalogues, since the latter rely on targets selected from photometric quasars. Therefore, not

just precise — but also accurate — cosmological inferences from quasar surveys require careful mitigation of systematics.

The first studies of the clustering of quasars in optical frequencies used both spectroscopic (e.g., Outram et al. 2003; Ross et al. 2009; Shen et al. 2007) and photometric (e.g., Myers et al. 2006) catalogues from early SDSS data, and were used to constrain numerous cosmological and astrophysical quantities of interest, such as the quasar bias, PNG, and the quasar luminosity function (Myers et al., 2007a; Richards et al., 2006; Serber et al., 2006; Slosar et al., 2008; Strand et al., 2008). They exhibited power excesses on large and small scales (Myers et al., 2007b), which were even more significant in the DR6 photometric quasar catalogue (Richards et al., 2009). Although cuts to the data were not sufficient to remove this excess power, thus pointing to PNG (Giannantonio et al., 2012; Karagiannis et al., 2013; Xia et al., 2009, 2011), recent work has demonstrated that the excess power was due to systematics (Pullen and Hirata, 2013) and could be eliminated using mode projection (Leistedt et al., 2013). Indeed, it is known that spatially-varying stellar contamination can combine with other systematics and generate such excess clustering power, mimicking significant levels of PNG (Giannantonio et al., 2014). Alternatively, other studies have focused on using the cross-correlations of the photometric quasars with other data (Giannantonio and Percival, 2014), thus extracting some of the information they contain while avoiding the main systematics. The resulting PNG constraints were competitive with those obtained using normal galaxies (e.g., Ross et al. 2013). The eighth data release of SDSS yielded a new catalogue of photometric quasars, XDQSO (Bovy et al., 2011a), relying on the extreme deconvolution technique (Bovy et al., 2011b); the latest spectroscopic data were used to provide an improved probabilistic selection of quasars, with greater control over completeness issues. XDQSO extends to fainter magnitudes, and was primarily used to select high redshift candidates scheduled for spectroscopic follow-up in the context of BOSS (Ross et al., 2012c). Its extension, coined XDQSOz (Bovy et al., 2012), provides estimates and probability density distributions for the photometric redshifts of the catalogued objects. XDQSOz was cross-correlated with the cosmic microwave background (CMB) lensing map from the Atacama Cosmology Telescope (Sievers et al., 2013) to constrain the quasar bias (Sherwin et al., 2012). Constraints on PNG were also derived from clustering measurements (Agarwal et al., 2014b; Ho et al., 2013), but required significant cuts and corrections to exploit measured power spectra, and corrected for systematics using methodologies introduced in Agarwal et al. (2014a); Ho et al. (2012); Ross et al. (2011, 2012a).

Quasar clustering studies therefore remain suboptimal and limited due to cuts and corrections needed to address the high levels of spurious correlations created by systematics. Nevertheless, most

of the potential systematics are actually known and can be mapped onto the sky. It is their complex impact on the data which is not known, and prevents the modelling of spurious correlations when estimating power spectra. Mode projection, however, can mitigate the impact of systematics in a robust manner, while minimising the need to throw out hard-won data through masking and cuts (Leistedt et al., 2013; Pullen and Hirata, 2013; Slosar et al., 2004; Tegmark et al., 1998). The purpose of this work is to extend the standard mode projection approach by designing a generic methodology to mitigate the impact of large numbers of known systematics in a flexible and robust manner. We apply this technique to the XDQSOz catalogue in order to precisely control the level of the contamination and accurately measure the power spectrum on the largest scales, which is essential for constraining PNG. Next generation photometric surveys, such as the Dark Energy Survey¹ (DES), will reach unprecedented precision and will require careful treatment of systematics. This *extended mode projection* approach will enable the full potential of such surveys to be exploited in the search for new physics.

This article is organised as follows. In Sec. 2 we recall the definitions and properties of quadratic power spectrum estimators, and introduce the extended mode projection technique to mitigate systematics when estimating power spectra. In Sec. 3 we turn to the XDQSOz catalogue of photometric quasars. We present our data samples, theory predictions and power spectrum measurements, and discuss the ability of the extended mode projection, in combination with blind analysis techniques, to mitigate systematics. The discussion and conclusions are presented in Sec. 4.

3.3. Theory and methods

3.3.1. Power spectrum estimation

The statistics of a Gaussian random field on the sphere are entirely characterised by its power spectrum C_{ℓ} . Any realisation of this field, denoted by x, has an observed power spectrum C_{ℓ} , defined

¹www.darkenergysurvey.org

as the variance of the spherical harmonic coefficients $a_{\ell m}$,

$$C_{\ell} = \sum_{m=-\ell}^{\ell} \frac{|a_{\ell m}|^2}{2\ell + 1},$$
(3.1)

which, on the full sky only, is an unbiased estimator of C_{ℓ} (i.e., over many realisations, $\langle C_{\ell} \rangle = C_{\ell}$) with cosmic variance

$$\operatorname{Var}(\mathcal{C}_{\ell}) = \frac{2C_{\ell}^2}{2\ell + 1}.$$
(3.2)

Cosmological models usually predict theory C_{ℓ} s which can then be confronted with observed C_{ℓ} s to constrain model parameters. However, when measuring C_{ℓ} with real data, several issues arise. Most datasets are only defined on a portion of the sky, and C_{ℓ} must then be estimated from a cut-sky – or masked – data vector x. This issue has been extensively studied in the context of the CMB (*e.g.*, Efstathiou 2004a,b, 2006; Gruetjen and Shellard 2012; Knox et al. 1998; Pontzen and Peiris 2010; Tegmark 1997), and also more specifically for LSS data (*e.g.*, Ho et al. 2012; Leistedt et al. 2013; Tegmark et al. 2002a). The two main quadratic angular power spectrum estimators are the pseudospectrum (PCL) and quadratic maximum likelihood (QML) estimators. In both cases, band-power estimates \hat{C}_{b} (*i.e.*, estimates of C_{ℓ} in multipole bins $\Delta \ell_{b}$) are quadratic in the data,

$$\hat{C}_b = \mathbf{x}^t \mathbf{E}^b \mathbf{x},\tag{3.3}$$

blwhere \mathbf{E}^{b} characterises the estimator under consideration. The covariance of the estimates, denoted by $\mathbf{V}_{bb'} = \text{Cov}(\hat{C}_{b}, \hat{C}_{b'})$, includes a cut-sky induced variance which is larger than the cosmic variance and specific to the estimator at hand, and a contribution from the noise and systematics in the data. The equations and implementation details we use are detailed in Leistedt et al. (2013), and in this paper we only provide a summary of the main characteristics of these estimators.

The PCL estimator deconvolves the mask-induced mode-coupling to produce unbiased estimates of C_{ℓ} . The estimator is straightforward to implement, not very computationally intensive, but only optimal (*i.e.*, it provides minimum-variance estimates) when C_{ℓ} is close to flat, and if the mask has a simple geometry. The QML estimator is unbiased in C_{ℓ} and is a minimum variance estimator, but is computationally intensive and requires priors on the power spectrum, noise and potential systematic uncertainties in the data. These priors go in a pixel-pixel covariance **C** modelled as the superposition of a signal part S calculated using a theory prior $\{C_{\ell}\}$ and noise N, *i.e.*,

$$\mathbf{C} = \langle \mathbf{x}\mathbf{x}^t \rangle = \mathbf{S} + \mathbf{N} = \sum_{\ell} C_{\ell} \mathbf{P}^{\ell} + \mathbf{N},$$
(3.4)

where $(\mathbf{P}^{\ell})_{ij} = (2\ell + 1)/4\pi P_{\ell}(\mathbf{n}_i \cdot \mathbf{n}_j)$ is a useful matrix notation (Tegmark, 1997) in which \mathbf{n}_i denotes the angular position of the *i*th pixel. In fact, it can be shown that PCL is equivalent to QML when assuming uncorrelated pixels, *i.e.*, using a flat power spectrum as prior. Further complications arise because power spectra of real data can only be accurately estimated in a multipole range $0 < \ell < \ell_{\text{max}}$, where ℓ_{max} mainly depends on the noise level and the resolution at which the pixelised maps are analysed. These parameters must be consistently adjusted when constructing priors and computing the power spectrum estimates. Rules that guarantee the correct implementation of the PCL and QML estimators for LSS studies are detailed in Leistedt et al. (2013).

3.3.2. Extended mode projection

As mentioned in the introduction, real data are often contaminated by systematic effects which create spurious correlations in the measured power spectra. A first order solution is to construct masks to ignore the regions of the sky which are not reliable, for example by rejecting pixels based on quality cuts, or where potential systematics are significant (*e.g.*, dust, seeing). However, this approach does not remove the spurious correlations due to spatially-varying depth or stellar contamination *in the unmasked region*. Hence, masking is insufficient when these systematics are not negligible compared to noise and the cosmic variance. In addition, several *ad hoc* choices are required to define what one means by unreliable regions, decide which systematics to include, and what cuts to apply in order to construct the sky mask. As a consequence, there is a risk of being subject to confirmation bias, *i.e.*, one may perform the analysis using an initial mask, then notice it is insufficient to remove the impact of systematics, and therefore perform another analysis with a more stringent mask, until the results are in accordance with expectation.

To illustrate more robust solutions to mitigate systematics, let us start with a toy model where the observed data y is the superposition of the systematics-free data vector x (of covariance C) and a set of N_{sys} systematics maps t_i , *i.e.*,

$$\mathbf{y} = \mathbf{x} + \sum_{i=1}^{N_{sys}} \alpha_i \mathbf{t}_i.$$
(3.5)

With a method to estimate and fix the α_i parameters at hand, one could correct the maps or the power spectra to remove the impact of the systematics (Ho et al., 2012; Ross et al., 2011, 2012a). However, if the systematics are not well-described by a linear model, the correction may be not be robust, and create biases in the measured power spectra.

Alternatively, one could sample the coefficients α_i , estimate the power spectra of the corrected data vector $\mathbf{y} - \sum \alpha_i \mathbf{t}_i$, and finally perform a Bayesian Monte Carlo marginalisation over the α_i s. This approach can be seen as an analytic marginalisation of systematics to linear order, and is more robust than attempting to correct the data or the power spectra. Interestingly, this approach can be implemented in the QML estimator to directly obtain power spectrum estimates for the systematics-free map \mathbf{x} . This *mode projection (e.g., Leistedt et al. 2013; Pullen and Hirata 2013; Slosar et al. 2004; Tegmark et al. 1998)* is implemented by using a prior covariance matrix

$$\mathbf{D} = \mathbf{C} + \lim_{\epsilon \to \infty} \sum_{i} \epsilon \mathbf{t}_{i} \mathbf{t}_{i}^{t},$$
(3.6)

and setting ϵ to large values, *i.e.*, by giving a large variance to the templates t_i . Hence, the spatial modes corresponding to these templates² are considered as noise and ignored when estimating the power spectra.

To gain intuition into the mode projection process and find its limitations, let us first notice that the results of the QML estimator are invariant under rotation, *i.e.*, transforming the data and the prior covariance matrix **D** using an arbitrary rotation matrix does not affect the band-power estimates (the rotation is a linear transformation of the data pixels and conserves the information content). Therefore, it is straightforward to show that using mode projection with only one template is equivalent to masking one pixel of the data vector **x** in an arbitrary rotated frame (defined by the template). With multiple templates, this masking interpretation is only valid if the templates are orthogonal on the patch of sky of interest, *i.e.*, if $\mathbf{t}_i^t \mathbf{t}_j = \delta_{ij}$. In this case only, projecting out N_{sys} templates is equivalent to masking N_{sys} pixels in some rotated frame. Hence, the increase in variance in the band-power estimates scales as the number of templates. However, real systematics are often correlated, and projecting out N_{sys} systematics may correspond to masking less than N_{sys} pixels, leaving no control over the effective increase in variance in $\mathbf{V}_{bb'}$. Moreover, without physically-motivated models for the contamination, one may want to include large numbers of templates and consider more generic data-driven models. Yet, in this case many modes may in fact not be relevant in treating systematics, and only contribute to increasing the variance of the band-

²In this paper, we will use the terms templates and spatial modes interchangeably.

power estimates. The standard mode projection framework is not appropriate in these situations, because it leaves no control over the increase in variance or the accepted level of contamination when projecting out numerous correlated modes.

To address these issues, we extend the mode projection framework and add two operations to it, prior to power spectrum estimation. Firstly, the systematics templates are decorrelated on the patch of sky of interest. This eliminates the redundant information and leads to a minimal representation of the input templates. Secondly, the resulting orthogonal modes are cross-correlated with the data, yielding null tests that can be used to select and only project out the modes which significantly contaminate the data. This selection can be done according to any criteria of interest, such as cuts based on reduced χ^2 for null tests, which is the metric we use in this work. Together, these two extensions provide control over the effective number of modes and the increase in variance due to mode projection. We now detail our implementation of *extended mode projection* in the context of galaxy and quasar survey analysis.

For the first step, we stack the systematics templates to construct a matrix \mathbf{T} where the rows contain the modes to be projected out. As detailed in the next sections about SDSS quasars, we will in fact use modes \mathbf{t}_i which are either base systematics templates, or products of templates, so that Eq. (3.5) will be extended to non-linear contamination models. The templates \mathbf{t}_i are then decorrelated using standard singular value decomposition (SVD) techniques to find a transformation matrix \mathbf{W} which, once applied to \mathbf{T} , yields uncorrelated templates contained in the rows of a matrix \mathbf{U} . This procedure is encapsulated in the set of equations³

$$\mathbf{T}^{t}\mathbf{T} = \mathbf{W}\mathbf{\Lambda}\mathbf{W}^{t} \tag{3.7}$$

$$\mathbf{U} = \mathbf{T}\mathbf{W} \tag{3.8}$$

$$\mathbf{U}^{t}\mathbf{U} = \mathbf{W}^{t}\mathbf{T}^{t}\mathbf{T}\mathbf{W} = \mathbf{\Lambda}, \qquad (3.9)$$

where Λ is a diagonal matrix containing the eigenvalues λ_i of \mathbf{T} . The advantage of this decorrelation step is twofold. Firstly, the output templates, denoted by \mathbf{u}_i , are orthogonal ($\mathbf{u}_i^t \mathbf{u}_j = \lambda_i \delta_{ij}$), and yield a minimal description of the input systematics. Secondly, the SVD provides the number and the significance of the orthogonal modes in this minimal description. In particular, it may find modes which are consistent with numerical noise⁴, and can safely be ignored since they are redundant or

 $^{^{3}}$ Note that one may modify these equations to incorporate a covariance matrix and use the Sherman-Morrison formula to analytically perform the mode projection. In this work, we have employed Eq. (3.6) instead, to avoid the need to reweight all templates using a large covariance matrix, which is computationally very intensive.

⁴A mode is considered as numerical noise if the ratio between its eigenvalue and the largest eigenvalue of the system is smaller than numerical precision.

not significant in the input system.

The second step aims to select which of these orthogonal modes u_i should be projected out. In this work, we resort to C_{ℓ} estimators to measure cross-correlation power spectra between the data and the systematics templates. In principle, any measure of correlation could be used, but resorting to power spectrum estimators captures scale information, which is desirable since the χ^2 will be used to project out modes for the auto-correlation power spectrum. On average, LSS data such as quasar and galaxy surveys should not be correlated with observational systematics, e.g., dust extinction, seeing variation or stellar density. Individual realisations may exhibit non-zero cross-spectra, but, in the absence of systematics, these should be consistent with zero within the statistical uncertainties resulting from the estimator and the noise. Therefore, the cross-spectra can be used for null tests, yielding one reduced χ^2 per systematics mode, to decide if the mode should be projected out. Large χ^2 indicates significant spurious correlation between the data and the mode, pointing to contamination by systematics. For a given set of uncorrelated systematics, one can impose a global χ^2_{cut} and only project out the modes with null test χ^2 above this value. The $\chi^2_{\rm cut}$ can be fixed according to absolute criteria, for example $\chi^2_{\rm cut} = 1.0$, which ensures that the remaining (i.e., non-projected) modes pass the null tests and are consistent with zero within the statistical uncertainties. Alternatively, one can also adjust $\chi^2_{
m cut}$ to control the number of projected modes — and thus, the increase in variance — in the estimated power spectra. The next section provides more details about the power spectrum estimator used for the null tests.

In conclusion, extended mode projection only relies on the ability to create systematics templates, decorrelate them, and cross-correlate them with the data of interest. It allows one to consider generic contamination models with large numbers and complex combinations of templates, from which the main orthogonal modes will be extracted. Under the reasonable assumption that these systematics — and thus the orthogonal modes — do not correlate on average with the data, one can efficiently find the modes which need to be projected out. Compared to other techniques previously applied to galaxy and quasar survey data (Agarwal et al., 2014b; Ho et al., 2013; Ross et al., 2011, 2012a), extended mode projection does not attempt to directly correct the data or the power spectra but rather, analytically marginalises over the systematics when computing band-power estimates. Beyond the parameters of the cross-power spectrum estimator (*e.g.*, band-power width), which are guided by the sky coverage and noise level, the only tunable parameters in this technique are the set of input systematics templates and χ^2_{cut} , which corresponds to the accepted level of correlation between a potential systematic and the data. Therefore, this framework is based on the principles of blind analysis, and can be used in a variety of situations where the data are non-linearly

contaminated by several systematics.

3.3.3. Fast cross-power spectra and null-tests

With our implementation of extended mode projection, the systematics modes to be projected out are selected based on null tests resulting from cross-power spectra with the data of interest. Due to the potentially large number of modes, it is essential to use a fast estimator to compute these cross-power spectra. The PCL estimator is a good candidate, but may be suboptimal when using complex masks, or if the cross spectra are significantly non-flat, which is likely when considering significantly contaminated data such as quasar surveys (Leistedt et al., 2013). Yet, using the QML estimator would require significant time and memory resources, despite these power spectra only being used for null tests, and thus not requiring the same level of accuracy as the main C_{ℓ} analysis. Therefore, we designed an approximate QML estimator which yields fast null tests which are quasi-optimal when compared with the full QML estimator. This new estimator simply replaces the Fisher matrix in $\mathbf{E}_{\text{QML}}^{b}$ by the inverse covariance matrix of the PCL estimator. More precisely, we compute cross spectra between one data sample \mathbf{x}_{map} and a systematics mode \mathbf{x}_{sys} using

$$\hat{C}_{b}^{\text{nulltest}} = \sum_{b'} \mathbf{G}_{bb'}^{-1} \frac{1}{2} [\mathbf{x}_{\text{map}}^{t} \mathbf{C}_{\text{map}}^{-1} \mathbf{P}^{b'} \mathbf{C}_{\text{sys}}^{-1} \mathbf{x}_{\text{sys}}], \qquad (3.10)$$

where $\mathbf{G}_{b_1b_2}$ is the covariance of the PCL estimator, which can be simplified to

$$\mathbf{G}_{b_1b_2} = \sum_{b_1b_2b_3b_4} \mathbf{M}_{b_1b_3}^{-1} \mathbf{M}_{b_2b_4}^{-1} C_{b_3}^{\mathrm{map}} C_{b_4}^{\mathrm{sys}} 4 \mathbf{\Pi}_{b_1b_2b_3b_4}, \qquad (3.11)$$

where M is the PCL coupling matrix. Here, $\Pi_{b_1b_2b_3b_4} = \text{Tr} \left[\mathbf{P}^{b_1}\mathbf{P}^{b_2}\mathbf{P}^{b_3}\mathbf{P}^{b_4}\right]$ is computationally intensive to evaluate but only depends on the mask under consideration, and can therefore be precalculated. In this formulation, C_b^{map} and C_b^{sys} are priors for the band-powers of the data and the systematics mode to be cross correlated, also used to construct the covariance matrices \mathbf{C}_{map} and \mathbf{C}_{sys} . In particular, C_b^{map} is the theoretical prior also used in the main optimal estimator, whereas we obtain C_b^{sys} using a fast PCL estimator applied to \mathbf{x}_{sys} . As a consequence, the approximate power spectrum estimator of Eq. (3.10) is fast and can be used to efficiently calculate the numerous cross-power spectra between the data and the systematics. In this work, we use null test χ^2 s which are simply calculated using a Gaussian likelihood comparing the $\hat{C}_b^{\text{nulltest}}$ with zero, as in



Figure 3.1: Redshift distributions of the photometric quasar samples selected with $P_{QSO} > 0.8$ and top hat windows on the photometric redshift estimates \hat{z}_p , highlighted by the grey regions. The black curves show redshift distribution estimates n(z) of the underlying quasar distributions, calculated as the sum of the posterior distributions on the individual redshift estimates. The dashed line shows a fit with Gaussian kernels, used for calculating the theoretical C_{ℓ} 's. The red and blue dashed curves are redshift histograms of the objects cross-matched to the SDSS DR7 and BOSS DR9 spectroscopic quasar catalogues. These are only used for qualitative purposes, since they provide estimates of n(z) which are biased towards low- and high-redshift, respectively, due to the magnitude and redshift selection effects (details in the text).

Leistedt et al. (2013). In a later section, we compare the χ^2 values obtained by the optimal and the approximate power spectrum estimators, and show that the latter is sufficiently accurate for the masks, data and systematics under consideration.

3.4. Application to SDSS photometric quasars

We now apply the extended mode projection framework to photometric quasars from the SDSS, and compare the resulting power spectrum measurements with theoretical predictions. We discuss these measurements, their suitability for cosmological inference, and the ability of extended mode projection to eliminate spurious correlations due to spatially-varying systematics.

3.4.1. Sample selection and redshift distributions

We consider the XDQSOz catalogue of photometric quasars (Bovy et al., 2012), which contains quasar candidates selected from the set of point sources of the SDSS DR8 imaging data, covering ~ 10^4 deg^2 of the southern and northern Galactic sky. For each object, XDQSOz includes the probability of being a quasar or a star, calculated from the observed magnitudes (*ugriz*) and their estimated errors, using a data-driven model of the density distribution of quasars in *ugriz* flux-redshift space. This model was constructed by applying the *extreme deconvolution* algorithm (Bovy et al., 2011b) to the spectroscopically-confirmed quasars in the SDSS DR7 quasar catalogue (Schneider et al., 2010). XDQSOz also includes photometric redshift estimates \hat{z}_p , defined as the highest peak in the posterior distribution of the individual photo-*z* estimates. Although the latter distributions are not directly released in the public version of XDQSOz, they can be straightforwardly recalculated using the flux-redshift space model and the publicly-available code.

We consider all objects with $P_{QSO} > 0.8$, but with no further cuts to the data, other than the sky masks described in the next section. We separate this catalogue into four samples by selecting objects with photometric redshifts \hat{z}_p in top-hat windows ranging [0.5, 1.35], [1.35, 1.7], [1.7, 2.2] and [2.2, 3.5], which have comparable numbers of objects, and thus similar shot noise. We reject the low-redshift objects ($\tilde{z}_p < 0.5$); this redshift range suffers from nontrivial incompleteness issues due to missed low-*z* quasars, which are resolved and therefore not processed by the point source star-quasar classifier.

In order to calculate theoretical predictions for the angular power spectra of the four samples, we follow Leistedt et al. (2013) and use CAMB_SOURCES (Challinor and Lewis, 2011) to compute and project the 3D matter power spectrum into angular auto- and cross-power spectra. We fix the cosmological parameters to *Planck* Λ CDM best-fit values⁵ (Planck Collaboration, 2013g). Importantly, although the matter power spectrum only depends on cosmological parameters, the projected power spectra require additional information about the samples under consideration: the unit-normalised redshift distribution of tracers n(z), and the logarithmic slope of the number counts s, to account for the effect of magnification due to lensing. We estimated s for the four samples by calculating the slope of the histogram of number counts in terms of the g-band PSF magnitude at g = 21. We found $\hat{s} = 0.361, 0.445, 0.493$, and 0.355, respectively, for these samples.

The redshift distributions are nontrivial to estimate because photometric redshift estimates of

⁵Specifically, $\Omega_c h^2 = 0.118$, $\Omega_m = 0.305$, $\Omega_b h^2 = 0.0221$, $H_0 = 67.7 \text{ kms}^{-1} \text{Mpc}^{-1}$, $\ln(10^{10} A_s) = 3.09$, $n_s = 0.961$, $\tau = 0.095$, $N_{\text{eff}} = 3.046$, and $\sum m_{\nu} = 0.05$.



Figure 3.2: Templates of the main systematics selected as potential contaminants in the photometric quasar catalogues, and used within the extended mode projection approach. Apart from the first row, all systematics are mapped in the *ugriz* bands and included in the analysis, but only templates for the *i* band are shown here. More details can be found in the text, and the full descriptions of the nature and units of these quantities (only referred to by their abbreviated names in this paper) can be found in the SDSS database and documentation.



Figure 3.3: A subset of systematics templates originating from the decorrelation (using Mask 1) of the basic set of systematics, a subset of which is presented in Fig. 3.2.

quasars suffer from large uncertainties (typically $\sigma(\hat{z}_{\text{phot}}) \sim 0.3$) and large fractions of outliers (see e.g., Bovy et al. 2012; Richards et al. 2009). In the present case, the actual n(z) of the quasar samples extend significantly beyond the windows used to construct them. This not only introduces uncertainties in the theoretical predictions, but also cosmological information in the cross-power spectra, since the n(z) of the samples are likely to overlap with each other, as detailed in the next sections. One solution to reduce the associated uncertainties is to employ low-resolution histograms, which are more robust to individual photo-z uncertainties, but do not capture all the structure of the underlying n(z). Alternative methods also exist to attempt to estimate high-resolution n(z)from the poorly-determined photometric redshifts (e.g., Matthews and Newman 2010; McQuinn and White 2013). In Leistedt et al. (2013), we used a cross-matching technique to obtain unbiased estimates for the n(z) of the Richards et al. (2009) catalogue SDSS DR6 photometric quasars. In our case the XDQSOz catalogue includes posterior distributions for the redshift estimates of individual objects, which are computed with the same data-driven model used for the classification, and trained with the best spectroscopic data available. Hence, a simple estimator for n(z) is the sum of the posterior distributions of the individual objects in each sample. Fig. 3.1 shows the resulting n(z), as well as a fit using a superposition of Gaussian kernels for use in CAMB_SOURCES. We also show the redshift histograms of cross-matched catalogues of quasars found in catalogues of spectroscopically-confirmed quasars from SDSS DR7 and BOSS DR9 (Paris et al., 2012; Schneider

et al., 2010). We applied the relevant flags and cuts to create spatially-uniform versions of these catalogues, *i.e.*, with constant depth, in order to avoid spurious selection effects due to depth variations. However, SDSS DR7 mostly contains bright objects, and BOSS DR9 is dominated by z > 2.2 quasars. Therefore the dashed redshift histograms shown in Fig. 3.1 are biased estimates of n(z), but nevertheless confirm the features found in the XDQSOz n(z). In particular, all four samples contain low-redshift quasars ($z \sim 0.7 - 1.0$), mainly because of their large scatter in *ugriz* colour space which overlaps with that of higher redshift quasars - hence the difficulty of separating them (Bovy et al., 2011a). The corrections required to follow Leistedt et al. (2013) and obtain unbiased n(z) estimates through cross-matching are more difficult to calculate for XDQSOz than the Richards et al. (2009) catalogue, because the former extends to fainter magnitudes and higher redshifts. In particular, using the full spectroscopic catalogue requires spatially-dependent corrections, which are much more uncertain due to sample variance. For these reasons, we did not attempt to derive and apply these corrections, but rather used the n(z) obtained using the stacked XDQSOz redshift posterior distributions. This should provide the most robust solution, since the XD data-driven model has not been demonstrated to suffer from any biases or systematics over the redshift ranges of interest.

In this paper, we focus on evaluating the quality of the measurements and the capacity of mode projection to mitigate systematics. Hence, we will fix the parameters used to calculate the theoretical angular power spectra, and also observational parameters such as the shot noise and the bias of quasars. The latter is assumed to be scale-independent and fixed to

$$b(z) = \left[1 + \left(\frac{1+z}{2.5}\right)^5\right],$$
 (3.12)

which is consistent with observations (*e.g.*, Porciani and Norberg 2006; Sherwin et al. 2012; Slosar et al. 2008). As shown in subsequent sections, these parameters provide a good fit to the measured power spectra. However, it is essential to consider their uncertainties when testing models and inferring cosmological parameters. Hence, all theoretical and observational parameters will be varied in the companion paper (Leistedt et al., 2014), where we constrain PNG and evaluate the robustness of these constraints to the underlying models and assumptions.

3.4.2. Masks and systematic uncertainties



Figure 3.4: Masks constructed for the power spectrum analysis of the XDQSOz samples. Mask 1 (grey + white) and Mask 2 (grey only) are created by applying dust extinction cuts E(B - V) > 0.1 and E(B - V) > 0.08 to the full SDSS DR8 coverage (black region), in addition to the common cuts Psfwidth_i > 2.0 and Score> 0.6 to avoid bad quality regions.

The main source of systematics in photometric quasar catalogues is stellar contamination, which arises because separating stars from quasars using imaging data is a difficult task. With the $P_{QSO} > 0.8$ cut, one can expect a significant fraction of the selected photometric quasars to be stars. However this proportion further depends on the redshift range of interest, since the colours of quasars, and thus the separability of the stellar and quasar loci, evolve with redshift. In addition, observing conditions unavoidably vary with time, and thus on the sky, as data are acquired. This affects both the number of detected point sources and the star-quasar separation, thus creating spatially-varying depth and stellar contamination. We have mapped as many of the potential systematics as possible in order to create templates for the mode projection framework. In what follows, we detail the construction of a base set and an extended set of systematics templates, which will be projected out when estimating the quasar power spectra.

For Galactic dust extinction, which affects the properties of point sources and photometric colours, we used the maps of E(B-V) from Schlegel et al. (1998) and *Planck* (Planck Collaboration, 2013e). They differ in some regions of the sky, especially near the Galactic plane, but are correlated, so we include both maps in the set of systematics prepared for extended mode projection. For the stellar density, we constructed a stellar density map from the SDSS DR6 point sources selected with 18.0 < r < 18.5 and i < 21.3, and a second map with an additional cut g - r > 1.4 to select red stars (Pullen and Hirata, 2013). We retrieved the data for the calibration and observing conditions from the FIELDS table in the SDSS CAS server. We mapped all quantities on the sky directly on the HEALPIX (Górski et al., 2005) grid, as the quasar maps are also manipulated in this format. This

solution does not use the exact geometry of the SDSS tiles (which would require the use of the MANGLE software (Hamilton and Tegmark, 2004; Swanson et al., 2008)) but is much faster while being sufficiently accurate on the scales of interest. Fig. 3.2 shows the base set of systematics we include in this study. Apart from the first row, all templates were constructed for the five *ugriz* bands, and only the *i* band is shown in Fig. 3.2. Also, the following quantities were not included on the figure for space reasons: ra, dec, run, fields, mjd, mustart, nustart, total, devapcorrection, devapcorrectionerr, devmodelapcorrection, devmodelapcorrectionerr, nfaintobj, nstaroffset. Full details of the nature, description and units of these quantities can be found in the SDSS database and documentation.

The total number of templates in this set of base systematics is 220, but after decorrelating these templates only ~ 100 have significant eigenvalues, the others being consistent with numerical noise, *i.e.*, redundant modes in the system of templates. A subset of these orthogonal modes is shown in Fig. 3. With such a small number of modes, all can be projected out when estimating power spectra, since they only generate a small increase in variance $\mathbf{V}_{bb'}$. However, this setting will only treat the spurious clustering due to linear combinations of these systematics. The actual contamination signal is likely to be non-linear and therefore may not be effectively eliminated. To consider non-linear contamination model, we create an extended set of templates, also including products of pairs of the base systematics. Once decorrelated, the $\sim 22,000$ maps were reduced to $\sim 3,700$ modes. Since they are orthogonal modes, projecting out all of them would be equivalent to removing $\sim 3,700$ data points in some rotated basis, thus significantly increasing the variance of the estimates. For this reason, we now use the extended mode projection approach to project out only those modes which are significantly correlated with the quasar data, as detailed in the next section.

Although extended mode projection considers non-linear contamination models, it is unlikely to fully clean the stellar contamination across the whole survey area, in particular in the worst regions where systematics are significant. Hence, we construct two masks, shown in Fig. 3.4. Both exclude pixels with PSFWIDTH_i > 2.0 and SCORE > 0.6. In addition, Mask 1 uses E(B - V) > 0.1and Mask 2 uses E(B - V) > 0.08, to remove more dusty regions in the South Galactic Cap. The masks were extended to avoid smoothing-induced contamination, as explained in Feeney et al. (2011b) and Leistedt et al. (2013). Note that previous SDSS quasar studies used more aggressive masks (Ho et al., 2013; Leistedt et al., 2013; Pullen and Hirata, 2013), and the point of extended mode projection precisely is to avoid extra sky cuts and obtain finer control over the elimination of contamination by projecting out the relevant systematics.

3.4.3. Estimation settings and blind mitigation of systematics

We used the QML estimator to simultaneously measure the auto- and cross- angular power spectra of the four XDQSOz samples. We used HEALPIX resolution of $N_{side} = 64$ for all quasar maps, masks and systematics templates in this work, in which case the power spectrum can be estimated accurately up to $\ell_{max} \sim 130$, which is a satisfactory band-limit given the noise levels and the sky cuts under consideration, and the computational requirements of the estimator. The choice of the multipole bin size is guided by the variance of the estimates. For most cases we adopted a bin size $\Delta \ell = 21$ to obtain a good compromise between variance and multipole resolution of the estimated power spectra. We also used $\Delta \ell = 15$ for the final runs in order to obtain a better resolution on the lowest multipoles, and therefore a better sensitivity to the PNG signal. The resulting constraints on PNG are presented in the companion paper (Leistedt et al., 2014), and are indeed more stringent in that setting. For the fiducial power spectrum priors required for QML, we used the theoretical power spectra detailed in the previous sections. The QML estimator is robust to small changes in these values, especially when measuring the auto- and cross- spectra, which are relatively featureless⁶, and simultaneously estimated from the four quasar maps and the full prior covariance matrix.

For the power spectrum measurements presented in the next section, we considered three settings for the treatment of systematics (in combination with the two masks of Fig. 3.4): no mode projection ("no mp"), projection of the base set of systematics "basic mp", ~ 100 uncorrelated modes), and projection of the extended set plus products of pairs of templates ("ext. mp", ~ 3700 uncorrelated modes). For the latter, we followed our implementation of extended mode projection, and only projected the modes which were significantly correlated with the data. To perform this selection, we cross-correlated the quasar maps with all orthogonal modes using the approximate QML estimator presented in Section 3.3.3, using the same setting as the main estimator (N_{side} = 64, $\ell_{max} = 130$, $\Delta \ell = 21$ or $\Delta \ell = 15$). The resulting cross-spectra were used as null tests, and we calculated a reduced χ^2 per systematics mode using a simple Gaussian likelihood (Leistedt et al., 2013). Fig. 3.5 shows the comparison of the reduced χ^2 obtained with the optimal and approximate QML estimators detailed in Section 3.3.3, and demonstrates that the latter is sufficiently accurate for the purpose of these null tests. We observe no significant bias, a reasonable scatter around the axis $\chi^2_{\text{opt.QML}} = \chi^2_{\text{approx.QML}}$, and a small number of outliers. In particular, we see that selecting modes

⁶Since quasars span large volumes, the projection integrals weaken features such as the baryon acoustic oscillations. The latter would only be detectable using narrow redshift samples, which cannot be done with XDQSOz due to the large uncertainties of the photometric redshift estimates.



Figure 3.5: Comparison of the reduced χ^2 obtained with the optimal and approximate QML crosscorrelation estimators. Each dot is a value obtained by using both estimators to calculate the cross-spectra between the four quasar samples (colour-coded) and the base set of 220 systematics templates.



Figure 3.6: Distributions of the reduced χ^2 arising from null tests performed with the cross-power spectra of the four quasar samples (four panels) with the ~ 3700 orthogonal modes of the extended set of systematics templates (~ 200 base templates plus products of pairs). The last step of extended mode projection consists of choosing a χ^2_{cut} to decide which modes will be projected out when estimating the angular power spectra of the quasar samples.

for mode projection with the cut $\chi^2_{approx.QML} > 1$ reproduces the selection using a cut $\chi^2_{opt.QML} > 1$ reasonably well, validating the use of the approximate estimator in the null tests.

Fig. 3.6 shows the distributions of the reduced χ^2 values obtained by cross-correlating the four quasar maps with the ~ 3700 modes. Interestingly, these distributions can be used to qualitatively evaluate the level of contamination of the samples, and also the effect of masking. In particular, the second and third samples are the least contaminated, and the χ^2 distributions are not improved by using the second mask instead of the first mask, in addition to being robust to the multipole resolution ($\Delta \ell = 21$ or $\Delta \ell = 15$). The null tests also demonstrate that the first and fourth samples are significantly contaminated by systematics, since a large fraction of the modes have large reduced χ^2 . Using resolution $\Delta \ell = 15$ yields even more modes with large χ^2 , because the cross spectra are able to extract features due to systematics which were not discernible with $\Delta \ell = 21$. In the next section, all power spectra calculated with extended mode projection were computed by projecting out modes with $\chi^2 > 1$. One can adjust this parameter to set the accepted level of correlation between with the data and the templates, above which modes are considered as contaminated by systematics and are projected out. One or multiple values for the χ^2 cut should be decided before the analysis, so that the systematics mitigation relies on objective criteria, and is performed in a blind fashion. Alternatively, iterating and refining the χ^2 cut during the analysis could improve the systematics mitigation, but may also yield fine-tuning and over-processing of the data if the results do not conform to expectation (confirmation bias), jeopardising the search for new physics.

We finally note that previous analyses of SDSS quasars (*e.g.*, Pullen and Hirata 2013) used samples with redshift distributions that were not overlapping: therefore their cross-power spectra could be used as null tests, *i.e.*, any systematics in common between the quasar samples would be detectable in the cross-spectra between redshift bins, which should be consistent with zero in the absence of systematics. In this work, the redshift distributions of the quasar samples overlap, some of them significantly, implying that the cross-spectra will contain cosmological information. We will demonstrate in the next section and the companion paper (Leistedt et al., 2014) that the fourth sample is not only the one with the most constraining power for PNG, but also the one with the redshift distribution overlapping the least with the other samples: its cross-power spectra can thus be used for approximate null tests.

3.4.4. Power spectrum measurements



Figure 3.7: Auto- and cross-angular power spectra of the four XDQSOz quasar samples measured with the quadratic maximum likelihood (QML) estimator, for the two masks in Fig. 3.4, without ("no mp") and with ("basic mp") mode projection of the base set of systematics (220 templates from the SDSS FIELDS table, shown in Fig. 3.2). The solid lines show our fiducial theoretical power spectra from a *Planck* best-fit cosmology and $f_{\rm NL} = 0$, and the shaded bands show the excursion region allowed when varying $f_{\rm NL}$ in [-50, 50] (see text for details).



Figure 3.8: Same as Fig. 3.7, but using extended mode projection. We used a non-linear contamination model incorporating the base set of systematics templates and also pairs of products of templates, yielding ~ 22,000 templates: but only ~ 3,700 orthogonal spatial modes remained after decorrelation. We calculated QML cross-angular power spectra to perform null-tests and calculated reduced χ^2 , as shown in Fig. 3.6.

Fig. 3.7 shows the results of the QML estimator applied to the quasar samples in three settings: Masks 1 and 2 without mode projection, and Mask 2 with mode projection of the base set of systematics templates. The black lines show our fiducial theoretical predictions, detailed in the previous section. The shaded band indicates the zone spanned by the theory curves when varying PNG in $-50 < f_{\rm NL} < 50$. Note that all spectra are dimensionless, since they relate to overdensity maps. At first examination, we see that all measured power spectra exhibit significant excess power on a range of scales compared to the theoretical expectation. In particular, the first and the fourth samples (the low and high redshift quasar samples) are in significant disagreement with the theoretical curve. The quasar cross-power spectra also exhibit power excesses, pointing to the presence of significant systematics and spatially-varying stellar contamination in these samples. Qualitatively, the discrepancy between the measurements and the prediction agrees with the levels of contamination measured by the null test χ^2 histograms of the previous section, indicating that at least some of the systematics responsible for the excesses are included in the sets of templates prepared for mode projection.

Indeed, the third set of spectra of Fig. 3.7 show that using the basic mode projection significantly decreases the power excesses in the auto-spectra, especially that of the fourth sample. The autopower spectra are now subject to an offset compared to the theoretical predictions. Although this could be due to inadequate cosmological or observational parameters in the predictions, such as underestimated biases, significant excess power in the cross-power spectra indicates that part of the remaining discrepancies must be due to systematics. In the formalism of Section 3.3.2, this also demonstrates that a linear model of the basic set of systematics (~ 100 uncorrelated modes) is insufficient to describe the contamination signal and spurious clustering. Moreover, since the error bars of Fig. 3.7 only include cosmic variance and shot noise, we deduce that the contamination level is so large that it prevents any reliable measurement of the quasar bias and PNG. In fact, constraining⁷ PNG from the measurements of Fig. 3.7 (ignoring the null test cross-spectra and thus the evidence for high levels of systematics) would yield very large levels of PNG, significantly above the limits set by previous studies of SDSS quasars, luminous red galaxies (LRGs), and in the CMB (e.g., Giannantonio et al. 2014; Planck Collaboration 2013h). The purpose of extended mode projection is precisely to address these issues, and reduce the levels of systematics below the statistical uncertainties by marginalising over the most contaminated modes.

⁷The error bars of Fig. 3.7 use the fiducial priors and are included to illustrate the large discrepancies and power excesses due to systematics. To exploit these power spectra in a likelihood analysis, one would need to recompute the covariance matrices around the best-fit theoretical power spectra, which were done for the cleanest power spectrum estimates of Fig. 3.8.

Fig. 3.8 shows three sets of estimates obtained when using extended mode projection with the more restrictive mask (Mask 2). The first set of spectra are calculated using band powers of width $\Delta \ell = 21$ for both quasar spectra and null tests, the third set uses $\Delta \ell = 15$ for both, and the second set is an intermediate setting where the quasar spectra use $\Delta \ell = 15$ and the null tests $\Delta \ell = 21$. This hybrid case is to check the robustness of the final spectra, since we observed in Fig. 3.6 that the null tests with $\Delta \ell = 21$ and $\Delta \ell = 15$ yielded slightly different χ^2 distributions for the first and fourth (most contaminated) quasar samples.

Firstly, we see that the three settings are in good overall agreement with each other, and also with the fiducial theoretical predictions, as demonstrated by the χ^2 values presented in Table 3.1. Indeed, despite a few band-powers that seem discrepant with the theory predictions at the $\sim 1 - 2\sigma$ level when examining Fig. 3.8, the values of the probability to exceed (PTE) show that the fiducial spectra are good fits to the measurements. This indicates that extended mode projection succeeds at mitigating the most important systematics and reduce the contamination to a level below the statistical uncertainties. Interestingly, the multipole resolution $\Delta \ell = 21$ yields better χ^2 than $\Delta \ell = 15$. This can be explained by the fact that some systematics had their correlation signatures averaged and smoothed out at $\Delta \ell = 21$, but resolved at $\Delta \ell = 15$. In other words, greater multipole resolution generally uncovers more systematics, and may require the use of more numerous or more complex combinations of systematics templates. But for the three sets of spectra shown in Fig. 3.8, none of the low- ℓ band powers in the auto- and cross-power spectra exhibit evidence for systematics. Therefore we conclude that they are suitable for cosmological inference and for constraining PNG, which produces a signal that is more pronounced on large scales. Indeed, the error bars and shaded bands of Fig. 3.8 demonstrate that the clean power spectra will be able to constrain PNG.

3.5. Conclusion

Photometric quasar surveys are deep and span extended redshift ranges, which allows us to probe the largest scales of the universe, and therefore test physics which is not well constrained by galaxy surveys. In particular, they can be used to constrain PNG, which is expected to enhance quasar clustering on large scales and leave a characteristic scale-dependent signature in the quasar bias. However, this requires accurate power spectrum measurements, which are compromised by the presence of numerous observational systematics, creating spurious correlations which can mimic

Table 3.1: The chi-square values for the power spectra measured using extended mode projection with Mask 2 with $\chi^2_{\text{cut}} = 1.0$, presented in Fig. 3.8. The theory curves use *Planck* Λ CDM cosmology, $f_{\text{NL}} = 0$, and fiducial linear bias $b(z) = 1 + [(1 + z)/2.5]^5$. The numbers of degrees of freedom are $\nu - p = 7$ and $\nu - p = 9$ for $\Delta \ell = 21$ and $\Delta \ell = 15$, respectively. The probability to exceed (PTE) the observed chi squares are shown in parentheses. Note that PTE < 1% corresponds to $\chi^2_7 = 18.4$ and $\chi^2_9 = 21.6$.

C_{ℓ} estimator:	$\Delta \ell = 21$	$\Delta \ell = 15$	$\Delta \ell = 15$
Ext. mp null tests:	$\Delta \ell = 21$	$\Delta \ell = 21$	$\Delta \ell = 15$
Bin 1 auto	4.81 (0.682)	7.42 (0.593)	7.66 (0.568)
Bin 2 $ imes$ Bin 1	8.15 (0.319)	6.13 (0.727)	5.67 (0.772)
Bin 2 auto	3.89 (0.792)	4.79 (0.852)	6.23 (0.716)
Bin 3 $ imes$ Bin 1	7.91 (0.341)	12.62 (0.180)	10.60 (0.304)
Bin 3 $ imes$ Bin 2	2.19 (0.948)	4.02 (0.909)	9.50 (0.392)
Bin 3 auto	0.76 (0.997)	6.39 (0.699)	9.16 (0.423)
Bin 4 $ imes$ Bin 1	8.43 (0.296)	13.60 (0.137)	10.71 (0.296)
Bin 4 \times Bin 2	9.33 (0.230)	13.00 (0.162)	17.24 (0.045)
Bin 4 \times Bin 3	5.83 (0.559)	7.55 (0.579)	12.94 (0.165)
Bin 4 auto	3.44 (0.8411)	3.14 (0.958)	4.07 (0.906)

the signatures of new physics.

We have introduced the *extended mode projection* technique to robustly mitigate the impact of large numbers of systematics when estimating angular power spectra, and applied it to the photometric quasars from the SDSS XDQSOz catalogue. This technique only relies on the ability to map known and potential sources of systematics on the sky, and cross-correlate them with the data of interest. Previous studies of XDQSOz data required stringent quality and sky cuts, and even the removal of band-powers in order to avoid excessive contamination by systematics. In our analysis, we have used minimal sky cuts, and applied the extended mode projection approach using a large number of systematics templates. Mode projection is equivalent to a Bayesian marginalisation over the amplitudes of the modes of the contamination model when estimating the power spectra. The base set of templates included 220 potential systematics found in the SDSS database, and we have also included products of pairs of templates, leading to a total of $\sim 22,000$ systematics templates, yielding a non-linear model for the contamination signal. We have then decorrelated the systematics, and cross-correlated the resulting $\sim 3,700$ orthogonal modes with the quasar samples to carry out null tests and detect the modes which most likely create spurious correlations in the data. We have finally estimated clean quasar power spectra by projecting out the modes which yielded reduced $\chi^2 > 1$ for the cross-spectrum null tests. Our pool of systematics and resulting contamination model was very general, and the sky masks minimal; thus, the reduced χ^2 cut is the only tuneable parameter in the extended mode projection approach. Our approach is therefore

based on the principles of blind analysis, since projecting out modes with reduced $\chi^2 > 1$ is a simple and pre-selected criterion for the accepted level of correlation between the systematics and the maps, which does not depend on the intrinsic clustering of quasars. Using various settings for both the power spectrum estimation and the systematics mitigation, we have tested that the power spectrum measurements are robust to these choices, and consistent with the theoretical predictions. In a companion paper (Leistedt et al., 2014), we show that these spectra, used in a combined likelihood function, yield stringent and robust constraints on PNG and on the bias of quasars. In particular, the constraints separately derived using the auto- and cross-spectra are consistent with each other, and robust to the underlying model and assumptions, for example to the uncertainties in the redshift distributions of the samples. This demonstrates that the remaining levels of correlations created by systematics are below the statistical uncertainties, and that the quasar power spectra are suitable for use in cosmological inferences.

Future galaxy and quasar survey data will reach unprecedented precision, and will require accurate mitigation of large numbers of systematics. For instance DES, *Euclid* (Amendola et al., 2013), and LSST (Abell et al., 2009) will observe hundreds of millions of objects, and probe extended redshift ranges and finer angular scales. This makes them very promising for testing new physics beyond the standard cosmological model, such as PNG, the neutrino sector, dark energy phenomenology, and modifications to General Relativity. At such precision levels, and given that these new physics signatures are typically small deviations from the standard model, power spectrum measurements will be highly sensitive to any systematics. The extended mode projection framework is a good candidate for mitigating such systematics in a robust and blind fashion, while extracting as much information as possible from the hard-won data.

3.6. Acknowledgements

We are grateful to Jo Bovy, David Hogg and Adam Myers for sharing the XDQSOz catalogue with us. We also thank Daniel Mortlock, Nina Roth, Filipe Abdalla, Aurélien Benoit-Lévy, and Blake Sherwin, for useful discussions and comments. BL is supported by the Perren Fund and the IMPACT Fund. HVP is supported by STFC and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007- 2013) / ERC grant agreement no 306478-CosmicDawn.

We acknowledge use of the following public software packages: HEALPIX (Górski et al., 2005) and CAMB_SOURCES (Challinor and Lewis, 2011). We acknowledge use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science.

This work is based on observations obtained with SDSS. Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is http://www.sdss.org/. The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

Constraints on primordial non-Gaussianity from 800,000 photometric quasars

"No one trusts a model except the man who wrote it; everyone trusts an observation, except the man who made it". H. Shapley

4.1. Abstract

We derive robust constraints on primordial non-Gaussianity (PNG) using the clustering of 800,000 photometric quasars from the Sloan Digital Sky Survey in the redshift range 0.5 < z < 3.5. These measurements rely on the novel technique of *extended mode projection* to control the impact of spatially-varying systematics in a robust fashion, making use of blind analysis techniques. This allows the accurate measurement of quasar halo bias at the largest scales, while discarding as little as possible of the data. The standard local-type PNG parameters $f_{\rm NL}$ and $g_{\rm NL}$ both imprint a k^{-2} scale-dependent effect in the bias. Constraining these individually, we obtain $-49 < f_{\rm NL} < 31$ and $-2.7 \times 10^5 < g_{\rm NL} < 1.9 \times 10^5$, while their joint constraints lead to $-105 < f_{\rm NL} < 72$ and $-4.0 \times 10^5 < g_{\rm NL} < 4.9 \times 10^5$ (all at 95% CL). Introducing a running parameter $n_{f_{\rm NL}}$ to constrain $b(k) \propto k^{-2+n_{f_{\rm NL}}}$ and a generalised PNG amplitude $\tilde{f}_{\rm NL}$, we obtain $-45.5 \exp(3.7 n_{f_{\rm NL}}) < \tilde{f}_{\rm NL} < 34.4 \exp(3.3 n_{f_{\rm NL}})$ at 95% CL. These results incorporate uncertainties in the cosmological parameters, redshift distributions, shot noise, and the bias prescription used to relate the quasar clustering to the underlying dark matter. These are the strongest constraints obtained to date on PNG using a single population of large-scale structure tracers, and are already at the level of pre-*Planck* constraints from the cosmic microwave background. A conservative forecast for a *Large Synoptic Survey Telescope*-like

survey incorporating mode projection yields $\sigma(f_{\rm NL}) \sim 5$ – competitive with the *Planck* result – highlighting the power of upcoming large scale structure surveys to probe the initial conditions of the universe.

4.2. Introduction

Canonical single-field slow-roll inflation predicts initial conditions for structure formation that are essentially Gaussian (Acquaviva et al., 2003; Allen et al., 1987; Babich et al., 2004; Bartolo et al., 2004; Falk et al., 1993; Gangui and Martin, 2000; Gangui et al., 1994; Komatsu and Spergel, 2001; Maldacena, 2003; Salopek and Bond, 1990; Verde et al., 2000; Wang and Kamionkowski, 2000). Any measurement of deviations from this prediction — summarised by the term primordial non-Gaussianity (PNG) — can thus provide evidence for non-standard inflationary physics. One of the most physically interesting forms of PNG is the so-called *local* model, where the primordial potential ϕ is modified by including higher order terms,

$$\Phi = \phi + f_{\rm NL}[\phi^2 - \langle \phi^2 \rangle] + g_{\rm NL}[\phi^3 - 3\phi \langle \phi^2 \rangle], \tag{4.1}$$

where all fields are evaluated at the same spatial coordinate, and $f_{\rm NL}$ and $g_{\rm NL}$ are real-valued constants (often called the *skewness* and *kurtosis* parameters).

The most stringent constraints on PNG currently come from higher-order statistics of the cosmic microwave background (CMB). Most recently, the *Planck* collaboration reported $-8.9 < f_{\rm NL} < 14.3$ (95% CL) (Planck Collaboration, 2013h), while constraints on the kurtosis have been obtained from the WMAP satellite: $-7.4 \times 10^5 < g_{\rm NL} < 8.2 \times 10^5$ (WMAP5, 95% CL) (Smidt et al., 2010), or $g_{\rm NL} = (-4.3 \pm 2.3) \times 10^5$ (WMAP9, 68% CL) (Regan et al., 2013), $g_{\rm NL} = (-3.3 \pm 2.2) \times 10^5$ (WMAP9, 68% CL) (Sekiguchi and Sugiyama, 2013). While these results are compatible with Gaussian initial conditions, their uncertainties still leave room for non-standard inflation models.

The unknown relation (*bias*) between the dark matter density field and a set of observed tracers (which inhabit dark matter halos) is generally considered to be a *complication* in constraining cosmological parameters from large-scale structure (LSS) data. However, in the case of PNG, the bias is actually an *advantage* that can be used to distinguish between non-Gaussianity in the initial conditions and that generated through late-time non-linear structure formation. PNG introduces a distinctive *k*-dependence into the halo bias; qualitatively, the bias for local-type PNG scales as

 $b(k) \sim k^{-2}$ (Dalal et al., 2008; Desjacques and Seljak, 2010b; Matarrese and Verde, 2008; Slosar et al., 2008; Smith et al., 2012).

This implies that the strongest signal can be expected on large scales (small *k*), accessible to wide-area galaxy surveys. At these scales, the bias can be well-approximated by a multiplicative factor between the dark matter- and galaxy power spectra. LSS clustering constraints on PNG provide an independent validation of the CMB results, and are predicted to improve significantly with on-going and future LSS surveys, eventually surpassing CMB constraints if systematic errors can be controlled (Carbone et al., 2008; Desjacques and Seljak, 2010a; Giannantonio et al., 2012; Hamaus et al., 2011; Lidz et al., 2013; Mao et al., 2013; Marian et al., 2011; Pillepich et al., 2012). Quasars – the bright nuclei at the centre of the most active galaxies – are highly-biased tracers of the LSS, spanning large volumes and covering extended redshift ranges: in principle, quasar surveys are ideal for constraining PNG Giannantonio et al. (2014); Slosar et al. (2008). However, previous analyses have been complicated by the presence of spurious excess power at large scales due to systematics, which mimic the signature of PNG (Giannantonio et al., 2014; Leistedt et al., 2013; Pullen and Hirata, 2013; Xia et al., 2009, 2011).

In this *Letter*, we use a large sample of quasars (Bovy et al., 2012) from the Sloan Digital Sky Survey (SDSS) (Gunn et al., 2006) to constrain PNG. Using a novel technique for blind mitigation of systematics described in Leistedt and Peiris (2014), we are able to significantly enhance the constraining power of the dataset, resulting in PNG constraints from a single LSS dataset which are competitive with those from the CMB. Our results represent a significant step toward achieving the exquisite control of systematics necessary to exploit future LSS surveys to measure PNG.

4.3. PNG with photometric quasars

The first LSS constraints on PNG were derived in Slosar et al. (2008) using a combination of tracers from early SDSS releases, leading to $f_{\rm NL} = 28^{+23}_{-24}$ (68% CL). Among these tracers, the photometric quasars — candidate quasars identified using imaging data only — have the highest bias and probe the largest volume. Therefore, they had the most constraining power ($f_{\rm NL} = 8^{+26}_{-37}$ at 68% CL), demonstrating their potential to constrain PNG and complement CMB experiments. However, subsequent analyses (Giannantonio et al., 2014; Xia et al., 2009, 2011) of photometric quasars from the Sixth SDSS Data Release (using the catalogue from Richards et al. 2009) also revealed
systematic effects, such as spatially-varying depth and stellar contamination, which could strongly bias the clustering measurements on the largest scales and jeopardise cosmological inferences if not properly mitigated (Leistedt et al., 2013; Pullen and Hirata, 2013). Giannantonio and Percival (2014); Giannantonio et al. (2014) obtained $f_{\rm NL} = 5 \pm 21$ (68% CL) using a range of LSS probes, discarding the auto-correlation of quasars to avoid the main systematic contamination. Ho et al. (2013) used the latest catalogue of SDSS photometric quasars, XDQSOz (Bovy et al., 2011a, 2012), to obtain $f_{\rm NL} = 103^{+148}_{-146}$ (68% CL), and $f_{\rm NL} = 2^{+65}_{-66}$ (68% CL) when combined with constraints from the clustering of Luminous Red Galaxies, relying however on stringent quality cuts on the survey maps and power spectra to limit the impact of systematics. Therefore, photometric quasars have thus far remained underexploited for PNG, as systematics are handled by removing hard-won data.

4.4. XDQSOz power spectrum measurements

A different approach was adopted by two of us in Leistedt and Peiris (2014), where the SDSS XDQSOz catalogue was analysed with minimal quality and sky cuts, leading to a sample of ~ 800,000 quasars covering ~ 8300 deg² (compared to ~ 400,000–500,000 objects used by previous analyses). This base sample was further separated into four redshift bins by selecting objects with photometric redshift estimates \hat{z}_p in top-hat windows [0.5, 1.35], [1.35, 1.7], [1.7, 2.2], [2.2, 3.5]. A quadratic maximum likelihood Bond et al. (1998); Tegmark (1997) method was used to simultaneously estimate the auto- and cross-angular power spectra of the *z*-binned data. To mitigate the impact of systematics in these power spectra, Leistedt and Peiris (2014) introduced a novel technique, *extended mode projection*, relying on the fact that most potential systematics (*e.g.*, observing conditions, calibration) were also measured during SDSS observations, and could therefore be mapped onto the sky.

We constructed a non-linear, data-driven model of systematics, using \sim 3,700 orthogonal templates obtained by decorrelating \sim 20,000 maps of potential contaminants, including 200 base templates constructed from SDSS data and products of pairs. These orthogonal templates were cross-correlated with the XQDSOz data, yielding null tests which are used to select the most significant systematics following the principles of blind analysis. These were then marginalised over via *mode projection* (Leistedt et al., 2013; Pullen and Hirata, 2013; Slosar et al., 2004; Tegmark,



Figure 4.1: Redshift distributions of the four quasar samples used in this analysis, parametrised as superpositions of normal distributions. The shaded regions show the 1, 2 and 3σ regions explored when adding 5% uncertainty to the parameters of these fits, which are included in the MCMC analysis.

1997; Tegmark et al., 1998; Wandelt et al., 2004) in the power spectrum estimator, self-consistently enhancing the estimator variance. This approach allows precision control over systematics, and yields robust measurements of the angular power spectra of XDQSOz quasars, even at the largest angular scales.

We now turn to the ingredients and models needed to connect the power spectrum measurements to theoretical predictions and constrain PNG.

The redshift distributions n(z) of the four quasar samples, shown in Fig. 4.1, were estimated in Leistedt and Peiris (2014) by stacking the posterior distributions of the individual photometric quasars, and then fit with a superposition of Gaussians. Here, we also added a 5% Gaussian uncertainty on the parameters of this fit (illustrated by the shaded bands in Fig. 4.1) which propagates into the final uncertainties in the PNG parameters.

The shot noise was measured in each sample from the average surface density of photometric quasars, but is also subject to uncertainties due to the unknown fraction of stars in the samples — between 0 and 20% from the quality cuts applied to XDQSOz (Leistedt and Peiris, 2014). Since non-zero stellar contamination reduces the shot noise, we also marginalised over this effect when constraining PNG.



Figure 4.2: Constraints on the quasar bias model described in Eq. 4.4. The solid line shows the fiducial model with $b_0 = 1, \beta = 5$, and the shaded bands show the 1σ constraints ($b_0 = 0.96 \pm 0.15$, 68% CL, β and $f_{\rm NL}$ marginalised) from the XDQSOz power spectra when varying the bias and PNG parameters, the shot noise, and the redshift distributions. The coloured bands show the results when fixing $\beta = 5$ and allowing a different bias amplitude in each redshift bin, to demonstrate the ability of the overall model to simultaneously describe the four samples.

4.5. Halo bias from PNG

The impact of local-type PNG is to modify the halo bias by adding a k-dependent term to the Gaussian bias $b^G(z)$ (Dalal et al., 2008; Desjacques and Seljak, 2010b; Matarrese and Verde, 2008; Slosar et al., 2008; Smith et al., 2012)

$$b^{\rm NG}(k,z) = b^{\rm G}(z) + \frac{\beta_f(z)f_{\rm NL} + \beta_g(z)g_{\rm NL}}{\alpha(k,z)}.$$
 (4.2)

Here, we neglect an additional small contribution induced by the effect of PNG on the halo mass function, which is independent of k and can thus be absorbed in $b^G(z)$. Note that we have also suppressed the implicit mass dependence of b^G , β_f , and β_g in the previous equation. The exact expression for $\alpha(k, z)$ and fitting functions for β_f and β_g can be found in Smith et al. (2012).

A simple extension of the local model is the introduction of a spectral index $n_{f_{\rm NL}}$ in the $f_{\rm NL}$ generated scale-dependent bias (Agarwal et al., 2014b; Agullo and Shandera, 2012; Dias et al., 2013; Giannantonio et al., 2014; Wagner and Verde, 2012), *i.e.*, changing its scaling from k^{-2} into $k^{-2+n_{f_{\rm NL}}}$ by using

$$\alpha(k,z) \rightarrow \alpha(k,z) \left(\frac{k}{k_{\text{piv}}}\right)^{-n_{f_{\text{NL}}}},$$
(4.3)



Figure 4.3: Constraints on local-type $f_{\rm NL}$ (in the $\Lambda \text{CDM} + f_{\rm NL}$ model, with $n_{f_{\rm NL}} = g_{\rm NL} = 0$) using the power spectrum analysis of XDQSOz quasars, for different bias models and incorporating uncertainties in the redshift distributions and cosmological parameters. The error bars show the 1 and 2σ constraints, the dashed line shows $f_{\rm NL} = 0$, and the shaded bands show the constraints from *Planck* (Planck Collaboration, 2013h).

where we choose $k_{piv} = 0.06 \text{ Mpc}^{-1}$. Note that this parametrisation is not equivalent to an intrinsically scale-dependent f_{NL} as described in Becker et al. (2011, 2012). Instead, it allows us to extend our analysis to other types of PNG, like that generated by single-field inflation with a modified initial state Agullo and Shandera (2012), or models with several light fields (Dias et al., 2013).

The quasar bias is known to evolve strongly with redshift (*e.g.*, Myers et al. 2006, 2007a,b; Shen et al. 2009; White et al. 2012), and thus one cannot use a constant linear bias per redshift bin due to the extended and complicated redshift distributions shown in Fig. 4.1. For the Gaussian bias b(z) in Eq. 4.2, we used

$$b(z) = b_0 \left[1 + \left(\frac{1+z}{2.5}\right)^{\beta} \right],$$
 (4.4)

which is in good agreement with previous studies of SDSS quasars (*e.g.*, Porciani and Norberg 2006).

4.6. Monte Carlo Markov Chain (MCMC) analysis

We built a Gaussian likelihood (Leistedt et al., 2013), jointly using the 10 auto- and cross-

angular power spectra (between redshift bins) estimated in Leistedt and Peiris (2014), at multipole resolution $\Delta \ell = 15$. The theoretical predictions were calculated using CAMB_sources (Challinor and Lewis, 2011), modified to support PNG and our quasar bias model. We used emcee (Foreman-Mackey et al., 2013) to run an MCMC analysis, and sample combinations of the following parameters: <u>Cosmological parameters</u> ('cosmo'): parameters of the base Λ CDM model, with fiducial values and uncertainties corresponding to the constraints from *Planck* combined with Baryon Acoustic Oscillations (BAO), as in Planck Collaboration (2013g). <u>Bias model</u>: the model described above, with uniform priors $b_0 \in [0, 2]$ and $\beta \in [4, 6]$. <u>Redshift distributions</u> ('n(z)'): the amplitude and width of the Gaussian functions used to fit the n(z) estimates, with Gaussian priors of 5% 1 σ uncertainties around the fiducial values. Additionally, we sampled the slope of number counts, which controls magnification bias, with Gaussian priors centred at the measured value with 5% 1 σ uncertainty. <u>Shot noise</u>: we marginalised over the shot noise with a prior [0.8, 1.0] times the value measured from the photometric quasar surface density, in order to account for the unknown (but bounded) amount of stellar contamination.

4.7. Results

We first test the robustness of the bias model by examining the bias measured in the four redshift samples individually and jointly. Therefore, in addition to the 'coupled' model presented above, used to connect all power spectra to the theory predictions, we consider an alternative, 'decoupled' case where the bias amplitude of each redshift sample is fit separately, using four parameters b_1, b_2, b_3, b_4 . In this case, we used $\beta = 5$ and uniform priors $b_i \in [0, 2]$. The constraints on the bias parameters from the XDQSOz power spectra are shown in Fig. 4.2, and demonstrate that the separate bias amplitudes $b_i, i = 1 \dots 4$ of the four samples are in good agreement with each other, with the fiducial model with $b_0 = 1$ and $\beta = 5$ (black line), and also with the results obtained with the coupled model (shaded band). Note that the slope parameter β is used to capture the uncertainty in the evolution of the bias at z > 2.5. This redshift range is not as strongly constrained by the data, but is nevertheless crucial since it has the greatest bias, and therefore is expected to produce the strongest PNG signature.

Having confirmed that Eq. 4.4 is an adequate model for the quasar bias, we can now advance towards constraining the PNG parameters. Unless stated otherwise, all values are quoted at 95% CL.



Figure 4.4: 1σ and 2σ joint constraints on $f_{\rm NL}$ and $g_{\rm NL}$ for the $(b_0, \beta) + n(z)$ +cosmo case, *i.e.*, marginalising over the uncertainties in the cosmological parameters, redshift distributions, and bias model.



Figure 4.5: 1σ and 2σ joint constraints on $\tilde{f}_{\rm NL}$ and $n_{f_{\rm NL}}$ for the extended model of Eq. 4.3, marginalising over the same parameters as Fig. 4.4.

Model	$f_{\rm NL}$	$g_{\rm NL}/10^5$	
$base + f_{NL}$	-4^{+27+8}_{-35-10}	_	
$base+g_{NL}$	_	$0.3^{+1.1+0.5}_{-2.2-0.8}$	
$base + f_{NL} + g_{NL}$	-10^{+54+28}_{-60-35}	$0.2^{+2.9+1.8}_{-2.7-1.5}$	

Table 4.1: PNG constraints for different models. The base model is (b_0, β) +cosmo+n(z). The first super/subscript corresponds to the 68% CL, and the 95% CL can be obtained by adding the second number.

Fig. 4.3 shows the constraints on $f_{\rm NL}$ only (with $n_{f_{\rm NL}} = g_{\rm NL} = 0$) for various combinations of parameters and sources of uncertainties. As expected, adding parameters (*e.g.*, using the 'decoupled' bias model) increases the error estimates. When only varying b_0 , β , we obtain $-26 < f_{\rm NL} < 34$ (top row), and adding the uncertainties in the cosmological parameters and redshift distributions relaxes the constraints to $-49 < f_{\rm NL} < 31$ (bottom row). Note that this case is more conservative but introduces degeneracies between parameters, in particular those affecting the overall amplitude of the power spectra, such as the bias b_0 , the amplitude of matter fluctuations σ_8 , and the relative amplitudes of the peaks in n(z). However, $f_{\rm NL}$ does not suffer from any direct degeneracies with these parameters given its specific signature on large scales, yielding constraints which are robust to the choice of model.

When only using $g_{\rm NL}$ as a source of scale-dependent bias, we obtain $-2.7 \times 10^5 < g_{\rm NL} < 1.9 \times 10^5$. However, as apparent from Eq. 4.2, and as shown in Roth and Porciani (2012), $f_{\rm NL}$ and $g_{\rm NL}$ leave similar signatures on the scale-dependent bias of LSS tracers, and are not easily separable. Despite the large volume and extended redshift range spanned by our data, this degeneracy is confirmed by their joint constraints, shown in Fig. 4.4, leading to $-105 < f_{\rm NL} < 72$ and $-4.0 \times 10^5 < g_{\rm NL} < 4.9 \times 10^5$.

For the extended model of Eq. 4.3, we can introduce a generalised parameter $\tilde{f}_{\rm NL}$, which only corresponds to $f_{\rm NL}$ as defined in Eq. 4.1 when $n_{f_{\rm NL}} = 0$. In this model, the scale-dependent bias scales as $b(k) \sim k^{-2+n_{f_{\rm NL}}}$, so the constraints on $\tilde{f}_{\rm NL}$ widen as $n_{f_{\rm NL}}$ increases, due to the less pronounced signature of PNG on large scales. Nevertheless, $\tilde{f}_{\rm NL}$ can be constrained at any $n_{f_{\rm NL}} < 2$. Fig. 4.5 shows the joint constraints on $\tilde{f}_{\rm NL}$ and $n_{f_{\rm NL}}$, summarised as $-45.5 \exp(3.7 n_{f_{\rm NL}}) < \tilde{f}_{\rm NL} < 34.4 \exp(3.3 n_{f_{\rm NL}})$, fully compatible with $\tilde{f}_{\rm NL} = 0$.

4.8. Future prospects

We performed a Fisher forecast of a *Large Synoptic Survey Telescope*-like photometric survey (LSST Dark Energy Science Collaboration, 2012), using angular power spectra of galaxies in 20 tomographic bins in the redshift range 0.5 < z < 3.5, with bias modelled by $b_0 = 1$, $\beta = 3$ in Eq. 4.4. This idealised forecast, without any contamination by systematics, yields a 1σ uncertainty on $f_{\rm NL}$ of \sim 1, consistent with the 3D power spectrum analysis of Carbone et al. (2008).

We then incorporated realistic estimates of the main 50 systematics identified in Leistedt and Peiris (2014). This translates into a ~ 10% clustering contamination for $\ell > 100$, up to factors of a few on the largest scales where the PNG signal is strongest. Without any mitigation in the subsequent analysis, the resulting estimate for $f_{\rm NL}$ is biased by $\Delta f_{\rm NL} \sim 30$, which is highly significant compared to the expected uncertainty. Including a harmonic-space mode projection in the forecast removes this bias, while consistently increasing the parameter uncertainty to $\sigma(f_{\rm NL}) = 5$, comparable to current *Planck* constraints. However, this increase in uncertainty will be significantly reduced when performing the mode projection in pixel space, using sky maps of the systematics which will be available for a real survey (Leistedt and Peiris, 2014; Leistedt et al., 2013).

4.9. Conclusion

Photometric quasar surveys, while potentially constituting ideal datasets for probing PNG, have thus far been systematics-limited. We derived constraints on PNG using the results of a novel power-spectrum estimation method applied to ~ 800,000 photometric quasars from SDSS. This approach self-consistently marginalises over a non-linear data-driven model of spatially-varying systematics. Our results, summarised in Table 4.1, incorporate uncertainties in the cosmological parameters and the parameters of a realistic bias model, while marginalising over uncertainties in the redshift distributions of the quasars. The resulting constraints on $f_{\rm NL}$ and $g_{\rm NL}$ are the tightest obtained using a single population of LSS tracers, and are at the level of pre-*Planck* CMB constraints. Our results demonstrate the potential of future LSS surveys to reach the $f_{\rm NL} \sim 1$ levels predicted by the simplest models of inflation. We are grateful to Jo Bovy, David Hogg and Adam Myers for sharing the XDQSOz catalogue with us, and Jo Bovy, Masahiro Takada for useful conversations.

BL is supported by the Perren Fund and the IMPACT Fund. HVP and NR are supported by STFC and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007- 2013) / ERC grant agreement no 306478-CosmicDawn.

We acknowledge use of the following public software packages: HEALPIX (Górski et al., 2005) and CAMB_SOURCES (Challinor and Lewis, 2011). We acknowledge use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science. This work is based on observations obtained with SDSS. Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is http://www.sdss.org/.

S2LET: A code to perform fast wavelet analysis on the sphere

"Energy is neither created nor destroyed. It just changes shape". S. Reynolds

5.1. Abstract

We describe S2LET, a fast and robust implementation of the scale-discretised wavelet transform on the sphere. Wavelets are constructed through a tiling of the harmonic line and can be used to probe spatially localised, scale-dependent features of signals on the sphere. The reconstruction of a signal from its wavelets coefficients is made exact here through the use of a sampling theorem on the sphere. Moreover, a multiresolution algorithm is presented to capture all information of each wavelet scale in the minimal number of samples on the sphere. In addition S2LET supports the HEALPix pixelisation scheme, in which case the transform is not exact but nevertheless achieves good numerical accuracy. The core routines of S2LET are written in C and have interfaces in Matlab, IDL and Java. Real signals can be written to and read from FITS files and plotted as Mollweide projections. The S2LET code is made publicly available, is extensively documented, and ships with several examples in the four languages supported. At present the code is restricted to axisymmetric wavelets but will be extended to directional, steerable wavelets in a future release.

5.2. Introduction

Signals defined or measured on the sphere arise in numerous disciplines, where analysis techniques defined explicitly on the sphere are now in common use. In particular, wavelets on the sphere (Antoine and Vandergheynst, 1998, 1999; Baldi et al., 2009; Marinucci et al., 2008; McEwen et al., 2006; Narcowich et al., 2006; Starck et al., 2006a; Wiaux et al., 2005, 2006, 2007, 2008; Yeo et al., 2008) have been applied very successfully to problems in astrophysics and cosmology, where data-sets are increasingly large and need to be analysed at high resolution in order to confront accurate theoretical predictions (e.g. Barreiro et al. 2000; Basak and Delabrouille 2012; Cayón et al. 2001; Deriaz et al. 2012; Labatie et al. 2012; Lan and Marinucci 2008; McEwen et al. 2006, 2007a,b, 2008; Pietrobon et al. 2008; Schmitt et al. 2010; Starck et al. 2006b; Vielva et al. 2004, 2006a,b, 2007; Wiaux et al. 2006, 2008).

While wavelet theory is well established in Euclidean space (see *e.g.* Daubechies 1992b), multiple wavelet frameworks have been developed on the sphere, only a fraction of which lead to exact transforms in both the continuous and discrete settings. In fact, discrete methodologies (Schröder and Sweldens, 1995; Sweldens, 1996, 1997) achieve exactness in practice but may not lead to a stable basis on the sphere (Sweldens, 1997). In the continuous setting several constructions are theoretically exact, and have been combined with sampling theorems on the sphere to enable exact reconstruction in the discrete setting also. In particular, scale-discretised wavelets (Wiaux et al., 2008) lean on a tiling of the harmonic line to yield an exact wavelet transform in both the continuous and discrete settings. In the axisymmetric case, the scale-discretised wavelets reduce to needlets (Baldi et al., 2009; Marinucci et al., 2008; Narcowich et al., 2006), which were developed independently using an analogous tiling of the harmonic line. Similarly, the isotropic undecimated wavelet transform (UWT) developed by (Starck et al., 2006a) exploits B-splines of order 3 to cover the harmonic line with filters with greater overlap but nevertheless compact support.

In this paper we describe the new publicly available S2LET¹ code to perform the scale-discretised wavelet transform of complex signals on the sphere. At present S2LET is restricted to axisymmetric wavelets (i.e. azimuthally symmetric when centred on the poles) and includes generating functions for axisymmetric scale-discretised wavelets (Wiaux et al., 2008), needlets (Baldi et al., 2009; Marinucci et al., 2008; Narcowich et al., 2006) and B-spline wavelets (Starck et al., 2006a). We intend to extend the code to directional, steerable wavelets and spin functions in a future release.

¹http://www.s2let.org/

The core routines of S2LET are written in C, exploit fast algorithms on the sphere, and have interfaces in Matlab, IDL and Java.

We note that many very useful public codes are already available to compute wavelet transforms on the sphere, including isotropic undecimated wavelet, ridgelet and curvelet transforms² (Starck et al., 2006a), invertible filter banks³ (Yeo et al., 2008), needlets (NeedATool⁴; Pietrobon et al. 2010) and scale-discretised wavelets (S2DW⁵; Wiaux et al. 2008). S2LET aims primarily to provide a fast and flexible implementation of the scale-discretised transform with exact reconstruction on the sphere using the sampling theorem of McEwen and Wiaux (2011), although it has also been extended to support some of the features of these other codes. Furthermore, particular attention has been paid in the development of S2LET to prove a user-friendly code, supporting multiple programming languages, and which is extensively documented.

The remainder of this article is organised as follows. In section 2 we detail the construction of scale-discretised axisymmetric wavelets and the corresponding exact scale-discretised wavelet transform on the sphere. In section 3 we describe the S2LET code, including implementation details, computational complexity and numerical performance. We present a number of simple examples using S2LET in section 4, along with the code to execute them. We conclude in section 5.

5.3. Wavelets on the sphere

We review the construction of scale-discretised wavelets on the sphere through tiling of the harmonic line (Wiaux et al., 2008). Directional, steerable wavelets were also considered by Wiaux et al. (2008), however we restrict our attention to axisymmetric wavelets here. Furthermore, the use of a sampling theorem on the sphere guarantees that spherical harmonic coefficients capture all the information content of band-limited signals, resulting in theoretically exact harmonic and wavelet transforms. One may alternatively adopt samplings of the sphere for which exact quadrature rules do not exist, such as HEALPix (Górski et al., 2005), but which nevertheless exhibit other useful properties, leading to numerically accurate but not theoretically exact transforms.

²http://jstarck.free.fr/mrs.html

³https://sites.google.com/site/yeoyeo02

⁴http://www.fisica.uniroma2.it/~pietrobon/

⁵http://www.spinsht.org/

5.3.1. Harmonic analysis on the sphere

The spherical harmonic decomposition of a square integrable signal $f \in L^2(S^2)$ on the twodimensional sphere S^2 reads

$$f(\omega) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} f_{\ell m} Y_{\ell m}(\omega), \qquad (5.1)$$

where $Y_{\ell m}$ are the spherical harmonic functions, which form the canonical orthogonal basis on S^2 . The spherical harmonic coefficients $f_{\ell m}$, with $\ell \in \mathbb{N}$ and $m \in \mathbb{Z}$ such that $|m| \leq \ell$, form a dual representation of the signal f in the harmonic basis on the sphere. The angular position $\omega = (\theta, \phi) \in S^2$ is specified by colatitude $\theta \in [0, \pi]$ and longitude $\phi \in [0, 2\pi)$. The spherical harmonic coefficients are given by

$$f_{\ell m} = \langle f | Y_{\ell m} \rangle = \int_{S^2} \mathrm{d}\Omega(\omega) f(\omega) Y_{\ell m}^*(\omega), \qquad (5.2)$$

with the surface element $d\Omega(\omega) = \sin\theta d\theta d\phi$. We consider band-limited signals in the spherical harmonic basis, with band-limit L if $f_{\ell m} = 0$, $\forall \ell \geq L$. For band-limited signals sampling theorems can be invoked so that both forward and inverse transforms can be reduced to finite summations that are theoretically exact. Sampling theorems effectively encode a quadrature rule for the exact evaluation of integrals on the sphere from a finite set of sampling nodes. Various sampling theorems exist in the literature (e.g. Driscoll and Healy 1994; Healy et al. 1996; McEwen and Wiaux 2011). In this work we adopt the McEwen and Wiaux (2011) sampling theorem (hereafter MW), which is based on an equiangular sampling scheme and, for a given band-limit L, requires the lowest number of samples on the sphere of all sampling theorems, namely $(L-1)(2L-1) + 1 \sim 2L^2$ samples (for comparison $\sim 4L^2$ samples are required by Driscoll and Healy (1994)). Fast algorithms to compute the corresponding spherical harmonic transform scale as $\mathcal{O}(L^3)$ and are numerically stable to band-limits of at least L = 4096 (McEwen and Wiaux, 2011). The GLESP pixelisation scheme (Doroshkevich et al., 2005) also provides a sampling theorem based on the Gauss-Legendre quadrature, and could be used in place of the MW sampling theorem. However, GLESP uses more samples than Gauss-Legendre quadrature requires, which may lead to an overhead when considering large band-limits and numerous wavelet scales. We focus on the MW sampling scheme to obtain a theoretically exact transform. Alternative sampling schemes that are not based on sampling theorems also exist such as HEALPix (Górski et al., 2005), which is supported by S2LET, MRS and Needatool. HEALPix does not lead to exact transforms on the sphere but the resulting approximate

transforms nevertheless achieve good accuracy and benefit from other practical advantages, such as equal-area pixels.

5.3.2. Scale-discretised wavelets on the sphere

The scale-discretised wavelet transform allows one to probe spatially localised, scale-dependent content in the signal of interest $f \in L^2(S^2)$. The *j*-th wavelet scale $W^{\Psi^j} \in L^2(S^2)$ is defined as the convolution of f with the wavelet $\Psi^j \in L^2(S^2)$:

$$W^{\Psi^{j}}(\omega) \equiv (f \star \Psi^{j})(\omega) \equiv \langle f | \mathcal{R}_{\omega} \Psi^{j} \rangle$$
$$\equiv \int_{S^{2}} \mathrm{d}\Omega(\omega') f(\omega') (\mathcal{R}_{\omega} \Psi^{j})^{*}(\omega'), \tag{5.3}$$

where * denotes complex conjugation. Convolution on the sphere is defined by the inner product of f with the rotated wavelet $\mathcal{R}_{\omega}\Psi^{j}$. We restrict our attention to axisymmetric wavelets, i.e. wavelets that are azimuthally symmetric when centred on the poles. Consequently, the rotation operator \mathcal{R}_{ω} is parameterised by angular position $\omega = (\theta, \phi)$ only and not also orientation⁶. For the axisymmetric case the spherical harmonic decomposition of $W^{\Psi^{j}}$ is then simply given by a weighted product in harmonic space:

$$W_{\ell m}^{\Psi^{j}} = \sqrt{\frac{4\pi}{2\ell+1}} f_{\ell m} \Psi_{\ell 0}^{j*}, \tag{5.4}$$

where $W_{\ell m}^{\Psi^{j}} = \langle W^{\Psi^{j}} | Y_{\ell m} \rangle$, $f_{\ell m} = \langle f | Y_{\ell m} \rangle$ and $\Psi_{\ell 0}^{j} \delta_{m 0} = \langle \Psi^{j} | Y_{\ell m} \rangle$, and where $\delta_{m 0}$ is the Kronecker delta symbol.

The wavelet coefficients extract the detail information of the signal only; a scaling function and corresponding scaling coefficients must be introduced to represent the low-frequency (low- ℓ), approximate information of the signal. The scaling coefficients $W^{\Phi} \in L^2(S^2)$ are defined by the convolution of f with the scaling function $\Phi \in L^2(S^2)$:

$$W^{\Phi}(\omega) \equiv (f \star \Phi)(\omega) = \langle f | \mathcal{R}_{\omega} \Phi \rangle, \qquad (5.5)$$

or in harmonic space,

$$W^{\Phi}_{\ell m} = \sqrt{\frac{4\pi}{2\ell+1}} f_{\ell m} \Phi^*_{\ell 0}, \tag{5.6}$$

⁶As already noted, the extension to directional scale-discretised wavelets has been derived by Wiaux et al. (2008). At present the S2LET code supports axisymmetric wavelets only; directional wavelets will be added in a later release.

where $W_{\ell m}^{\Phi} = \langle W^{\Phi} | Y_{\ell m} \rangle$ and $\Phi_{\ell 0} \delta_{m 0} = \langle \Phi | Y_{\ell m} \rangle$.

Provided the wavelets and scaling function satisfy an admissibility property (defined below), the function f may be reconstructed exactly from its wavelet and scaling coefficients by

$$f(\omega) = \int_{S^2} d\Omega(\omega') W^{\Phi}(\omega') (\mathcal{R}_{\omega'} \Phi)(\omega) + \sum_{j=J_0}^J \int_{S^2} d\Omega(\omega') W^{\Psi^j}(\omega') (\mathcal{R}_{\omega'} \Psi^j)(\omega),$$
(5.7)

or equivalently in harmonic space by

$$f_{\ell m} = \sqrt{\frac{4\pi}{2\ell+1}} W^{\Phi}_{\ell m} \Phi_{\ell 0} + \sqrt{\frac{4\pi}{2\ell+1}} \sum_{j=J_0}^{J} W^{\Psi^j}_{\ell m} \Psi^j_{\ell 0}.$$
 (5.8)

The parameters J_0 , J define the lowest and highest scales j of the wavelet decomposition and must be defined consistently to extract and reconstruct all the information content of f. These parameters depend on the construction of the wavelets and scaling function and are defined explicitly in the next paragraphs. The admissibility condition under which a band-limited function f can be decomposed and reconstructed exactly is given by the following resolution of the identity:

$$\frac{4\pi}{2\ell+1} \left(|\Phi_{\ell 0}|^2 + \sum_{j=J_0}^J |\Psi_{\ell 0}^j|^2 \right) = 1, \quad \forall \ell.$$
(5.9)

We are now in a position to define wavelets and a scaling function that satisfy the admissibility property. In this paper, we use the smooth generating functions defined by Wiaux et al. (2008) in order to tile the harmonic line. Alternative definitions are also supported by S2LET, as presented at the end of this section. Consider the C^{∞} Schwartz function with compact support on [-1, 1]:

$$s(t) \equiv \begin{cases} e^{-\frac{1}{1-t^2}}, & t \in [-1,1] \\ 0, & t \notin [-1,1] \end{cases},$$
(5.10)

for $t\in\mathbb{R}.$ We introduce the positive real parameter $\lambda\in\mathbb{R}^+_*$ to map s(t) to

$$s_{\lambda}(t) \equiv s\left(\frac{2\lambda}{\lambda - 1}(t - 1/\lambda) - 1\right),\tag{5.11}$$

which has compact support in $[1/\lambda, 1]$. We then define the smoothly decreasing function k_{λ} by

$$k_{\lambda}(t) \equiv \frac{\int_{t}^{1} \frac{\mathrm{d}t'}{t'} s_{\lambda}^{2}(t')}{\int_{1/\lambda}^{1} \frac{\mathrm{d}t'}{t'} s_{\lambda}^{2}(t')},\tag{5.12}$$

which is unity for $t < 1/\lambda$, zero for t > 1, and is smoothly decreasing from unity to zero for $t \in [1/\lambda, 1]$. We finally define the wavelet generating function by

$$\kappa_{\lambda}(t) \equiv \sqrt{k_{\lambda}(t/\lambda) - k_{\lambda}(t)}$$
(5.13)

and the scaling function generating function by

$$\eta_{\lambda}(t) \equiv \sqrt{k_{\lambda}(t)}.$$
(5.14)

The wavelets and scaling function are constructed from their generating functions to satisfy the admissibility condition given by Eqn. (6.35). A natural approach is to define $\Psi_{\ell m}^{j}$ from the generating functions κ_{λ} to have support on $[\lambda^{j-1}, \lambda^{j+1}]$, yielding

$$\Psi_{\ell m}^{j} \equiv \sqrt{\frac{2\ell+1}{4\pi}} \kappa_{\lambda} \left(\frac{\ell}{\lambda^{j}}\right) \delta_{m0}.$$
(5.15)

For these wavelets Eqn. (6.35) is satisfied for $\ell \ge \lambda^{J_0}$, where J_0 is the lowest wavelet scale used in the decomposition. The scaling function Φ is constructed to extract the modes that cannot be probed by the wavelets (i.e. modes with $\ell < \lambda^{J_0}$):

$$\Phi_{\ell m} \equiv \sqrt{\frac{2\ell+1}{4\pi}} \,\eta_{\lambda} \left(\frac{\ell}{\lambda^{J_0}}\right) \delta_{m0}. \tag{5.16}$$

To satisfy exact reconstruction, J is set to ensure the wavelets reach the band-limit of the signal of interest, yielding $J = \lceil \log_{\lambda}(L-1) \rceil$. The choice of the lowest wavelet scale J_0 is arbitrary, provided that $0 \le J_0 < J$. The wavelets and scaling function may then be reconstructed on the sphere through an inverse spherical harmonic transform. The harmonic tiling and real space representation of these wavelets are shown in Figure 5.1 and Figure 5.2 respectively.

In addition to the scale-discretised generating functions (Wiaux et al., 2008), S2LET also supports the needlet functions (Marinucci et al., 2008)⁷, which yield a similar tiling of the harmonic line, as shown in Figure 5.1. The B-spline filters used to construct the isotropic undecimated wavelet

⁷In our implementation of needlets we use a scaling function to represent the approximate information in the signal, which is not always included (e.g., NeedAtool; Pietrobon et al. 2010).





(b) Angular profiles of the scaling function and the first wavelets

Figure 5.1: Wavelets and scaling function constructed with the scale-discretised (SD), needlet and B-spline generating functions (Marinucci et al., 2008; Starck et al., 2006a; Wiaux et al., 2008) with parameters $\lambda = 3$ and $J_0 = 2$ and for band-limit L = 128. The tiling is shown in the top panel, and the profiles of the reconstructed wavelets in the bottom panel.

transform (Starck et al., 2006a) are also supported, as also shown in Figure 5.1.⁸ With these three constructions, the wavelets and scaling functions are well-localised both spatially on the sphere and also in harmonic space. Consequently, the associated wavelet transforms on the sphere can be used to extract spatially localised, scale-dependent features in signals of interest.

$$k_{\lambda}(t) = \frac{3}{2} B_3(2 \frac{t \lambda^{J-1}}{L})$$
(5.17)

$$B_3(x) = \frac{1}{12}(|x-2|^3 - 4|x-1|^3 + 6|x|^3 - 4|x+1|^3 + |x+2|^3),$$
(5.18)

 $^{^{8}}$ For the B-spline-based construction to probe approximately the same scales as the scale-discretised and needlet ones, we defined the generating functions as

so that the *j*th filter has (compact) support $[0, L/\lambda^{J-j-2}]$ and peaks at the same scales as the *j*-th scale-discretised and needlet filters obtained with the same parameters.



Figure 5.2: Wavelets for scales $j \in \{2, 3, 4, 5\}$ and scaling function constructed through a tiling of the harmonic line using scale-discretised functions, with parameters $\lambda = 3$ and $J_0 = 2$ and for band-limit L = 128. This plot was produced with a Matlab demo included in S2LET.

5.4. The S2LET code

In this section we describe the S2LET code. We first introduce a multiresolution algorithm to capture each wavelet scale in the minimum number of samples on the sphere, which follows by taking advantage of the reduced band-limit of the wavelets for scales j < J - 1. This multiresolution algorithm reduces the computation cost of the transform considerably. We then provide details of the implementation, the computational complexity and the numerical accuracy of the scale-discretised wavelet transform supported in S2LET. We finally outline planned future extensions of the code.

5.4.1. Multiresolution algorithm

In harmonic space, the wavelet coefficients are simply given by the weighted product of the spherical harmonic coefficients of f and the wavelets, as expressed in Eqn. (6.30). Although the wavelet coefficients can be analysed at the same resolution as the signal f (i.e., at full resolution), by construction they have different band-limits for different scales j, as shown in Figure 5.1. The reconstruction can thus be performed at lower resolution, without any loss of information if a sampling theorem is used. This approach yields a multiresolution algorithm where the wavelet coefficients are reconstructed with the minimal number of samples on the sphere: the j-th wavelet

coefficients have band-limit $k = \lambda^{j+1}$ when using the scale-discretised and needlet kernels, and $k = L/\lambda^{J-j-2}$ when using the B-splines. When the MW sampling theorem is used, the wavelets are recovered on (k - 1)(2k - 1) + 1 samples on the sphere. This approach leads to significant improvements in terms of speed and memory use compared to the full-resolution case, as shown in the next section. Figure 5.3 illustrates the use of the full-resolution and multiresolution transforms on a map of Earth topography data with the scale-discretised filters and the MW scheme. When adopting the HEALPix sampling of the sphere, multiresolution can also be used. However HEALPix does not rely on a sampling theorem and therefore the resolution for the reconstruction of each wavelet scale must be chosen heuristically and adapted to the desired accuracy. For example, in the MRS code (Starck et al., 2006a) it is chosen such that $N_{side}^{j} = k/2$. More detail on the accuracy of the wavelet transform with HEALPix are provided below.



(a) Full-resolution scale-discretised wavelet transform

(b) Multiresolution scale-discretised wavelet transform

Figure 5.3: Scale-discretised wavelet transform of a band-limited topography map of the Earth for $\lambda = 3$, $J_0 = 2$ and L = 128, i.e. with the scale-discretised wavelets shown in Figure 5.2. The wavelet transform decomposes the band-limited signal into wavelet coefficients that extract spatially localised, scale-dependent features. Since the wavelets for different scales j have different band-limits, the wavelet coefficients can be reconstructed at lower resolution on the sphere for lower scales j. Panel (a) shows the full-resolution wavelet transform of the topography map. The original Earth topography map is shown in the top-left plot, the scaling coefficients are shown in the top-right plot, while the wavelet coefficients at scales $j \in \{2, 3, 4, 5\}$ are shown left-to-right, top-to-bottom respectively in the remaining plots. Panel (b) shows the same decomposition but using the multiresolution algorithm. The signals shown in panel (b) contain the same information as in panel (a) but represented in the minimal number of samples on the sphere. These plots were produced by one of the many Matlab demos provided with S2LET.

5.4.2. Implementation

The core numerical routines of S2LET are implemented in C. By adopting a low level programming language such as C for the implementation of the core algorithms, computational efficiency is optimised. The C library includes the full-resolution and multiresolution wavelet transforms, with specific optimisations for real signals in order to take advantage of all symmetries of the spherical harmonic transform. The wavelet transform is computed in harmonic space through Eqn. (6.30) and Eqn. (6.32), for the input parameters (L, λ, J_0) . To reconstruct signals on the sphere, by default S2LET uses the exact spherical harmonic transform of the MW sampling theorem (McEwen and Wiaux, 2011) implemented in the SSHT⁹ code. In this case all transforms are theoretically exact and one can analyse and synthesise real and complex signals at floating-point precision. S2LET has been extended to also support the HEALPix sampling scheme, in which case the transform is not theoretically exact but nevertheless achieves good numerical accuracy.

We provide interfaces for the C library in three languages: Matlab, IDL and Java. The Matlab and IDL codes also include routines to read/write signals on the sphere stored in either HEALPix FITS¹⁰ files or the FITS file format used to stored MW sampled signals. In addition, functionality to plot the Mollweide projection of real signals for both MW or HEALPix samplings is included. The Java interface includes an object-oriented representation of sampled maps, spherical harmonics and wavelet transforms. All routines and interfaces are well documented and illustrated with several examples for both the MW and HEALPix samplings. These examples cover multiple combinations of parameters and types of signals. S2LET requires SSHT, which implements fast and exact algorithms to perform the forward and inverse spherical harmonic transforms corresponding to the MW sampling theorem (McEwen and Wiaux, 2011). SSHT in turn requires the FFTW¹¹ package for the computation of fast Fourier transforms. The fast spherical harmonic functions, through efficient recursion using either the method of Trapani and Navaza (2006) or Risbo (1996). Here we present results using the recursion of Risbo (1996). The fast spherical harmonic transform algorithms implemented in SSHT scale as $O(L^3)$ (McEwen and Wiaux, 2011).

Although primarily intended to perform the scale-discretised wavelet transform of Wiaux et al. (2008), S2LET also supports the needlet and spline-based wavelet transforms developed by Marinucci

⁹http://www.spinsht.org/

¹⁰http://fits.gsfc.nasa.gov/

¹¹http://www.fftw.org/

et al. (2008) and Starck et al. (2006a). As shown in Figure 5.1, these generating functions yield the same number of wavelet scales (for the parameter choices described previously). However, with the scale-discretised and needlet generating functions the *j*-th wavelet scale has compact support in $[\lambda^{j-1}, \lambda^{j+1}]$, whereas the support is much wider with the B-splines, i.e. $[0, L/\lambda^{J-j-2}]$ in the S2LET implementation. As a consequence, when using the multiresolution algorithm the wavelet coefficients must be captured on a greater number of pixels than with the scale-discretised or needlet kernels, while probing approximately the same scales, as shown in Figure 5.1.

The complexity of the axisymmetric wavelet transform is dominated by spherical harmonic transforms since the wavelet transforms are computed efficiently in harmonic space, through Eqn. (6.30) and Eqn. (6.32) for the forward transform and through Eqn. (6.35) for the inverse transform. Given a band-limit L and wavelet parameters (λ, J_0) , recall that the maximum scale is given by $J = \lceil \log_{\lambda}(L-1) \rceil$ and hence the wavelet transform (forward or inverse) involves $(J - J_0 + 3)$ spherical harmonic transforms (one for the original signal, one for the scaling coefficients and $(J - J_0 + 1)$ for the wavelet coefficients). If the scaling coefficients and all wavelet coefficients are reconstructed at full-resolution in real space, the axisymmetric wavelet transform scales as $\mathcal{O}((J - J_0 + 3)L^3)$. However, in the previous section we established a multiresolution algorithm that takes advantage of the reduced band-limit of the wavelets for scales $j \in \{J - 1, J\}$ are computed at maximal resolution corresponding to the band-limit of the signal. The complexity of the overall multiresolution wavelet transform is then dominated by these operations and effectively scales as $\mathcal{O}(L^3)$.

5.4.3. Numerical validation

We first evaluate the performance of S2LET in terms of accuracy and complexity using the MW sampling theorem, for which all transforms are theoretically exact. We show that S2LET achieves floating-point precision and scales as detailed in the previous section.

We consider band-limits $L = 2^i$ with $i \in \{2, ..., 10\}$ and generate sets of spherical harmonic coefficients $f_{\ell m}$ following independent Gaussian distributions $\mathcal{N}(0, 1)$. We then perform the wavelet decomposition and reconstruct the harmonic coefficients, denoted by $f_{\ell m}^{\text{rec}}$. We evaluate the accuracy of the transform using the error metric $\epsilon = \max |f_{\ell m} - f_{\ell m}^{\text{rec}}|$, which is theoretically zero since all signals are band-limited by construction. The complexity is quantified by observing how the



Figure 5.4: Numerical accuracy and computation time of the scale-discretised wavelet transform computed with S2LET. We consider $L = 2^i$ with $i \in \{2, ..., 10\}$, with parameters $\lambda = 2$, $J_0 = 0$. These results are averaged over many realisations of random band-limited signals and were found to be very stable. The scale-discretised transform is either performed at full-resolution (solid lines) or with the multiresolution algorithm (dashed lines). Very good numerical accuracy is achieved by both the full-resolution and multiresolution algorithms (which achieve indistinguishable accuracy), with numerical errors comparable to floating-point precision, found empirically to scale as O(L) as shown by the red line in panel (a). Computation time scales as $O(L^3)$ for both algorithms as shown by the red line in panel (b), in agreement with theory. The multiresolution algorithm is four to five times faster than the full-resolution approach for the band-limits considered.

computation time $t_c = [t_{synthesis} + t_{analysis}]/2$ scales with band-limit, where the synthesis and analysis computation times are specified by $t_{\text{synthesis}}$ and t_{analysis} respectively. Since we evaluate the wavelet transform in real space, a preliminary step is required to reconstruct the signal f from the randomly generated $f_{\ell m}$. This step is not included in the computation time since its only purpose is to generate a valid band-limited test signal on the sphere. The analysis then denotes the decomposition of f into wavelet coefficients $W^{\Psi^{j}}$ and scaling coefficients W^{Φ} on the sphere. The synthesis refers to recovering the signal f^{rec} from these coefficients. The final step, which is not included in the computation time either, is to decompose $f^{\rm rec}$ into harmonic coefficients $f_{\ell m}^{\rm rec}$ in order to compare them with $f_{\ell m}$. The stability of both ϵ and t_c is checked by averaging over hundreds of realisations of $f_{\ell m}$ for $L = 2^i$ with $i \in \{2, ..., 8\}$ and a few realisations with $i \in \{9, 10\}$. The results proved to be very stable, i.e. the variances of the error and timing metrics are lower than 5%. Recall that for given band-limit L the number of samples on the sphere required by the exact quadrature is (2L-1)(L-1) + 1. All tests were run on an Intel 2.0GHz Core i7 processor with 8GB of RAM. On this machine, precision of floating point numbers is of the order of $\sim 10^{-16}$, and errors are expected to add up and accumulate when considering linear operations such as the spherical harmonic and wavelet transforms. The accuracy and timing performance of the scale-discretised wavelet transform implemented in S2LET are presented in Figure 5.4. S2LET achieves very good numerical accuracy, with numerical errors comparable to accumulated floating-point errors only¹². Moreover, the full-resolution and multiresolution algorithms are indistinguishable in terms of accuracy. However, the latter is four to five times faster than the former for the band-limits considered since only the wavelet coefficients for $j \in \{J - 1, J\}$ are computed at full-resolution. As shown in Figure 5.4, computation time scales as $\mathcal{O}(L^3)$ for both algorithms, in agreement with theory.

S2LET can also be used with HEALPix, in which case the accuracy of the spherical harmonic transform is critical to the accuracy of the wavelet transform (since HEALpix does not rely on a sampling theorem it does not exhibit theoretically exact harmonic transforms, unlike SSHT or GLESP). The performances of the spherical harmonic transforms in HEALPix and GLESP have been widely studied in the past (see, e.g., Doroshkevich et al. 2011; Reinecke 2011; Reinecke and Seljebotn 2013), and that of the MW sampling were presented in McEwen and Wiaux (2011). We do not compile the entirety of these results here, but we have reproduced the essential results on our machine; Table 5.1 summarises the orders of accuracy of the HEALPix iterative spherical harmonic transform. Using the same setup as previously, we calculated the maximum error on the spherical

¹²The GLESP sampling adopted in MRS also achieves floating point accuracy, although using many more pixels to capture the wavelet scales due to the greater band-limits of the spline-based kernels and the oversampling of the GLESP scheme.

$\max f_{\ell m} - f_{\ell m}^{\rm rec} $	$L={\rm N_{side}}/2$	$L = N_{side}$	$L = 2 \mathcal{N}_{\text{side}}$	$L = 3 \mathrm{N}_{\mathrm{side}}$
0 iteration	$\sim 10^{-6}$	$\sim 10^{-4}$	$\sim 10^{-2}$	$\sim 10^{-1}$
1 iterations	$\sim 10^{-10}$	$\sim 10^{-7}$	$\sim 10^{-3}$	$\sim 10^{-1}$
2 iterations	$\sim 10^{-14}$	$\sim 10^{-10}$	$\sim 10^{-5}$	$\sim 10^{-1}$
3 iterations	$\sim 10^{-14}$	$\sim 10^{-13}$	$\sim 10^{-6}$	$\sim 10^{-1}$
4 iterations	$\sim 10^{-14}$	$\sim 10^{-14}$	$\sim 10^{-7}$	$\sim 10^{-1}$

Table 5.1: Order of magnitude of the accuracy of the HEALPix spherical harmonic transform, averaged over the parameter $N_{\rm side}$.

harmonic coefficients when performing the transform back and forth, averaged over the values of N_{side} , since the results were found to be sensitive only to the ratio L/N_{side} . Even with several iterations, which multiplies the number of transforms and thus computation time, the spherical harmonic transform in HEALPix remains at least an order of magnitude less accurate than the MW and GLESP counterparts (which, being both theoretically exact, achieve comparable performances, see Doroshkevich et al. 2011; McEwen and Wiaux 2011; Reinecke 2011; Reinecke and Seljebotn 2013). Since the wavelet transforms implemented in MRS and Needatool are also computed in harmonic space, their complexity and accuracy are dominated by that of the underlying spherical harmonic transforms. As a consequence, when adopting the HEALPix scheme, S2LET, MRS and Needatool achieve similar performances, resulting from the computation time and accumulated errors of $(J - J_0 + 1)$ HEALPix spherical harmonic transforms. In the multiresolution case, the results depend on the resolution chosen to reconstruct each wavelet scale.

5.4.4. Future extensions

In future work we plan to extend S2LET to support directional, steerable wavelets on the sphere (Wiaux et al., 2008). We also plan to exploit recent ideas leading to fast (spin) spherical harmonic transforms (McEwen and Wiaux, 2011) to yield faster algorithms than those developed by Wiaux et al. (2008) and McEwen et al. (2013) to compute directional wavelet transforms on the sphere. Finally, we intend to add support to analyse spin signals on the sphere (c.f. Geller and Marinucci, 2010; Geller et al., 2008; Starck et al., 2009). In a future release, the code will also be parallelised, which will lead to further speed improvements. The S2LET code will thus be under active development with future releases forthcoming. In any case, we hope this first version of the S2LET code will prove useful for axisymmetric scale-discretised wavelet analysis on the sphere. Indeed, the code has already been used as an integral part of the new exact flaglet wavelet transform

on the ball (Leistedt and McEwen, 2012), the spherical space constructed by augmenting the sphere with the radial line.

5.5. Examples

The S2LET code is extensively documented and ships with several examples in the four languages supported. In this section we present a subset of short examples, along with the code to execute them in order to demonstrate the ease of using S2LET to perform wavelet transforms¹³. All examples were run with the scale-discretised wavelet generating functions.

5.5.1. Wavelet transform from the command line

S2LET includes ready-to-use high-level programs to directly decompose a real signal into wavelet coefficients. The inputs are a FITS file containing the signal of interest and the parameters for the transform. The program writes the output coefficients in FITS files in the same directory as the input file and with a consistent naming scheme. These commands are available for both HEALPix and MW sampling schemes. For the MW sampling case illustrated in Example 5.1, the wavelet transform is theoretically exact and the band limit corresponds to the resolution of the input map, which will be read automatically from the file. The transform may be performed in full-resolution or multiresolution by adjusting the multiresolution flag specified by the last parameter (respectively 0 and 1), and the output wavelet coefficients are computed at full and minimal resolution accordingly. For the case of a HEALPix map, as illustrated in Example 5.2, the band-limit must be supplied as the last parameter in the command. The output scaling and wavelet coefficients of a HEALPix map are reconstructed and stored in FITS files at the same resolution as the input map. For both MW and HEALPix samplings the output coefficients may be read and plotted using the Matlab or IDL routines.

Example 5.1: Performing the forward (analysis) and inverse (synthesis) wavelet transform of a real signal (MW sampling) from the command line.

¹³Note that the code uses a slightly different notation compared to the equations of this article: *B* refers to the wavelet scaling parameter (denoted λ herein) and J_{\min} to the first scale of the transform (denoted J_0 herein).

Example 5.2: Performing the forward (analysis) and inverse (synthesis) wavelet transform of a real signal (HEALPix sampling) from the command line.

5.5.2. Wavelet transform in Matlab and IDL

Examples 5.3 and 5.4 read real signals on the sphere from FITS files, calculate the wavelet coefficients and plot them using a Mollweide projection. The first case is a Matlab example where the input map is a simulation of the cosmic microwave background in the HEALPix sampling. The second case is a IDL example where the input map is a topography map of the Earth in MW sampling. S2LET ships with versions of these two examples in C, Matlab and IDL.

5.5.3. Wavelet denoising in C

Example 5.5 illustrates the use of the wavelet transform to denoise a signal on the sphere. The input noisy map is a band-limited topography map of the Earth in MW sampling at resolution L = 128. It is read from a FITS file, decomposed into wavelet coefficients (for given parameters λ and J_0) which are then denoised by thresholding. The denoised signal is reconstructed from the denoised wavelet coefficients and written to a FITS file.

In this example we consider a noisy signal $y = s + n \in L^2(S^2)$, where the signal of interest $s \in L^2(S^2)$ is contaminated with noise $n \in L^2(S^2)$. We consider zero-mean white Gaussian noise on the sphere, where the variance of the harmonic coefficients of the noise is specified by

$$\mathbb{E}\left(|n_{\ell m}|^2\right) = \sigma^2, \quad \forall \ell, m. \tag{5.19}$$

A simple way to evaluate the fidelity of the observed signal y is through the signal-to-noise ratio (SNR), define on the sphere by

$$SNR(y) \equiv 10 \log_{10} \frac{\|s\|_2^2}{\|y - s\|_2^2},$$
(5.20)

```
% Example: Wavelet transform in Matlab
lambda = 3; J0 = 2; L = 192;
Jmax = s2let_jmax(L, lambda);
\% Read a real HEALPix map from a FITS file
inputfile = 'data/somecmbsimu_hpx_128.fits';
[f, nside] = s2let_hpx_read_real_map(inputfile);
% Perform the wavelet transform
[f_wav, f_scal] = s2let_axisym_hpx_analysis (f,'B',lambda,'L',L,'J_min',J0);
% Plot the map and the wavelet coefficients
figure; ns = ceil(sqrt(2+Jmax-J0+1));
subplot(ns, ns, 1);
s2let_hpx_plot_mollweide(f);
title('Initial band-limited data')
subplot(ns, ns, 2);
s2let_hpx_plot_mollweide(f_scal);
title('Scaling fct')
for j = J0:Jmax
  subplot(ns, ns, j-J0+3);
   s2let_hpx_plot_mollweide(f_wav{j-J0+1});
   title(['Wavelet scale : ',int2str(j)-J0+1])
end
```

Example 5.3: Performing the wavelet transform of a real signal (HEALPix sampling) using the Matlab interface.

```
; Example: Wavelet transform in IDL
lambda = 3
J0 = 2
; Read a real MW map from a FITS file
file = 'data/earth_tomo_mw_128.fits'
f = s2let_mw_read_real_map(file)
L = s2let_get_mw_bandlimit(f)
Jmax = s2let_j_max(L, lambda)
; Perform the wavelet transform
f_wav = s2let_axisym_mw_wav_analysis_real (f, lambda, J0)
f_rec = s2let_axisym_mw_wav_synthesis_real (f_wav)
; Plot the map and the wavelet coefficients
ns = ceil(sqrt(3+Jmax-J0))
!P.MULTI=[0,ns,ns]
s2let_mw_plot_mollweide, f_rec, title='Band-limited map'
s2let_mw_plot_mollweide, f_wav.scal, title='Scaling map'
for j=0, Jmax-J0 do begin
   s2let_mw_plot_mollweide, f_wav.(j), title='Wavelet map '+strtrim(j+1,2)+'
       on '+strtrim(Jmax-J0+1,2)
endfor
!P.MULTI=0
```

Example 5.4: Performing the wavelet transform of a real signal (MW sampling) using the IDL interface.

where the signal energy is defined by

$$\|y\|_{2}^{2} \equiv \langle y|y\rangle = \int_{S^{2}} d\Omega(\omega)|y(\omega)|^{2} = \sum_{\ell m} |y_{\ell m}|^{2}.$$
 (5.21)

We seek a denoised version of y, denoted by $d \in L^2(S^2)$, with large SNR(d) so that d isolates the informative signal s. When taking the wavelet transform of the noisy signal y, one expects the energy of the informative part to be concentrated in a small number of wavelet coefficients, whereas the noise energy should be spread over various wavelet scales. In this particular toy example, the signal has significant power on large scales, as shown in Figure 5.3, which are well described in the wavelet basis and less affected by the random white noise. Since the transform is linear, the wavelet coefficients of the j-th scale are simply given by the sum of the individual contributions:

$$Y^{j}(\omega) = S^{j}(\omega) + N^{j}(\omega), \qquad (5.22)$$

where capital letters denote the wavelet coefficients, i.e. $Y^j \equiv y \star \Psi^j$, $S^j \equiv s \star \Psi^j$ and $N^j \equiv n \star \Psi^j$. For the zero-mean white Gaussian noise defined by Eqn. 6.46, the noise in wavelet space is also zero-mean and Gaussian, with variance

$$\mathbb{E}\left(|N^{j}(\omega)|^{2}\right) = \sigma^{2} \sum_{\ell} |\Psi_{\ell 0}^{j}|^{2} \equiv \left(\sigma^{j}\right)^{2}.$$

Denoising is performed by hard-thresholding the wavelet coefficients Y^j , where the threshold is taken as $T^j = 3\sigma^j$. The denoised wavelet coefficients $D^j \equiv d \star \Psi^j$ are thus given by

$$D^{j}(\omega) = \begin{cases} 0, & \text{if } Y^{j}(\omega) < T^{j}(\omega) \\ Y^{j}(\omega), & \text{otherwise} \end{cases}$$
(5.23)

The denoised signal $d \in L^2(S^2)$ is reconstructed from its wavelet coefficients D^j and the scaling coefficients of y, which are not thresholded. The denoising procedure outlined above is particularly simple and more sophisticated denoising strategies can be developed; we adopt this simple denoising strategy merely to illustrate the use of the S2LET code. In this example we perform the wavelet transform with parameters $\lambda = 2$ and $J_0 = 0$. For a noisy signal y with SNR(y) = 11.78dB, the scale-discretised wavelet denoising recovers a denoised signal d with SNR(d) = 14.66dB. The initial, noisy and denoised maps are shown in Figure 5.5. When switching to needlets and B-spline wavelets while keeping λ and J_0 unchanged, the denoised signals have SNR(d) = 14.68dB and

5.6. Summary

In the era of precision astrophysics and cosmology, large and complex data-sets on the sphere must be analysed at high precision in order to confront accurate theoretical predictions. Scalediscretised wavelets are a powerful analysis technique where spatially localised, scale-dependent signal features of interest can be extracted and analysed. Combined with a sampling theorem, this framework leads to an exact multiresolution wavelet analysis, where signals on the sphere can be reconstructed from their scaling and wavelet coefficients exactly.

We have described S2LET, a fast and robust implementation of the scale-discretised wavelet transform. Although the first public release of S2LET is restricted to axisymmetric wavelets, the generalisation to directional, steerable wavelets will be made available in a future release. The core numerical routines of S2LET are written in C and have interfaces in Matlab, IDL and Java. Both MW and HEALPix pixelisation schemes are supported. In this article we have presented a number of examples to illustrate the ease of use of S2LET for performing wavelet transform of real signals stored as FITS files and to plot scaling and wavelet coefficients on Mollweide projections of the sphere. We have also detailed a denoising example where denoising is performed through simple hard-thresholding in wavelet space. Although only a simple denoising strategy was performed to illustrate the use of the S2LET code, it nevertheless performed very well, highlighting the effectiveness of the scale-discretised wavelet transform on the sphere.

5.7. Acknowledgements

BL is supported by the Perren Fund and the Impact Fund. JDM is supported in part by a Newton International Fellowship from the Royal Society and the British Academy. YW is supported by the Center for Biomedical Imaging (CIBM) of the Geneva and Lausanne Universities, EPFL and the Leenaards and Louis-Jeantet foundations.



(a) Band-limited signal



(b) Noisy signal with SNR(y) = 11.8 dB



(c) Denoised signal with ${\rm SNR}(d)=14.66{\rm dB}$

Figure 5.5: Wavelet denoising by hard-thresholding, using parameters $\lambda = 2$ and $J_0 = 0$ and scale-discretised generating functions. When using needlets and B-spline wavelets, the denoised signals have SNR(d) = 14.68dB and 14.46dB respectively. This example is included in S2LET as a documented demo program.

```
// Example: Wavelet denoising in C
int lambda = 2, JO = O;
// Read a real MW map from a FITS file % \mathcal{M} = \mathcal
char inputfile[100] = "..."
double *f;
int L = s2let_fits_read_mw_bandlimit(file);
s2let_mw_allocate_real(&f, L);
s2let_fits_read_mw_map(f, file, L);
// Perform multiresolution wavelet analysis
double *f_wav, *f_scal;
s2let_axisym_mw_allocate_f_wav_multires_real (&f_wav, &f_scal, lambda, L, J0);
s2let_axisym_mw_wav_analysis_multires_real (f_wav, f_scal, g, lambda, L, J0);
// Threshold the wavelets with a noise model
s2let_axisym_wav_hardthreshold_multires_real (f_wav, threshold, lambda, L, J0);
// Reconstruct the denoised signal
double *f_denoised;
s2let_mw_allocate_real(&f_denoised, L);
s2let_axisym_mw_wav_synthesis_multires_real (f_denoised, f_wav, f_scal,
                    lambda,L,J0);
// Write the denoised signal
char outputfile[100] = "..."
s2let_fits_write_mw_map(outfile,f_denoised,L);
```

Example 5.5: Denoising a real signal (MW sampling) in C through hard-thresholding of the wavelet coefficients.

Exact Wavelets on the Ball

"When you change the way you look at things, the things you look at change". M. Planck

6.1. Abstract

We develop an exact wavelet transform on the three-dimensional ball (i.e. on the solid sphere), which we name the *flaglet* transform. For this purpose we first construct an exact transform on the radial half-line using damped Laguerre polynomials and develop a corresponding quadrature rule. Combined with the spherical harmonic transform, this approach leads to a sampling theorem on the ball and a novel three-dimensional decomposition which we call the Fourier-Laguerre transform. We relate this new transform to the well-known Fourier-Bessel decomposition and show that bandlimitedness in the Fourier-Laguerre basis is a sufficient condition to compute the Fourier-Bessel decomposition exactly. We then construct the flaglet transform on the ball through a harmonic tiling, which is exact thanks to the exactness of the Fourier-Laguerre transform (from which the name flaglets is coined). The corresponding wavelet kernels are well localised in real and Fourier-Laguerre spaces and their angular aperture is invariant under radial translation. We introduce a multiresolution algorithm to perform the flaglet transform rapidly, while capturing all information at each wavelet scale in the minimal number of samples on the ball. Our implementation of these new tools achieves floating-point precision and is made publicly available. We perform numerical experiments demonstrating the speed and accuracy of these libraries and illustrate their capabilities on a simple denoising example.

6.2. Introduction

A common problem in data analysis is the extraction of non-trivial patterns and structures of interest from signals. This problem can be addressed by projecting the data onto an appropriate basis. Whereas Fourier analysis focuses on oscillatory features, wavelets extract the contributions of scale-dependent features in both real and frequency space simultaneously. Initially defined in Euclidean space, wavelets have been extended to various manifolds and are now widely used in numerous disciplines. In particular, spherical wavelets (Antoine and Vandergheynst, 1998, 1999; Baldi et al., 2009; Marinucci et al., 2008; McEwen et al., 2006; Starck et al., 2006a; Wiaux et al., 2005, 2006, 2008; Yeo et al., 2008) have been extremely successful at analysing data on the sphere and have now become a standard tool in geophysics (*e.g.*, Audet 2011; Audet and Johnson 2011; Charléty et al. 2012; Loris et al. 2010; Simons et al. 2011; Simons et al. 2001; Deriaz et al. 2012; Faÿ et al. 2008; Labatie et al. 2012; Lan and Marinucci 2008; McEwen et al. 2006, 2007a,b, 2008; Pietrobon et al. 2008; Schmitt et al. 2010; Starck et al. 2006b; Vielva et al. 2004, 2006a). Naturally, data may also be defined on the three-dimensional ball when radial information (such as depth, redshift or distance, for example) is associated with each spherical map.

First approaches to perform wavelet-type transforms on the ball were developed by Fengler et al. (2006); Michel (2005) in the continuous setting only, which thus cannot be used for exact reconstruction in practice. The spherical Haar transform (Lessig, 2007; Lessig and Fiume, 2008) was extended to the ball by Chow (2010) to support exact analysis and synthesis. However, this framework is very restrictive and may not necessarily lead to a stable continuous basis Schröder and Sweldens (1995); Sweldens (1996, 1997). The first wavelet transform on the ball to tackle both the continuous and discrete settings was developed in the influential work of Lanusse et al. (2012). This wavelet transform is based on an isotropic undecimated wavelet construction, built on the Fourier-Bessel transform. Since these wavelets are isotropic, their angular aperture depends on the distance to the origin. Although the wavelet transform on the ball is exact in Fourier-Bessel space, wavelet coefficients must be recovered on the ball from their Fourier-Bessel coefficients (in order to extract spatially localised information). However, there exists no exact quadrature formula for the spherical Bessel transform (the radial part of the Fourier-Bessel transform) (Lemoine, 1994), and thus no way to perform the Fourier-Bessel transform exactly. Consequently, the undecimated wavelet transform on the ball is not theoretically exact when wavelet coefficients are recovered on the ball. Nevertheless, the isotropic undecimated wavelet transform does achieve good numerical accuracy, which may be sufficient for many applications.¹ Wavelets on the ball have also been discussed in geophysics by Simons et al. (2011); Simons et al. (2011), who espoused a philosophy of separability in the three Cartesian coordinates of a ball-to-"cubed-sphere-ball" mapping, although in Simons et al. (2011) examples are shown where wavelet transforms have been performed on each spherical shell only but not in the radial direction. In ongoing work, these same authors have extended their approach to the ball, where the wavelet transform in the radial direction is tailored to seismological applications by honouring certain major discontinuities in the seismic wavespeed profile of the Earth (Charléty et al., 2012; Loris et al., 2010). At present, to the best of our knowledge, there does not exist an exact wavelet transform of a band-limited signal defined on the ball.

One reason there is no exact wavelet transform on the ball is due to the absence of an exact harmonic transform. We resolve this issue by deriving an exact spherical Laguerre transform on the radial half-line, leading to a new Fourier-Laguerre transform on the ball which is theoretically exact. Furthermore, this gives rise to a sampling theorem on the ball, where all information of a band-limited function is captured in a finite number of samples. With an exact harmonic transform on the ball in hand, we construct exact wavelets through a harmonic tiling, which we call *flaglets* (since they are built on the Fourier-LAGuerre transform). Each wavelet kernel is localised in real and Fourier-Laguerre spaces, and probes a characteristic angular scale which is invariant under radial translation. Flaglets allow one to probe three-dimensional spherical data in position and scale simultaneously. Moreover, their exactness properties guarantee that the flaglet transform captures and preserves all the information contained in a band-limited signal.

The remainder of this article is organised as follows. In Section 6.3 we define the spherical Laguerre transform on the radial half-line and the Fourier-Laguerre transform on the ball. In Section 6.4 we construct the exact flaglet transform on the ball. In Section 6.5 we present a multiresolution algorithm to compute the flaglet transform and evaluate our algorithms numerically. A simple denoising example is presented in Section 6.6. Concluding remarks are made in Section 6.7.

6.3. Harmonic Analysis on the Ball

¹The accuracy of the Fourier-Bessel transform, and thus the isotropic undecimated wavelet transform on the ball, may be improved by numerical iteration, although this can prove problematic for certain applications.

The aim of this section is to construct a novel three-dimensional transform which is appropriate for spherical coordinates and admits an exact quadrature formula. For this purpose, we first set out a radial one-dimensional transform inspired by the Laguerre polynomials and we derive a natural sampling scheme and quadrature rule on the radial half-line. We relate this novel spherical Laguerre transform to the spherical Bessel transform and show that the latter can be evaluated exactly if the signal is band-limited in the spherical Laguerre basis. We combine the spherical Laguerre transform with the spherical harmonics to form the Fourier-Laguerre transform on the ball, yielding a novel sampling theorem and an exact harmonic transform.²

6.3.1. The spherical Laguerre transform

The Laguerre polynomials, solutions to the Laguerre differential equation Pollard (1947); Weniger (2008), are well known for their various applications in engineering and physics, notably in the quantum-mechanical treatment of the hydrogen atom Dunkl (2003), as well as in modern optics Bond et al. (2011); Siegman (1973). They form a natural orthogonal basis on the interval $[0, \infty)$ (*i.e.* non-negative reals \mathbb{R}^+) with respect to an exponential weight function. In this work, since we use this expansion along the radial half-line, we define the spherical Laguerre basis function $K_p(r)$ with $r \in \mathbb{R}^+$ as

$$K_p(r) \equiv \sqrt{\frac{p!}{(p+2)!}} \frac{e^{-r/2\tau}}{\sqrt{\tau^3}} L_p^{(2)}\left(\frac{r}{\tau}\right),$$
(6.1)

where $L_p^{(2)}$ is the *p*-th generalised Laguerre polynomial of order two, defined as

$$L_p^{(2)}(r) \equiv \sum_{j=0}^p \binom{p+2}{p-j} \frac{(-r)^j}{j!},$$
(6.2)

and $\tau \in \mathbb{R}^+$ is a scale factor that adds a scaling flexibility and shall be defined at the end of this section. The basis functions K_p are orthonormal on \mathbb{R}^+ with respect to a radial inner product:

$$\langle K_p | K_q \rangle = \int_{\mathbb{R}^+} \mathrm{d}r r^2 K_p(r) K_q^*(r) = \delta_{pq}.$$
(6.3)

Note that the complex conjugate * is facultative since we use real basis functions. Any square-

²A harmonic transform is typically associated with basis functions which are eigenvalues of the Laplacian operator (e.g. the Fourier transform). In this paper our basis functions on the radial half-line (and thus on the ball) are not solutions of the Laplacian, hence harmonic analysis on the ball is interpreted in a broader sense. Nonetheless, these basis functions form orthonormal transforms and define valid dual spaces. We define band-limited signals to have bounded support in the transform space of these orthogonal basis functions.

integrable real signal $f \in L^2(\mathbb{R}^+)$ may be expanded as

$$f(r) = \sum_{p=0}^{\infty} f_p K_p(r),$$
 (6.4)

for natural $p \in \mathbb{N}$, where f_p is the projection of f onto the p-th basis function:

$$f_p = \langle f | K_p \rangle = \int_{\mathbb{R}^+} \mathrm{d}r r^2 f(r) K_p^*(r).$$
(6.5)

The decomposition follows by the orthogonality and completeness of the spherical Laguerre basis functions: orthonormality is given by Eqn. (6.3), while the completeness relation is obtained by applying the Gram-Schmidt orthogonalisation process to the basis functions and exploiting the completeness of polynomials on $L^2(\mathbb{R}^+, r^2e^{-r}dr)$.

When it comes to calculating the transform, one must evaluate the integral of Eqn. (6.5) numerically. We consider functions f band-limited at P in the spherical Laguerre basis, such that $f_p = 0$, $\forall p \ge P$. It is straightforward to show that if f is band-limited, then both $e^{r/2\tau}K_p(r)$ and $e^{r/2\tau}f(r)$ are polynomials of maximum degree P - 1. In this case, Eqn. (6.5) is the integral of a polynomial of order 2P - 2 on \mathbb{R}^+ with weight function r^2e^{-r} . Thus, applying Gaussian quadrature (*e.g.*, Graça and Esmeralda Sousa-Dias 2012; Watson 2007) with P sampling nodes is sufficient to evaluate this integral exactly. The resulting quadrature formula is known as the Gauss-Laguerre quadrature and is commonly used to evaluate numerical integrals on \mathbb{R}^+ . Hence, Eqn. (6.5) reduces to a weighted sum:

$$f_p = \sum_{i=0}^{P-1} w_i f(r_i) K_p^*(r_i),$$
(6.6)

where $r_i \in \mathbb{R}^+$ is the *i*-th root of the *P*-th generalised Laguerre polynomial of order two, and

$$w_i = \frac{(P+2)r_i e^{r_i}}{(P+1)[L_{P+1}^{(2)}(r_i)]^2} \in \mathbb{R}^+$$
(6.7)

is the corresponding weight. Any *P*-band-limited function f can be decomposed and reconstructed exactly using the spherical Laguerre transform. All information content of the function is captured in *P* samples located in the interval $[0, r_{P-1}]$ where r_{P-1} is the largest root of the sampling. Since r_{P-1} increases with *P*, one may wish to rescale the sampling so that the spherical Laguerre transform contains samples in any interval of interest [0, R], with $R \in \mathbb{R}^+$, while the underlying continuous function is nevertheless defined on \mathbb{R}^+ . The scale factor τ is then chosen such that $\tau = R/r_{P-1}$. Figure 6.1 shows the resulting spherical Laguerre sampling constructed on $r \in [0, 1]$ (*i.e.* rescaled


Figure 6.1: Spherical Laguerre sampling scheme on $r \in [0,1]$ for increasing band-limit P. If a function f is P-band-limited then f and the basis functions need only be evaluated on P points for the spherical Laguerre transform to be exact. For a particular P, the associated sampling is denser near the origin since the quadrature is constructed on \mathbb{R}^+ with measure $e^{-r}dr$.

with τ) for increasing band-limit *P*. Figure 6.2 shows the first six basis functions constructed on $r \in [0, 1]$ and the sampling nodes used to obtain an exact transform.³ Note that the spherical Laguerre transform is a real transform that can be extended to complex signals by considering the real and imaginary parts separately.

6.3.2. Relation to the spherical Bessel transform

The spherical Bessel transform is a fundamental radial transform arising from the resolution of the Laplacian operator in spherical coordinates. It is central to the Fourier-Bessel transform, commonly used in cosmology (Abramo et al., 2010; Leistedt et al., 2012; Rassat and Refregier, 2012) to analyse the spectral properties of galaxy surveys in three dimensions. In this section we derive an analytical formula to exactly compute the spherical Bessel transform of a function whose

³If one preferred to consider the measure dr rather than the spherical measure $r^2 dr$, then the basis functions $rK_p(r)$ shown in Figure 6.2 (c) could be used in place of the spherical Laguerre basis functions defined here.



Figure 6.2: First six spherical Laguerre basis functions $K_p(r)$ constructed on $r \in [0,1]$ and the associated sample positions (circles). A function f with band-limit P = 6 can be decomposed and reconstructed exactly using these six basis functions only. In that case, f and the basis functions are solely evaluated at the sampling points. Functions $rK_p(r)$ can be viewed as basis functions in cartesian coordinates satisfying the usual orthogonality relation $\int_{\mathbb{R}^+} dr(rK_p(r))(rK_q(r)) = \delta_{pq}$.

spherical Laguerre transform is band-limited. This section is optional to the reader interested in wavelets only.

The spherical Bessel transform of $f\in L^2(\mathbb{R}^+)$ reads

$$\tilde{f}_{\ell}(k) = \langle f|j_{\ell} \rangle = \sqrt{\frac{2}{\pi}} \int_{\mathbb{R}^+} \mathrm{d}r r^2 f(r) j_{\ell}^*(kr),$$
(6.8)

for $k \in \mathbb{R}^+$, $\ell \in \mathbb{N}$, and where $j_\ell(kr)$ is the ℓ -th order spherical Bessel function. Note again that the complex conjugate is facultative since the spherical Bessel functions are real. The reconstruction formula is given by

$$f(r) = \sqrt{\frac{2}{\pi}} \int_{\mathbb{R}^+} \mathrm{d}k k^2 \tilde{f}_\ell(k) j_\ell(kr).$$
(6.9)

The spherical Bessel transform is thus symmetric and the problem is reduced to the calculation of a similar inner product for the decomposition and the reconstruction. However, to our knowledge, there exists no method to compute such an integral exactly for a useful class of functions, and finding a quadrature formula for the spherical Bessel functions on \mathbb{R}^+ is a non-trivial issue. Moreover, the use of numerical integration methods does not always guarantee good accuracy because of the oscillatory nature of the spherical Bessel functions.

To find a tractable expression to compute Eqn. (6.8), we first express f by its spherical Laguerre expansion, giving

$$\tilde{f}_{\ell}(k) = \sqrt{\frac{2}{\pi}} \sum_{p} f_p j_{\ell p}(k), \qquad (6.10)$$

which is a finite sum if f is band-limited in spherical Laguerre space. In this expression $j_{\ell p}(k)$ is the projection of K_p onto $j_{\ell}(kr)$, i.e.

$$j_{\ell p}(k) \equiv \langle K_p | j_\ell \rangle = \int_{\mathbb{R}^+} \mathrm{d}r r^2 K_p(r) j_\ell^*(kr).$$
(6.11)

Consequently, the problem of computing the spherical Bessel decomposition of f is recast as evaluating Eqn. (6.10), through the computation of the inner product of Eqn. (6.11). But unlike the initial problem of Eqn. (6.8), $j_{\ell p}(k)$ admits an analytic formula. Starting from the definition of Laguerre polynomials in Eqn. (6.2), one can show that

$$j_{\ell p}(k) = \sqrt{\frac{p!}{(p+2)!}} \sum_{j=0}^{p} c_j^p \mu_{j+2}^{\ell}(k),$$
(6.12)

where the c_i^p satisfy the following recurrence:

$$c_{j}^{p} \equiv \frac{(-1)^{j}}{j!} \binom{p+2}{p-j} = -\frac{p-j+1}{j(j+2)} c_{j-1}^{p}.$$
(6.13)

The functions $\mu_j^\ell(k)$ are the moments of $j_\ell(kr)e^{-\frac{r}{2\tau}}$, *i.e.*

$$\mu_{j}^{\ell}(k) \equiv \frac{1}{\tau^{j-\frac{1}{2}}} \int_{\mathbb{R}^{+}} \mathrm{d}r r^{j} j_{\ell}(kr) e^{-\frac{r}{2\tau}}.$$
(6.14)

From Watson (1995) we find an analytical solution for the latter integral:

$$\mu_{j}^{\ell}(k) = \sqrt{\pi} 2^{j} \tilde{k}^{\ell} \tau^{\frac{3}{2}} \frac{\Gamma(j+\ell+1)}{\Gamma(\ell+\frac{3}{2})}$$

$$\times {}_{2}F_{1}\left(\frac{j+\ell+1}{2}; \frac{j+\ell}{2}+1; \ell+\frac{3}{2}; -4\tilde{k}^{2}\right)$$
(6.15)

where $\tilde{k} = \tau k$ is the rescaled k scale and $_2F_1$ is the Gaussian hypergeometric function. Since either $(j + \ell + 1)/2$ or $(j + \ell)/2$ is a positive integer, the latter reduces to a polynomial of \tilde{k}^2 and it is possible to compute the quantity $j_{\ell p}(k)$ exactly using Eqn. (6.12) to (6.15). Consequently, the inverse spherical Bessel transform $\tilde{f}_{\ell}(k)$ may then be calculated analytically through Eqn. (6.10), which is computed exactly if f is band-limited in the spherical Laguerre basis.

6.3.3. The spherical harmonic transform

Whereas the spherical Laguerre transform is specifically designed for analysing functions on the radial half-line, the spherical harmonic transform is a natural choice for the angular part of a consistent three-dimensional analysis. For a function $f \in L^2(S^2)$ on the two-dimensional sphere, the transform reads

$$f(\omega) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} f_{\ell m} Y_{\ell m}(\omega),$$
(6.16)

where $\omega = (\theta, \phi) \in S^2$ are spherical coordinates of the unit sphere S^2 , with colatitude $\theta \in [0, \pi]$ and longitude $\phi \in [0, 2\pi)$. Thanks to the orthogonality and completeness of the spherical harmonics $Y_{\ell m}(\omega)$, the inverse transform is given by the following inner product on the sphere:

$$f_{\ell m} = \langle f | Y_{\ell m} \rangle = \int_{S^2} \mathrm{d}\omega f(\omega) Y_{\ell m}^*(\omega), \qquad (6.17)$$

with surface element $d\omega = \sin \theta d\theta d\phi$. For a function which is band-limited in this basis at L, *i.e.* $f_{\ell m} = 0$, $\forall \ell \ge L$, the decomposition and reconstruction operations can be performed with a finite summation over the harmonics. This is usually resolved by defining an appropriate sampling theorem on the sphere with nodes $\omega_j = (\theta_j, \phi_j)$, associated with a quadrature formula. Various sampling theorems exist in the literature (Driscoll and Healy, 1994; Healy et al., 1996; McEwen and Wiaux, 2011); the main features of all sampling theorems are (i) the number of nodes required to capture all information in a band-limited signal and (ii) the complexity of the related algorithms to compute forward and inverse spherical harmonic transform), we adopt the McEwen & Wiaux (hereafter MW) sampling theorem (McEwen and Wiaux, 2011) which is equiangular and has the lowest number of samples for a given band-limit L, namely $(L - 1)(2L - 1) + 1 \sim 2L^2$. The corresponding algorithms to compute the spherical harmonic transforms scale as $\mathcal{O}(L^3)$ and are numerically stable to band-limits of at least L = 4096 (McEwen and Wiaux, 2011). Further technical details are provided in Section 6.5.2.

6.3.4. The Fourier-Laguerre transform

We define the Fourier-Laguerre basis functions on $B^3 = \mathbb{R}^+ \times S^2$ as the product of the spherical Laguerre basis functions and the spherical harmonics: $Z_{\ell mp}(\mathbf{r}) = K_p(r)Y_{\ell m}(\omega)$ with the 3D spherical coordinates $\mathbf{r} = (r, \omega) \in B^3$. The orthogonality and completeness of the Fourier-Laguerre basis functions follow from the corresponding properties of the individual basis functions, where the orthogonality relation is given explicitly by the following inner product on B^3 :

$$\langle Z_{\ell m p} | Z_{\ell' m' p'} \rangle = \int_{B^3} d^3 \mathbf{r} Z_{\ell m p} Z_{\ell' m' p'}^*(\mathbf{r})$$

$$= \delta_{\ell \ell'} \delta_{m m'} \delta_{p p'},$$
(6.18)

where $d^3\mathbf{r} = r^2 \sin\theta dr d\theta d\phi$ is the volume element in spherical coordinates. Any three-dimensional signal $f \in L^2(B^3)$ can be decomposed as

$$f(\mathbf{r}) = \sum_{p=0}^{P-1} \sum_{\ell=0}^{L-1} \sum_{m=-\ell}^{\ell} f_{\ell m p} Z_{\ell m p}(\mathbf{r}),$$
(6.19)

with L and P the angular and radial band-limits, respectively, *i.e.* f is such that $f_{\ell mp} = 0, \forall \ell \geq L$,

 $\forall p \geq P$. The inverse relation is given by the projection of f onto the basis functions:

$$f_{\ell m p} = \langle f | Z_{\ell m p} \rangle = \int_{B^3} \mathrm{d}^3 \mathbf{r} f(\mathbf{r}) Z^*_{\ell m p}(\mathbf{r}).$$
(6.20)

The Fourier-Laguerre transform may also be related to the Fourier-Bessel transform using the results of Section 6.3.2.⁴

In practice, calculating the Fourier-Laguerre transform requires the evaluation of the integral of Eqn. (6.20). For this purpose, combining the quadrature rules on the sphere and on the radial halfline leads to a sampling theorem on B^3 . For a signal with angular and radial band-limits L and P, respectively, all of the information content of the signal is captured in $N = P[(2L-1)(L-1)+1] \sim 2PL^2$ samples, yielding an exact Fourier-Laguerre transform on B^3 . The three-dimensional sampling consists of spherical shells, discretised according to a sampling theorem (where here we adopt the MW sampling theorem), located at the nodes of the radial sampling. The radial sampling may furthermore be rescaled to any spherical region of interest $[0, R] \times S^2$ using the parameter τ to dilate or contract the radial quadrature rule.

6.4. Wavelets on the Ball

The exactness of the Fourier-Laguerre transform supports the design of an exact wavelet transform on the ball. In this section we first define a three-dimensional convolution operator on the ball, derived from the convolutions defined on the sphere and on the radial half-line. We then construct flaglets through an exact tiling of Fourier-Laguerre space, leading to wavelet kernels which are spatially localised and form a tight frame. Furthermore, each kernel projects onto an angular scale which is invariant under radial translation. We finally introduce a multiresolution algorithm to compute the flaglet transform and capture the information of each wavelet scale in the minimal number of samples on the ball, while optimising the computational cost of the transform.

$$\tilde{f}_{\ell m}(k) = \sqrt{\frac{2}{\pi}} \sum_p f_{\ell m p} j_{\ell p}(k)$$

⁴The Fourier-Laguerre and the Fourier-Bessel transforms of f are related through

If *f* is band-limited in terms of its Fourier-Laguerre decomposition, the latter sum is finite and both transforms can be calculated exactly since $j_{\ell p}(k)$ admit the exact analytic formula Eqn. (6.12).

6.4.1. Convolutions

The convolution of two functions f and h in a (Hilbert) space of interest is often defined by the inner product of f with a transformed version of h. In standard Fourier analysis this transformation is the natural translation. Likewise, for two signals on the sphere $f, h \in L^2(S^2)$, the convolution is constructed from the rotation operator \mathcal{R}_{ω} :

$$(f \star h)(\omega) \equiv \langle f | \mathcal{R}_{\omega} h \rangle = \int_{S^2} \mathrm{d}\omega' f(\omega') \left(\mathcal{R}_{\omega} h \right)^* (\omega').$$
(6.21)

where, here and henceforth, we restrict ourselves to axisymmetric kernels h, so that the rotation is only parameterised by an angle $\omega = (\theta, \phi)$ (Wiaux et al., 2007, 2008). The spherical harmonic decomposition of $f \star h$ is given by the product of the individual transforms:

$$(f \star h)_{\ell m} = \langle f \star h | Y_{\ell m} \rangle = \sqrt{\frac{4\pi}{2\ell + 1}} f_{\ell m} h_{\ell 0}^*,$$
 (6.22)

with $f_{\ell m} = \langle f | Y_{\ell m} \rangle$ and $h_{\ell 0} \delta_{m 0} = \langle h | Y_{\ell m} \rangle$.

Similarly, we introduce a translation operator \mathcal{T}_r to construct the convolution of two functions on the radial half-line $f, h \in L^2(\mathbb{R}^+)$:

$$(f \star g)(r) \equiv \langle f | \mathcal{T}_r h \rangle = \int_{\mathbb{R}^+} \mathrm{d}r' r'^2 f(r') \left(\mathcal{T}_r h\right)^* (r').$$
(6.23)

The convolution in Laguerre space (GÂólich and Markett, 1982; Kanjin, 1986; Markett, 1986) is defined such that the action T_r on the basis functions is

$$(\mathcal{T}_r K_p)(r') \equiv K_p^*(r) K_p(r'),$$
 (6.24)

in which case $f \star h$ simplifies to a product in spherical Laguerre space, yielding

$$(f \star h)_p = \langle f \star h | K_p \rangle = f_p h_p^*, \tag{6.25}$$

where $f_p = \langle f | K_p \rangle$ and $h_p = \langle h | K_p \rangle$. Consequently any function f which is translated by a distance r on the radial half-line has each coefficient f_p transformed into $f_p K_p(r)$. This operation corresponds to a translation with a damping factor, which is illustrated on a wavelet kernel in

Figure 6.3 (the wavelet kernel itself is defined in Section 6.4.3).⁵

Finally, we define the convolution of two functions on the ball $f, h \in L^2(B^3)$, where h is again assumed to be axisymmetric in the angular direction, by combining the convolution operators defined on the sphere and radial half-line, yielding

$$(f \star h)(\mathbf{r}) \equiv \langle f | \mathcal{T}_r \mathcal{R}_\omega h \rangle$$
 (6.26)

$$= \int_{B^3} \mathrm{d}^3 \mathbf{r}' f(\mathbf{r}') \left(\mathcal{T}_r \mathcal{R}_\omega h \right)^* (\mathbf{r}'). \tag{6.27}$$

The convolution is given in harmonic space by the product

$$(f \star h)_{\ell m p} = \langle f \star h | Z_{\ell m p} \rangle = \sqrt{\frac{4\pi}{2\ell + 1}} f_{\ell m p} h^*_{\ell 0 p}, \tag{6.28}$$

with $f_{\ell m p} = \langle f | Z_{\ell m p} \rangle$ and $h_{\ell 0 p} \delta_{m 0} = \langle h | Z_{\ell m p} \rangle$.

6.4.2. Exact flaglet transform

With an exact harmonic transform and a convolution operator defined on the ball in hand, we are now in a position to construct the exact flaglet transform on the ball. For a function of interest $f \in L^2(B^3)$, we define its jj'-th wavelet coefficient $W^{\Psi^{jj'}} \in L^2(B^3)$ as the convolution of f with the flaglet (*i.e.* wavelet kernel) $\Psi^{jj'} \in L^2(B^3)$:

$$W^{\Psi^{jj'}}(\mathbf{r}) \equiv (f \star \Psi^{jj'})(\mathbf{r}) = \langle f | \mathcal{T}_r \mathcal{R}_\omega \Psi^{jj'} \rangle.$$
(6.29)

The scales j and j' respectively relate to angular and radial spaces. Since we restrict ourselves to axisymmetric kernels, the wavelet coefficients are given in Fourier-Laguerre space by the product

$$W_{\ell m p}^{\Psi^{jj'}} = \sqrt{\frac{4\pi}{2\ell+1}} f_{\ell m p} \Psi_{\ell 0 p}^{jj'*}, \tag{6.30}$$

where $W_{\ell m p}^{\Psi^{jj'}} = \langle W^{\Psi^{jj'}} | Z_{\ell m p} \rangle$, $f_{\ell m p} = \langle f | Z_{\ell m p} \rangle$ and $\Psi_{\ell 0 p}^{jj'} \delta_{m0} = \langle \Psi^{jj'} | Z_{\ell m p} \rangle$. The wavelet coefficients contain the detail information of the signal only; a scaling function and corresponding scaling coefficients must be introduced to represent the low-frequency, approximate information of the signal. The scaling coefficients $W^{\Phi} \in L^2(B^3)$ are defined by the convolution of f with the scaling

⁵Note that this translation operator may also be viewed in real space as a convolution with a delta function, similarly to the Euclidian convolution.





Figure 6.3: Slices of an axisymmetric flaglet wavelet kernel constructed on the ball of radius R = 1, translated along the radial half-line. The chosen kernel has j = j' = 5 and is constructed at resolution P = L = 64. For clarity we zoomed on the range $r \in [0, 0.5]$ (the slice hence relates to a ball of radius r = 0.5). The three-dimensional wavelet can be visualised by rotating this slice around the vertical axis passing through the origin. The translation on the radial half-line not only translates the main feature (the wavelet peak) but also accounts for a damping factor. Flaglets are well localised in both real and Fourier-Laguerre spaces and their angular aperture is invariant under radial translation.

function $\Phi \in L^2(B^3)$:

$$W^{\Phi}(\mathbf{r}) \equiv (f \star \Phi)(\mathbf{r}) = \langle f | \mathcal{T}_r \mathcal{R}_{\omega} \Phi \rangle, \qquad (6.31)$$

or in Fourier-Laguerre space,

$$W^{\Phi}_{\ell m p} = \sqrt{\frac{4\pi}{2\ell + 1}} f_{\ell m p} \Phi^*_{\ell 0 p}, \tag{6.32}$$

where $W^{\Phi}_{\ell m p} = \langle W^{\Phi} | Z_{\ell m p} \rangle$ and $\Phi_{\ell 0 p} \delta_{m 0} = \langle \Phi | Z_{\ell m p} \rangle$.

Provided the flaglets and scaling function satisfy an admissibility property, a function f may be reconstructed exactly from its wavelet and scaling coefficients by

$$f(\mathbf{r}) = \int_{B^3} \mathrm{d}^3 \mathbf{r}' W^{\Phi}(\mathbf{r}') (\mathcal{T}_r \mathcal{R}_\omega \Phi)(\mathbf{r}')$$

$$+ \sum_{j=J_0}^J \sum_{j'=J'_0}^{J'} \int_{B^3} \mathrm{d}^3 \mathbf{r}' W^{\Psi^{jj'}}(\mathbf{r}') (\mathcal{T}_r \mathcal{R}_\omega \Psi^{jj'})(\mathbf{r}'),$$
(6.33)

or equivalently in harmonic space by

$$f_{\ell m p} = \sqrt{\frac{4\pi}{2\ell+1}} W^{\Phi}_{\ell m p} \Phi_{\ell 0 p}$$

$$+ \sqrt{\frac{4\pi}{2\ell+1}} \sum_{j=J_0}^{J} \sum_{j'=J'_0}^{J'} W^{\Psi^{jj'}}_{\ell m p} \Psi^{jj'}_{\ell 0 p}.$$
(6.34)

The parameters J_0 , J'_0 , J and J' defining the minimum and maximum scales must be defined consistently to extract and reconstruct all the information contained in f. They depend on the construction of the flaglets and scaling function and are defined explicitly in the next section.

Finally, the admissibility condition under which a band-limited function f can be decomposed and reconstructed exactly is given by the following resolution of the identity:

$$\frac{4\pi}{2\ell+1} \left(|\Phi_{\ell 0p}|^2 + \sum_{j=J_0}^J \sum_{j'=J'_0}^{J'} |\Psi_{\ell 0p}^{jj'}|^2 \right) = 1, \quad \forall \ell, p.$$
(6.35)

We may now construct flaglets and scaling functions that satisfy this admissibility property and thus lead to an exact wavelet transform on the ball.

6.4.3. Flaglets and scaling functions

We extend the notion of harmonic tiling (Marinucci et al., 2008; Pietrobon et al., 2010; Wiaux

et al., 2008) to the Fourier-Laguerre space and construct axisymmetric wavelets (flaglets) well localised in both real and Fourier-Laguerre spaces. We first define the flaglet and scaling function generating functions, before defining the flaglets and scaling function themselves.

We start by considering the C^∞ Schwartz function with compact support

$$s(t) \equiv \begin{cases} e^{-\frac{1}{1-t^2}}, & t \in [-1,1] \\ 0, & t \notin [-1,1] \end{cases},$$
(6.36)

for $t \in \mathbb{R}$. We introduce the positive real parameter $\lambda \in \mathbb{R}^+_*$ to map s(t) to

$$s_{\lambda}(t) \equiv s\left(\frac{2\lambda}{\lambda - 1}(t - 1/\lambda) - 1\right),\tag{6.37}$$

which has compact support in $[\frac{1}{\lambda}, 1]$. We then define the smoothly decreasing function k_{λ} by

$$k_{\lambda}(t) \equiv \frac{\int_{t}^{1} \frac{\mathrm{d}t'}{t'} s_{\lambda}^{2}(t')}{\int_{1/\lambda}^{1} \frac{\mathrm{d}t'}{t'} s_{\lambda}^{2}(t')},\tag{6.38}$$

which is unity for $t < 1/\lambda$, zero for t > 1, and is smoothly decreasing from unity to zero for $t \in [1/\lambda, 1]$. Axisymmetric flaglets are constructed in a two-dimensional space corresponding to the harmonic indices ℓ and p. We associate λ with ℓ -space and we introduce a second parameter ν associated with p-space, with the corresponding functions s_{ν} and k_{ν} . We define the flaglet generating function by

$$\kappa_{\lambda}(t) \equiv \sqrt{k_{\lambda}(t/\lambda) - k_{\lambda}(t)} \tag{6.39}$$

and the scaling function generating function by

$$\eta_{\lambda}(t) \equiv \sqrt{k_{\lambda}(t)},\tag{6.40}$$

with similar expressions for κ_{ν} and η_{ν} , complemented with a hybrid scaling function generating function

$$\eta_{\lambda\nu}(t,t') \equiv \begin{bmatrix} k_{\lambda}(t/\lambda)k_{\nu}(t') \\ + k_{\lambda}(t)k_{\nu}(t'/\nu) \\ - k_{\lambda}(t)k_{\nu}(t') \end{bmatrix}^{1/2}.$$
(6.41)

The flaglets and scaling function are constructed from their generating functions to satisfy the admissibility condition given by Eqn. (6.35). A natural approach is to define $\Psi_{\ell m p}^{jj'}$ from the generating functions κ_{λ} and κ_{ν} to have support on $[\lambda^{j-1}, \lambda^{j+1}] \times [\nu^{j'-1}, \nu^{j'+1}]$, yielding

$$\Psi_{\ell m p}^{jj'} \equiv \sqrt{\frac{2\ell+1}{4\pi}} \kappa_{\lambda} \left(\frac{\ell}{\lambda^{j}}\right) \kappa_{\nu} \left(\frac{p}{\nu^{j'}}\right) \delta_{m0}.$$
(6.42)

With these kernels, Eqn. (6.35) is satisfied for $\ell > \lambda^{J_0}$ and $p > \nu^{J'_0}$, where J_0 and J'_0 are the lowest wavelet scales used in the decomposition. The scaling function Φ is constructed to extract the modes that cannot be probed by the flaglets:⁶

$$\Phi_{\ell m p} \equiv \begin{cases} \sqrt{\frac{2\ell+1}{4\pi}} \eta_{\nu} \left(\frac{p}{\nu^{J_0'}}\right) \delta_{m0}, & \text{if } \ell > \lambda^{J_0}, \ p \le \nu^{J_0'} \\ \sqrt{\frac{2\ell+1}{4\pi}} \eta_{\lambda} \left(\frac{\ell}{\lambda^{J_0}}\right) \delta_{m0}, & \text{if } \ell \le \lambda^{J_0}, \ p > \nu^{J_0'} \\ \sqrt{\frac{2\ell+1}{4\pi}} \eta_{\lambda\nu} \left(\frac{\ell}{\lambda^{J_0}}, \frac{p}{\nu^{J_0'}}\right) \delta_{m0}, \text{ if } \ell < \lambda^{J_0}, \ p < \nu^{J_0'} \\ 0, & \text{elsewhere.} \end{cases}$$

To satisfy exact reconstruction, J and J' are defined from the band-limits by $J = \lceil \log_{\lambda}(L-1) \rceil$ and $J' = \lceil \log_{\nu}(P-1) \rceil$. The choice of J_0 and J'_0 is arbitrary, provided that $0 \le J_0 < J$ and $0 \le J'_0 < J'$. This framework generalises the notion of the harmonic tiling used to construct exact wavelets on the sphere (Marinucci et al., 2008; Wiaux et al., 2008); in fact, the flaglets defined here reduce in angular part to the wavelets defined in Wiaux et al. (2008) for the axisymmetric case. The flaglets and scaling function tiling of the Fourier-Laguerre space of the ball is illustrated in Figure 6.4. Flaglets and the scaling function may be reconstructed in the spatial domain from their harmonic coefficients. In Figure 6.5 flaglets are plotted in the spatial domain for a range of different scales; translated flaglets are plotted in Figure 6.3. The flaglets are well localised in both real and Fourier-Laguerre spaces and their angular aperture is invariant under radial translation.

6.5. Multiresolution Algorithm

In this section we discuss our implementation of the Fourier-Laguerre and flaglet transforms. We notably introduce a multiresolution algorithm for the flaglet transform to capture each wavelet scale in the minimal number of samples on the ball, thereby reducing the computational cost of

⁶Note that despite its piecewise definition $\Phi_{\ell m p}$ is continuous along and across the boundaries $p = \nu^{J'_0}$ and $\ell = \lambda^{J_0}$.



Figure 6.4: Tiling of Fourier-Laguerre space at resolution L = N = 64 for flaglet parameters $\lambda = \nu = 2$, giving J = J' = 7. Flaglets divide Fourier-Laguerre space into regions corresponding to specific scales in angular and radial space. The scaling part, here chosen as $J_0 = J'_0 = 4$, is introduced to cover the low frequency region and insures that large scales are also represented by the transform.



Figure 6.5: Slices of four successive axisymmetric flaglet wavelet kernels, probing different scales in angular and radial space. The flaglet parameters are $\lambda = \nu = 2$ and the kernels are constructed at resolution L = N = 92 on a ball of radius R = 1. For visualisation purposes we show the flaglets corresponding to $j \in \{4, 5\}$ and $j' \in \{5, 6\}$, translated to r = 0.3 and zoomed on the range $r \in [0, 0.4]$. Kernels of angular order j = 4 (first row) probe large angular scales compared to those of order j = 5 (second row). Similarly, kernels of radial order j' = 5 (first column) probe large radial scales compared to those of order j' = 5 (second column).

the transform. We finally provide accuracy and complexity tests for our implementation of both transforms, which we make publicly available.

6.5.1. Algorithm

In our framework, each flaglet $\Psi^{jj'}$ has compact support in Fourier-Laguerre space on $\ell \times p \in [\lambda^{j-1}, \lambda^{j+1}] \times [\nu^{j'-1}, \nu^{j'+1}]$, as shown in Figure 6.4. Thus, $\Psi^{jj'}$ has band-limits in ℓ and p of λ^{j+1} and $\nu^{j'+1}$ respectively. For a band-limited function $f \in L^2(B^3)$, recall that the jj'-th wavelet contribution is given by the simple product of Eqn. (6.30) in harmonic space. Consequently, the band-limits of $W^{\Psi^{jj'}}$ are given by the minimum of the band-limits of f and $\Psi^{jj'}$. Thus, for j < J or j' < J' the wavelet scale $W^{\Psi^{jj'}}$ can be represented in fewer samples than f, without any loss of information. We exploit this property by designing a multiresolution approach where each wavelet scale is represented in real space with the smallest number of samples necessary. Note that the scaling function must be used at full resolution algorithm, although f is decomposed at full resolution, the wavelets coefficients are reconstructed in real space with the minimum number of samples supporting their band-limits. This leads to a significant reduction in computation time, which is then dominated by the small number of full resolution Fourier-Laguerre transforms.

6.5.2. Fast implementation

Our implementation of the algorithms of this article is made available in the following three packages, which are written in C and include MATLAB interfaces for most high-level features, and are described in turn:

- FLAG: spherical Laguerre transform and Fourier-Laguerre transforms on the ball (exact spherical Bessel and Fourier-Bessel decompositions are optional features that additionally require the GNU Math Library⁷).
- S2LET: axisymmetric wavelet transform on the sphere through harmonic tiling.
- FLAGLET: axisymmetric flaglet transform on the ball, combining FLAG and S2LET to construct flaglets in Fourier-Laguerre space through harmonic tiling.

⁷http://www.gnu.org/software/gsl/

We make these three packages publicly available.⁸ All packages require SSHT⁹, which implements fast and exact algorithms to perform the forward and inverse spherical harmonic transforms corresponding to the MW sampling theorem (McEwen and Wiaux, 2011). SSHT requires the FFTW¹⁰ package.

Since the naive spherical harmonic transform scales as $\mathcal{O}(L^4)$ and the spherical Laguerre transform scales as $\mathcal{O}(P^2)$, the naive complexity of the Fourier-Laguerre transform is $\mathcal{O}(P^2L^4)$. However, rather than computing triple integrals/sums over the ball directly, it is straightforward to show that the Fourier-Laguerre transform can be performed separately on the sphere and on the radial half-line, like the Fourier-Bessel transform (Leistedt et al., 2012). Since the angular and radial samplings are separable, the related transforms can be computed independently through a separation of variables, so that the complexity reduces to $\mathcal{O}(Q^5)$ for $Q \sim P \sim L$. The separation of variables also means we are able to exploit high-performance recurrences and algorithms that exist for both the spherical Laguerre and spherical harmonic transforms. In particular, the radial basis functions $K_p(r)$ are calculated using a normalised recurrence formula derived from the recurrence on the Laguerre polynomials. Moreover, a critical point for the accuracy of the Fourier-Laguerre transform is the computation of the Gauss-Laguerre quadrature, for which we use the previous normalised recurrence complemented with an appropriate root-finder algorithm. The fast spherical harmonic transforms implemented in the SSHT package use the Trapani & Navaza method (Trapani and Navaza, 2006) to efficiently compute Wigner functions (which are closely related to the spherical harmonics) through recursion.¹¹ These fast spherical harmonic transform algorithms (McEwen and Wiaux, 2011) scale as $\mathcal{O}(L^3)$. The final complexity achieved by the Fourier-Laguerre transform is thus $\mathcal{O}(Q^4)$.

The flaglet transform (forward and inverse) is calculated in a straightforward manner in Fourier-Laguerre space, thus its computation is dominated by the Fourier-Laguerre transform of the signal, approximation coefficients, and wavelets coefficients at all scales, requiring $[(J + 1 - J_0)(J' + 1 - J'_0) + 2]$ Fourier-Laguerre transforms. If all wavelet contributions are reconstructed at full resolution in real space, the overall wavelet transform scales as $\mathcal{O}([(J + 1 - J_0)(J' + 1 - J'_0) + 2]Q^4)$. Note that J and J' depend on the band-limits L and P and the parameters λ and ν , respectively. However, in the previous section we established a multiresolution algorithm that takes advantage of the band-limits of the individual flaglets. With this algorithm, only the scaling function and the finest wavelet

⁸http://www.flaglets.org/

⁹http://www.jasonmcewen.org/

¹⁰http://www.fftw.org/

¹¹Alternatively, Risbo's method could also be used to compute Wigner functions (Risbo, 1996).

scales (i.e. $j \in \{J - 1, J\}$ and $j' \in \{J' - 1, J'\}$) are computed at maximal resolution corresponding to band-limits *L* and *P*. The complexity of the overall multiresolution flaglet transform is then dominated by these operations and scales as $\mathcal{O}(Q^4)$.

6.5.3. Numerical validation

In this section we evaluate FLAG and FLAGLET in terms of accuracy and complexity. We show that they achieve floating-point precision and scale as detailed in the previous section. In both cases we consider band-limits $L = P = 2^i$ with $i \in \{2, ..., 9\}$ and generate sets of harmonic coefficients $f_{\ell mp}$ following independent Gaussian distributions $\mathcal{N}(0, 1)$. We then perform either the Fourier-Laguerre or the flaglet decomposition, before reconstructing the harmonic coefficients, therefore denoted by $f_{\ell mp}^{\text{rec}}$. We evaluate the accuracy of the transforms using the error metric $\epsilon = \max |f_{\ell mp} - f_{\ell mp}^{\text{rec}}|$, which is theoretically zero for both transforms since all signals are band-limited by construction. The complexity is quantified by observing how the computation time $t_c = [t_{\text{synthesis}} + t_{\text{analysis}}]/2$ scales with the band-limits, where the synthesis and analysis computation times, $t_{\text{synthesis}}$ and t_{analysis} respectively, are defined explicitly for the two transforms in the paragraphs that follow. The stability of both ϵ and t_c is checked by averaging over hundreds of realisations of $f_{\ell mp}$ in the cases $i \in \{2, ..., 7\}$ and a small number of realisations for $i \in \{8, 9\}$. Recall that for given band-limits L and P the number of samples on the ball required by the exact quadrature is N = P[(2L - 1)(L - 1) + 1]. All tests were run on a 2.5GHz Core i5 processor with 8GB of RAM.

The results of these tests for the Fourier-Laguerre transform are presented on Figure 6.6. The indicators ϵ and t_c are plotted against the number of samples N. Each test starts from coefficients $f_{\ell mp}$ randomly generated. The synthesis refers to constructing the band-limited signal f from the decomposition $f_{\ell mp}$. The analysis then corresponds to decomposing f into Fourier-Laguerre coefficients $f_{\ell mp}^{\text{rec}}$. As shown in Figure 6.6, FLAG achieves very good numerical accuracy, with numerical errors comparable to floating-point precision, and computation time scales as $\mathcal{O}(Q^4)$, in agreement with theory.

The results of similar tests for the flaglet transform (entirely performed in real space) are presented on Figure 6.7. As previously, the indicators ϵ and t_c are plotted against the number of samples N. Since we evaluate the flaglet transform in real space, a preliminary step is required to construct a band-limited signal f from the randomly generated $f_{\ell mp}$. This step is not included in the computation time since its only purpose is to generate a valid band-limited test signal in real space. The analysis then denotes the decomposition of f into wavelet coefficients $W^{\Psi^{jj'}}$ and scaling coefficients W^{Φ} on the ball. The synthesis refers to recovering the signal f^{rec} from these coefficients. The final step, which is not included in the computation time, is to decompose f^{rec} into Fourier-Laguerre coefficients $f^{\text{rec}}_{\ell mp}$ in order to compare them with $f_{\ell mp}$. As shown in Figure 6.7, FLAGLET achieves very good numerical accuracy, with numerical errors comparable to floating-point precision. Moreover, the full resolution and multiresolution algorithms are indistinguishable in terms of accuracy. However, the latter is ten times faster than the former since only the scaling function and a small number of wavelet coefficients are computed at full resolution. As shown in Figure 6.7, computation time scales as $\mathcal{O}(Q^4)$ for both algorithms, in agreement with theory.

6.6. Denoising Illustration

In this section we illustrate the use of the flaglet transform in the context of a simple denoising problem. We consider two datasets naturally defined on the ball and contaminate them with bandlimited noise. We compute the flaglet transform of the noisy signal and perform simple denoising by hard-threshold the wavelet coefficients. We reconstruct the signal from the thresholded wavelet coefficients and examine the improvement in signal fidelity.

6.6.1. Wavelet denoising

Consider the noisy signal $y = s + n \in L^2(B^3)$, where the signal of interest $s \in L^2(B^3)$ is contaminated with noise $n \in L^2(B^3)$. A simple way to evaluate the fidelity of the observed signal yis to examine the signal-to-noise ratio, which we define on the ball by

$$SNR(y) \equiv 10 \log_{10} \frac{\|s\|_2^2}{\|y - s\|_2^2}.$$
 (6.43)

The signal energy is given by

$$\|y\|_{2}^{2} \equiv \langle y|y\rangle = \int_{B^{3}} \mathrm{d}^{3}\mathbf{r}|y(\mathbf{r})|^{2} = \sum_{\ell m p} |y_{\ell m p}|^{2}, \qquad (6.44)$$

where the final equality follows from a Parseval relation on the ball (which follows directly from



(b) Computation time of the Fourier-Laguerre transform

Figure 6.6: Numerical accuracy and computation time of the Fourier-Laguerre transform computed with FLAG, where N corresponds to the number of samples on the ball required to capture all the information contained in the band-limited test signal. We consider $L = P = 2^i$ with $i \in \{2, ..., 9\}$. These results are averaged over many realisations of random band-limited signals and were found to be very stable. Very good numerical accuracy is achieved, with numerical errors comparable to floating-point precision, found empirically to scale as $\mathcal{O}(Q^2)$ as shown by the red line in panel (a), where $Q \sim P \sim L$. Computation time scales as $\mathcal{O}(Q^4)$ as shown by the red line in panel (b), in agreement with theory.



(b) Computation time of the flaglet transform

Figure 6.7: Numerical accuracy and computation time of the flaglet transform computed with FLAGLET, where *N* corresponds to the number of samples on the ball required to capture all the information contained in the band-limited test signal. We consider $L = P = 2^i$ with $i \in \{2, ..., 9\}$, with parameters $\lambda = \nu = 2$, $J_0 = J'_0 = 0$. These results are averaged over many realisations of random band-limited signals and were found to be very stable. The flaglet transform is either performed at full-resolution (dashed lines) or with the multiresolution algorithm (solid lines). Very good numerical accuracy is achieved by both the full resolution and multiresolution algorithms (which achieve indistinguishable accuracy), with numerical errors comparable to floating-point precision, found empirically to scale as O(PL) as shown by the red line in panel (a). The multiresolution algorithm is ten times faster than the full-resolution approach. Computation time scales as $O(PL^3)$ for both algorithms as shown by the red line in panel (b), in agreement with theory.

the orthogonality of the Fourier-Laguerre basis functions). In practice, we compute signal energies through the final Fourier-Laguerre space expression to avoid the necessity of an explicit quadrature rule.

We seek a denoised version of y, denoted by $d \in L^2(B^3)$, such that SNR(d) is as large as possible in order to extract the informative signal s. We take the flaglet transform of the noisy signal since we intend to denoise the signal in wavelet space, where we expect the energy of the informative signal to be concentrated in a small number of wavelet coefficients while the noise energy will be spread over many wavelet coefficients. Since the flaglet transform is linear, the wavelet coefficients of the jj'-th scale of the noisy signal is simply the sum of the individual contributions:

$$Y^{jj'}(\mathbf{r}) = S^{jj'}(\mathbf{r}) + N^{jj'}(\mathbf{r}),$$
(6.45)

where capital letters denote the wavelet coefficients, *i.e.* $Y^{jj'} \equiv y \star \Psi^{jj'}$, $S^{jj'} \equiv s \star \Psi^{jj'}$ and $N^{jj'} \equiv n \star \Psi^{jj'}$.

In the illustrations performed here, we assume the noise model

$$\mathbb{E}\left(|n_{\ell m p}|^2\right) = \sigma^2 \left(\frac{p}{P}\right)^2 \delta_{\ell \ell'} \delta_{m m'} \delta_{p p'}, \qquad (6.46)$$

which corresponds to a white noise for the angular space with a dependence on the radial mode p, where $\mathbb{E}(\cdot)$ denotes ensemble averages. We do not opt for a white noise in radial space (*i.e.* $\mathbb{E}(|n_{\ell mp}|^2) = \sigma^2 \delta_{\ell\ell'} \delta_{mm'} \delta_{pp'}$) because the latter has its energy concentrated in the centre of the ball due to the shape of the spherical Laguerre basis functions. The *p*-dependence gives a greater weight to small-scale radial features and hence yields a more homogeneous noise on the ball, which is more useful for visualisation purposes. For this noise model one can show that the expected covariance of the wavelet coefficients of the *jj'*-th scale reads

$$\mathbb{E}\left(|N^{jj'}(r,\omega)|^2\right) = \sigma^2 \sum_{\ell p} \left(\frac{p}{P}\right)^2 |\Psi_{\ell 0p}^{jj'}|^2 |K_p(r)|^2 \qquad (6.47)$$
$$\equiv \left(\sigma^{jj'}(r)\right)^2.$$

Denoising is performed by hard-thresholding the wavelet coefficients $Y^{jj'}$, where the threshold is taken as $T(\mathbf{r}) = 3\sigma^{jj'}(r)$. The wavelet coefficients of the denoised signal $D^{jj'} \equiv d \star \Psi^{jj'}$ are thus given by

$$D^{jj'}(\mathbf{r}) = \begin{cases} 0, & \text{if } Y^{jj'}(\mathbf{r}) < T(\mathbf{r}) \\ Y^{jj'}(\mathbf{r}), & \text{otherwise.} \end{cases}$$
(6.48)

The denoised signal d is then reconstructed from its wavelet coefficients and scaling coefficients (the latter are not thresholded and thus not altered). To assess the effectiveness of this simple flaglet denoising strategy when the informative signal s is known, we compute the SNR of the denoised signal and compare it to the SNR of the original noisy signal. In what follows we apply this simple denoising technique to two datasets naturally defined on the ball.

6.6.2. Examples

The first dataset we consider is the full-sky Horizon simulation (Teyssier et al., 2009): an N-body simulation covering a 1Gpc periodic box of 70 billion dark matter particles generated from the concordance model cosmology derived from 3-year Wilkinson Microwave Anisotropy Probe (WMAP) observations (Spergel et al., 2007). The purpose of such a simulation is to reproduce the action of gravity (and to a minor extent galaxy formation) on a large system of particles, with the initial conditions drawn from a cosmological model of interest. The outcome is commonly used to confront astrophysical models with observations. For simplicity we only consider a ball of 1MPc radius centered at the origin so that the structures are of reasonable size. Figure 6.8 shows the initial data, band-limited at L = P = 128, as well as their wavelet coefficients with $\lambda = \nu = 2$, J = J' = 7and scaling coefficients for $J_0 = J'_0 = 6$ since the lower scale indices do not contain a great deal of information. We see that the filamentary distribution of matter is naturally suited to a flaglet analysis on the ball since the informative signal is likely to be contained in a reduced number of wavelet coefficients. The original data are corrupted by the addition of random noise defined by Eqn. (6.46) for an SNR of 5dB. The wavelet denoising procedure described previously is then applied. The denoised signal is recovered with an SNR of 11dB, highlighting the effectiveness of this very simple flaglet denoising strategy on the ball. The results of this denoising illustration are presented in Figure 6.9.

The second dataset we consider is Ritsema's seismological Earth model of shear wavespeed perturbations in the mantle, known as S40RTS (Ritsema et al., 2011; Simons et al., 2011; Simons et al., 2011).¹² The model supplies spherical harmonic coefficients in the angular dimension and

¹²http://www.earth.lsa.umich.edu/~jritsema/



Figure 6.8: Flaglet decomposition of the N-body simulation dataset considered for the first denoising example. The initial dataset was pixelised and band-limited at L = P = 128. The flaglet parameters are $\lambda = \nu = 2$ (giving J = J' = 7) and the scaling coefficients correspond to $J_0 = J'_0 = 6$ since the lower scale indices do not contain a great deal of information. The four wavelet coefficients together with the scaling coefficients decompose the initial dataset exactly, *i.e.* the original signal can be recovered perfectly from these wavelet and scaling coefficients. All signals were oversampled on L = P = 256 for visualisation purposes.



(a) Band-limited data

(b) Noise



Figure 6.9: Denoising of an N-body simulation. The data are contaminated with a band-limited noise and decomposed into wavelet coefficients. Denoising is performed by a simple hard-thresholding of the wavelet coefficients, following a noise model. The denoised signal is then reconstructed from the thresholded wavelet coefficients. In this example, for an initial SNR of 5dB, the flaglet denoised signal is recovered with SNR of SNR = 11dB (with resolution L = P = 128, oversampled on L = P = 256 and using flaglet parameters $\lambda = \nu = 2$, $J_0 = J'_0 = 0$, giving J = J' = 7).

radial spline coefficients in the depth dimension to define a signal on the ball, which we band-limit. Contrarily to the first example, Ritsema's model does not contain a lot of structure at the smallest scales but essentially contains large-scale features. As previously, the original data are corrupted by the addition of random noise defined by Eqn. (6.46) for an SNR of 5dB. The flaglet denoising procedure described previously is then applied. The denoised signal is recovered with an SNR of 17dB, again highlighting the effectiveness of this very simple flaglet denoising strategy on the ball. As expected, the improvement in SNR is better than for the previous dataset since the informative signal is mainly captured by a few large wavelet scales. The results of this denoising illustration are presented in Figure 6.10.

6.7. Conclusions

One reason an exact wavelet transform of a band-limited signal on the ball has not yet been derived is due to the absence of an exact harmonic transform on the ball. We have taken advantage of the orthogonality of the Laguerre polynomials on \mathbb{R}^+ to define the spherical Laguerre transform, a novel radial transform that admits an exact quadrature rule. Combined with the spherical harmonics, we used this to derive a sampling theorem and an exact harmonic transform on the ball, which we call the Fourier-Laguerre transform. A function that is band-limited in Fourier-Laguerre space can be decomposed and reconstructed at floating-point precision, and its Fourier-Bessel transform can be calculated exactly. For radial and angular band-limits *P* and *L*, respectively, the sampling theorem guarantees that all the information of the band-limited signal is captured in a finite set of N = P[(2L - 1)(L - 1) + 1] samples on the ball.

We have developed an exact wavelet transform on the ball, the so-called flaglet transform, through a tiling of the Fourier-Laguerre space. The resulting flaglets form a tight frame and are well localised in both real and Fourier-Laguerre spaces. Their angular aperture is invariant under radial translation. We furthermore established a multiresolution algorithm to compute the flaglet transform, capturing all the information contained in each wavelet scale in the minimal number of samples on the ball, thereby reducing the computation cost of the flaglet transform considerably.

Flaglets are a promising new tool for analysing signals on the ball, particularly for extracting spatially localised features at different scales of interest. Exactness of both the Fourier-Laguerre and the flaglet transforms guarantees that any band-limited signal can be analysed and decomposed



Figure 6.10: Denoising of a seismological Earth model. The data are contaminated with a bandlimited noise and decomposed into wavelet coefficients. Denoising is performed by a simple hard-thresholding of the wavelet coefficients, following a noise model. The denoised signal is then reconstructed from the thresholded wavelet coefficients. In this example, for an initial SNR of 5dB, the flaglet denoised signal is recovered with SNR of 17dB (with resolution L = P = 128 and using flaglet parameters $\lambda = \nu = 3$, $J_0 = J'_0 = 0$, giving J = J' = 7).

into wavelet coefficients and then reconstructed without any loss of information. To illustrate these capabilities, we considered the denoising of two different datasets which were contaminated with synthetic noise. A very simple flaglet denoising strategy was performed by hard-thresholding the wavelet coefficients of the noisy signal, before reconstructing the denoised signal from the threshol-ded wavelet coefficients. In these illustrations a considerable improvement in SNR was realised by this simple flaglet denoising strategy, demonstrating the effectiveness of flaglets for the analysis of data defined on the ball. Our implementation of all of the transforms and examples detailed in this article is made publicly available. In future work we intend to revoke the axisymmetric constraint by developing directional flaglets.

No new cosmological concordance with massive sterile neutrinos

"I have done a terrible thing, I have postulated a particle that cannot be detected". W. Pauli

7.1. Abstract

It has been claimed recently that massive sterile neutrinos could bring about a new concordance between observations of the cosmic microwave background (CMB), the large-scale structure (LSS) of the Universe, and local measurements of the Hubble constant, H_0 . We demonstrate that this apparent concordance results from combining datasets which are in significant tension, even within this extended model, possibly indicating remaining systematic biases in the measurements. We further show that this tension remains when the cosmological model is further extended to include significant tensor modes, as suggested by the recent BICEP2 results. Using the Bayesian evidence, we show that the minimal Λ CDM model is strongly favoured over its neutrino extensions by various combinations of datasets. Robust data combinations yield stringent limits of $\sum m_{\nu} \leq 0.3$ eV and $m_{\nu,\text{sterile}}^{\text{eff}} \leq 0.3$ eV at 95% CL for the sum of active and sterile neutrinos, respectively.

7.2. Introduction

The temperature fluctuations of the cosmic microwave background (CMB), as measured by the *Planck* satellite (Planck Collaboration, 2013k), have yielded sub-percent level constraints on the

cosmological parameters of the vanilla Λ CDM model. However, the primary CMB temperature fluctuations only indirectly probe the growth of cosmic structure, and it is therefore essential to complement it with observations large-scale structure (LSS) such as galaxy clusters, weak lensing, and clustering measurements. The first cosmological results from the *Planck* satellite have revealed a ~ 2 σ tension between CMB temperature measurements and the Sunyaev-Zel'dovich (SZ) cluster abundances (Planck Collaboration, 2013j), mainly in terms of σ_8 , the linear-theory mass dispersion on a scale of $8h^{-1}$ Mpc. A similar tension is observed with the X-ray cluster counts (Vikhlinin et al., 2009).

Massive neutrinos can potentially alleviate this tension because they suppress power in the clustering of matter at late times. They are an appealing solution since solar and atmospheric experiments have already provided evidence for their mass, with room for extra sterile species, supported by anomalies in short baseline and reactor neutrino experiments (for reviews of particle physics constraints, see *e.g.*, Abazajian et al. (2012); Beringer et al. (2012); Conrad et al. (2013); Gonzalez-Garcia et al. (2012). Cluster abundances, galaxy surveys and weak lensing are sensitive to the total neutrino mass, either from active neutrinos $\sum m_{\nu}$ (the total mass from active species), or sterile neutrinos $m_{\nu,\text{sterile}}^{\text{eff}}$ (an effective parameter that connects to actual neutrino masses in the context of specific models, see *e.g.*, Lesgourgues and Pastor (2012). In addition, an extra parameter N_{eff} can be introduced to denote the effective number of relativistic species, in which case $N_{\text{eff}} > 3.046$ (the standard number) is referred to as "dark radiation", and is also appealing as it could alleviate the tension between *Planck* and local H_0 measurements (Verde et al., 2013).

A number of recent studies have carried out joint analyses of various data combinations to conclude that these tensions are resolved within a new concordance model which implies non-standard neutrino parameters (Battye and Moss, 2014; Beutler et al., 2014; Giusarma et al., 2014; Hamann and Hasenkamp, 2013; Wyman et al., 2013). Battye and Moss (2014) argued that combining the CMB with lensing or SZ cluster measurements reveals evidence for non-zero neutrino mass in both the active and sterile neutrino scenarios. Hamann and Hasenkamp (2013); Wyman et al. (2013) claimed that sterile neutrinos could reconcile *Planck* with LSS data, in particular with the X-ray cluster abundances (Vikhlinin et al., 2009) and the latest constraints on H_0 (Riess et al., 2011b). By combining the CMB with shear and redshift space distortion (RSD) measurements, Beutler et al. (2014) found hints of non-zero masses for active neutrinos. Finally, Dvorkin et al. (2014); Zhang et al. (2014) further claimed that sterile neutrinos could resolve a potential tension between *Planck* and BICEP (BICEP2 Collaboration et al., 2014) constraints on $r_{0.002}$, the tensor-to-scalar ratio at k = 0.002 Mpc⁻¹.



Figure 7.1: Constraints on the Λ CDM+N_{eff}+ $m_{\nu,\text{sterile}}^{\text{eff}}$ model, showing that non-zero sterile neutrino mass is only favoured as a result of a tension between the CMB and cluster data (PlaSZ, Xray) in the σ_8 - Ω_m plane, and the degeneracy between σ_8 and neutrino mass.

Although these conclusions are not universally accepted (Efstathiou, 2013; Feeney et al., 2013; Hu et al., 2014; Planck Collaboration, 2013k; Verde et al., 2013), tension between the datasets may indeed point to new physics. Alternatively, tension may also indicate remaining systematic biases in the measurements, which can have substantial impact on cosmological parameter measurements at the level of precision achieved by current data. Consequently, new physics in the neutrino sector is only a viable solution if the extra parameters eliminate the tension between datasets seen in the standard concordance cosmology, and is robustly confirmed by a variety of datasets. In this *Letter*, we show that sterile neutrinos do not relieve the tension between *Planck* and X-ray and SZ clusters, or with local measurements of H_0 . Further, we show that the extended neutrino models are not preferred over the minimal model by any data combination, and that robust combinations of current measurements prefer low neutrino masses $\sum m_{\nu}$, $m_{\nu,\text{sterile}} \lesssim 0.3$ eV.

7.3. Data and methods.

We use CosmoMC (Lewis and Bridle, 2002) to constrain the parameters of the Λ CDM model extended with active $(+N_{\text{eff}}, \sum m_{\nu})$ and sterile $(+N_{\text{eff}}, m_{\nu,\text{sterile}}^{\text{eff}})$ neutrinos, using combinations of the following datasets. <u>CMB</u>: the *Planck* CMB temperature likelihood (Planck Collaboration, 2013i), combined with Wilkinson Microwave Anisotropy Probe (WMAP) polarisation (Bennett et al., 2012), and high- ℓ temperature spectra from Atacama Cosmology Telescope (ACT) and South Pole Telescope (SPT) (Das et al., 2013; Keisler et al., 2011; Reichardt et al., 2012). Lensing: the CMB lensing

likelihood from *Planck* (Planck Collaboration, 2013d). <u>BAO</u>: the Baryon Acoustic Oscillations (BAO) measurements from 6dF (Beutler et al., 2011), Sloan Digital Sky Survey (SDSS) DR7 (Padmanabhan et al., 2012), WiggleZ (Blake et al., 2011), and Baryon Oscillation Spectroscopic Survey (BOSS) DR11 (Anderson et al., 2013). <u>Shear</u>: the weak lensing tomographic analysis from Canada-France Hawaii Telescope Lensing Survey (CFHTLenS) (Kilbinger et al., 2013). <u>PlaSZ</u>: the *Planck* SZ cluster abundances (Planck Collaboration, 2013j). <u>RSD</u>: the RSD measurements from BOSS (Beutler et al., 2013, 2014). <u>Xray</u>: X-ray cluster mass function constraints (Vikhlinin et al., 2009). <u>HST</u>: the H_0 measurement using supernovae by the *Hubble Space Telescope* (Riess et al., 2011b). <u>Clustering</u>: the 3D galaxy power spectrum from WiggleZ (Blake et al., 2010; Parkinson et al., 2012), and the power spectrum of the reconstructed halo density field derived from Luminous Red Galaxies (LRG) in SDSS DR7 (Reid et al., 2010), both up to $k = 0.2 h Mpc^{-1}$. Note that we only use either the power spectrum or the BAO measurement from each dataset.

Finally, we use the Evidence ratio (or Bayes factor), which gives the relative odds of two models correctly describing the observations, under the assumption of equal *a priori* model probabilities (see *e.g.*, Cox 1946; Verde et al. 2013 and references therein). We calculate $\ln[E_{\Lambda CDM}/E_{ext.}]$, the logarithm of the Evidence ratio of the Λ CDM model divided by that of the the extended neutrino models; thus, positive numbers favour the minimal model. In practice, since the models are nested, we compute Evidence ratios with the Savage-Dickey Density Ratio, and we use Kernel Density Estimation (KDE) to process Monte-Carlo Markov chains and reliably compute the marginalised posterior distributions at the Λ CDM values ($\sum m_{\nu} = 0.06 \text{ eV}, m_{\nu,\text{sterile}}^{\text{eff}} = 0.0 \text{ eV}, N_{\text{eff}} = 3.046$). The errors are calculated by jackknifing the KDE parameters. For all parameters, we consider the same prior ranges as the official *Planck* analysis (Planck Collaboration, 2013k). However, the Bayes factors only depend on the neutrino parameters since we consider nested models. Specifically, we assume uniform priors in [0, 5], [0, 3] and [3.046, 10] for $\sum m_{\nu}, m_{\nu,\text{sterile}}^{\text{eff}}$ and N_{eff} , respectively, and we impose $m_{\nu,\text{sterile}}^{\text{eff}}/(N_{\text{eff}} - 3.046) < 7$ eV to avoid a degeneracy between very massive neutrinos and cold dark matter.

7.4. No new concordance with sterile neutrinos.

Fig. 7.1 shows constraints on the $\sigma_8-m_{\nu,\text{sterile}}^{\text{eff}}$ plane for several data combinations, including those used by Battye and Moss (2014); Hamann and Hasenkamp (2013); Wyman et al. (2013). Our

Table 7.1: Evidence ratios $\ln[E_{\Lambda CDM}/E_{ext.}]$ between the minimal ΛCDM model and the extended neutrino models, in the active and sterile scenarios, showing that the extended models are not favoured by any data combination. In particular, the upper part refers to the "tension" data combinations of Fig. 7.1, whereas the lower part corresponds to more robust data combinations (details in text), for which marginalised constraints are presented in Tables 7.2 and 7.3.

	Active	Sterile
CMB+BAO+PlaSZ+Xray+HST	$1.52_{-0.33}^{+0.16}$	$-0.16^{+0.39}_{-0.35}$
CMB+Lensing+BAO+Shear+PlaSZ	$3.77^{+0.10}_{-0.09}$	$1.05^{+0.26}_{-0.55}$
CMB+BAO	$4.42_{-0.05}^{+0.04}$	$3.10^{+0.07}_{-0.14}$
CMB+Lensing+BAO	$4.64_{-0.09}^{+0.03}$	$2.99^{+0.06}_{-0.05}$
CMB+Lensing+BAO+Clustering	$4.70_{-0.00}^{+0.02}$	$3.35\substack{+0.09\\-0.13}$
CMB+Lensing+BAO+Clusters	$4.65_{-0.19}^{+0.10}$	$2.61^{+0.21}_{-0.23}$
CMB+Lensing+BAO+Shear	$4.32_{-0.16}^{+0.10}$	$2.10^{+0.21}_{-0.41}$
CMB+Lensing+BAO+RSD	$4.14_{-0.19}^{+0.10}$	$1.81\substack{+0.11\\-0.09}$

minimal dataset is CMB+BAO, since adding BAO to CMB does not shift the contours but constrains the matter density Ω_m and reduces the error-bars (as expected for consistent datasets). However, the addition of the PlaSZ or Xray clusters, which prefer lower σ_8 , shifts the contours significantly (by more than 2σ) outside the region allowed by CMB+BAO. This clearly indicates that the addition of sterile neutrinos to the Λ CDM model does not bring the CMB and cluster measurements into agreement. Note that the active scenario (not shown here) leads to similar results and tension, and does not yield concordance within the extended model either. Thus we may conclude that the tension must be resolved either by considering systematics in one or more of the relevant datasets, or else by new physics other than the introduction of massive (active or sterile) neutrinos. This is confirmed by the Bayes factor, presented in the first section of Table 7.1, showing that the extended models are not preferred over the minimal Λ CDM model even in the presence of a tension.

Cluster cosmology is currently limited by modelling rather than statistical uncertainties (Planck Collaboration, 2013j); thus, error-bars on the X-ray, SZ and optical clusters data used in Fig. 7.1 and in Battye and Moss (2014); Beutler et al. (2014); Dvorkin et al. (2014); Giusarma et al. (2014); Hamann and Hasenkamp (2013); Wyman et al. (2013); Zhang et al. (2014) may need to be significantly increased to account for additional potential systematics. The calibration of the mass-observable relation is critical for deriving robust cosmological constraints from clusters, and is complicated by uncertainties in mass measurements and the selection functions (see *e.g.*, Rozo et al. 2013; Vikhlinin et al. 2009). Constraints on σ_8 from PlaSZ clusters are sensitive to assumptions and uncertainties in the modelling, as investigated in Planck Collaboration (2013j), and there are indications of a systematic mismatch between masses obtained via weak lensing compared with SZ

masses (von der Linden et al., 2014). The error bars on $\sigma_8(\Omega_m)^\beta$ from X-ray clusters used in Battye and Moss (2014) should be enlarged to account for confirmed sources of systematic uncertainties (Vikhlinin et al., 2009). Interestingly, it was shown that the mass calibration by Rozo et al. (2014) from a self-consistent analysis of X-ray, SZ, and optical scaling relations is consistent with a minimal flat Λ CDM model with no massive neutrinos (1.7 σ), and is a better fit to additional data (*e.g.* H_0). Finally, the model dependence of these cluster constraints in the context of non-standard models has not been investigated; therefore it is unclear whether they can be used in a joint analysis in the context of such extended models.

If, after further investigation of such systematic effects, PlaSZ and Xray clusters remain in tension with CMB+BAO, this tension cannot be simply resolved by adding sterile neutrinos.

7.5. Constraints on neutrino masses from robust datasets.

We now investigate the constraints obtained on neutrino masses when combining datasets which are compatible and have been demonstrated to be robust to modelling uncertainties. Recent works using galaxy power spectra have obtained tight constraints on the mass of active neutrinos (e.g., Giusarma et al. 2013; Riemer-Sorensen et al. 2013; Thomas et al. 2010), and also showed that it could help in breaking degeneracies with the freedom in the primordial power spectrum from inflation (de Putter et al., 2014). For Clustering data, we use the power spectra from SDSS DR7 (reconstructed halo power spectrum) and WiggleZ (galaxy power spectrum), truncated at k = 0.2hMpc⁻¹ in order to avoid non-linear scales, marginalising over the galaxy bias. For <u>Shear</u> data, we use the tomographic weak gravitational lensing analysis by the CFHTLenS (Kilbinger et al., 2013), which were shown to be usable in neutrino extensions of Λ CDM (Beutler et al., 2014). For the Clusters data, we use the thermal SZ measurements from cross-correlation of the CMB with X-ray clusters (Hajian et al., 2013), which are the most recent cluster-derived cosmological constraints. They rely on cross-correlations, and were also demonstrated to be robust to the choices in the modelling and data (tested with *Planck* and WMAP). We jointly use the *Planck* CMB temperature and Lensing power spectra (to probe the growth of structure with the CMB) with the BAO constraints (to constrain Ω_m). Finally, we also use the <u>RSD</u> measurements from BOSS (Beutler et al., 2013).

Tables 7.2 and 7.3 summarise the constraints on neutrino masses in the active and sterile neutrino scenarios, respectively, *i.e.*, $\Lambda \text{CDM} + \text{N}_{\text{eff}} + m_{\nu,\text{sterile}}^{\text{eff}}$ and $\Lambda \text{CDM} + \text{N}_{\text{eff}} + m_{\nu,\text{sterile}}^{\text{eff}}$ models,

Table 7.2: Marginalised 95% CL constraints on the Λ CDM+N_{eff} + $\sum m_{\nu}$ model from a variety of robust LSS datasets with the *Planck* CMB temperature and lensing measurements. These datasets are not in tension and tightly constrain the mass of active neutrinos.

	$\sum m_{\nu}$ [eV]	$N_{\rm eff}$
CMB+BAO	< 0.23	<3.88
CMB+Lensing+BAO	< 0.25	<3.84
CMB+Lensing+BAO+Clustering	< 0.26	<3.80
CMB+Lensing+BAO+Clusters	< 0.29	<3.78
CMB+Lensing+BAO+Shear	< 0.34	<3.79
CMB+Lensing+BAO+RSD	< 0.37	<3.75

Table 7.3: Same as Table 7.2, but for the Λ CDM+N_{eff} + $m_{\nu,\text{sterile}}^{\text{eff}}$ model, showing tight constraints on the mass of sterile neutrinos.

	$m_{\nu, \text{sterile}}^{\text{eff}}$ [eV]	$N_{\rm eff}$
CMB+BAO	< 0.28	<3.91
CMB+Lensing+BAO	< 0.35	<3.84
CMB+Lensing+BAO+Clustering	< 0.24	<3.87
CMB+Lensing+BAO+Clusters	< 0.33	<3.83
CMB+Lensing+BAO+Shear	< 0.51	<3.82
CMB+Lensing+BAO+RSD	< 0.59	<3.70

arising from a variety of data combinations. We see that multiple combinations yield similar constraints, and tend to small neutrino masses, *e.g.*, $\sum m_{\nu}$, $m_{\nu,\text{sterile}}^{\text{eff}} \lesssim 0.3 \text{ eV}$ at 95% CL. Note that some of these constraints may be relaxed by adding freedom to the model, for example to the primordial power spectrum (de Putter et al., 2014). Interestingly, as also noted by Beutler et al. (2014), the Shear and RSD data prefer lower σ_8 and thus, larger neutrino mass. However, the Bayes factors presented in the second section of Table 7.1 indicate a preference for the minimal Λ CDM model in all cases, even with the Shear and RSD data. Note that Beutler et al. (2014) marginalised over the lensing information which, as is well-known (Planck Collaboration, 2013k), leads to a preference for higher σ_8 ; conversely, our analysis combined the CMB temperature and lensing information.

Fig. 7.2 illustrates the persistence of the tension between the CMB+BAO, HST, PlaSZ and Xray data, as one extends the minimal Λ CDM model in the neutrino sector. The tension with local measurements of H_0 is alleviated by N_{eff} because of the degeneracy between these parameters (Feeney et al., 2013; Verde et al., 2013), but the tension with PlaSZ and Xray clusters persists despite the addition of both N_{eff} and neutrino masses. The levels of tension are comparable in minimal and extended models when adding Lensing and Clustering data. We note that the PlaSZ and Xray constraints were derived for the Λ CDM model, and it is unclear whether they can be used



Figure 7.2: Persistence of the tension as the minimal Λ CDM model is extended in the neutrino sector, *i.e.*, as N_{eff} and massive active or sterile neutrinos are added.

in the context of the extended models. In contrast, the datasets used in Tables 7.2 and 7.3 all relied on uncompressed likelihoods or constraints shown to be usable within the extended models.

Finally, sterile neutrinos were claimed (Dvorkin et al., 2014; Zhang et al., 2014) to also resolve the tension in the *Planck* measurements of the tensor-to-scalar ratio ($r_{0.002} < 0.11$ at 95% CL) and the recent BICEP result, $r_{0.002} = 0.2^{+0.07}_{-0.05}$ (BICEP2 Collaboration et al., 2014). However, the tension in the $\sigma_8-\Omega_m$ plane detailed previously persists in the extended model Λ CDM+ $r_{0.002}$ + N_{eff} + $m_{\nu,\text{sterile}}^{\text{eff}}$, as shown in Fig. 7.3. Hence, the relaxed constraints on $r_{0.002}$ from this data combination originates from a compromise between datasets in tension, not a new concordance. This is confirmed by the Bayes factors, presented in Table 7.4, showing that the extended model is not favoured over Λ CDM.

7.6. Conclusions.

The need for extra parameters yielding a new cosmological concordance can only be convincing if the combined datasets are in tension in the minimal model, and in agreement in extended model. We show that massive sterile neutrinos do not bring about a new cosmic concordance, but rather highlight the tension between the CMB+BAO and SZ or X-ray clusters. A compilation of current LSS data which have been demonstrated to be robust to modelling uncertainties, when combined



Figure 7.3: Constraints on the Λ CDM+ $r_{0.002}$ + N_{eff} + $m_{\nu,\text{sterile}}^{\text{eff}}$ model, illustrating the persisting tension between X-ray clusters and CMB+BAO in the σ_8 - Ω_m plane, despite an apparent reconciliation of the BICEP and *Planck* results on $r_{0.002}$.

Table 7.4: Evidence ratios $\ln[E_{\Lambda CDM}/E_{ext.}]$ between the minimal ΛCDM model and the $\Lambda CDM + r_{0.002} + N_{eff} + m_{\nu, sterile}^{eff}$ model, showing that sterile neutrinos are not favoured by the data, even when adding the BICEP results.

	Sterile
CMB+Lensing+BAO+Clustering	$2.89^{+0.13}_{-0.19}$
CMB+BAO+Xray+HST	$-0.70^{+0.07}_{-0.02}$
CMB+BAO+Xray+HST+BICEP	$-0.66\substack{+0.05\\-0.04}$

with *Planck*, tend to small masses $\sum m_{\nu}$, $m_{\nu,\text{sterile}}^{\text{eff}} \lesssim 0.3$ eV at 95% CL in the context of the Λ CDM model extended with N_{eff} and neutrino mass parameters. Similarly, as found in Feeney et al. (2013); Verde et al. (2013) the data cannot distinguish between $N_{\text{eff}} \sim 3$ and 4, and do not favour extra neutrinos over the standard 3 families. These conclusions are corroborated by the Bayesian evidence: the more complex models are not preferred, even when using datasets in tension. We conclude that current cosmological constraints do not provide evidence for large neutrino masses or extra neutrinos, even in the presence of the tension between *Planck* CMB and SZ and X-ray clusters. If this tension does not resolve after further investigation of systematic effects, new physics beyond massive neutrinos will be necessary to reconcile these datasets.
7.7. Acknowledgments

We thank Ofer Lahav, Stephen Feeney, Nina Roth, Florian Beutler, August Evrard, Raphael Flauger, Marta Spinelli, and Filipe Abdalla, for useful discussions. BL is supported by the Perren Fund and the IMPACT Fund. HVP is supported by STFC and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007- 2013) / ERC grant agreement no 306478-CosmicDawn. LV is supported by supported by the European Research Council under the European Community's Seventh Framework Programme FP7-IDEAS-Phys.LSS 240117 and Mineco grant FPA2011-29678- C02-02. Based on observations obtained with *Planck* (http://www.esa.int/Planck), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada.

Bibliography

- C. E. Aalseth, P. S. Barbeau, N. S. Bowden, B. Cabrera-Palmer, J. Colaresi, J. I. Collar, S. Dazeley, P. de Lurgio, J. E. Fast, N. Fields, C. H. Greenberg, T. W. Hossbach, M. E. Keillor, J. D. Kephart, M. G. Marino, H. S. Miley, M. L. Miller, J. L. Orrell, D. C. Radford, D. Reyna, O. Tench, T. D. van Wechel, J. F. Wilkerson, and K. M. Yocum. Results from a Search for Light-Mass Dark Matter with a p-Type Point Contact Germanium Detector. *Physical Review Letters*, 106(13):131301, Apr. 2011a. doi: 10.1103/PhysRevLett.106.131301.
- C. E. Aalseth, P. S. Barbeau, J. Colaresi, J. I. Collar, J. Diaz Leon, J. E. Fast, N. Fields, T. W. Hossbach, M. E. Keillor, J. D. Kephart, A. Knecht, M. G. Marino, H. S. Miley, M. L. Miller, J. L. Orrell, D. C. Radford, J. F. Wilkerson, and K. M. Yocum. Search for an Annual Modulation in a p-Type Point Contact Germanium Dark Matter Detector. *Physical Review Letters*, 107(14):141301, Sept. 2011b. doi: 10.1103/PhysRevLett.107.141301.
- K. Abazajian, M. Acero, S. Agarwalla, A. Aguilar-Arevalo, C. Albright, et al. Light Sterile Neutrinos: A White Paper. 2012.
- P. A. Abell et al. LSST Science Book, Version 2.0. ArXiv e-prints, 2009.
- L. R. Abramo, P. H. Reimberg, and H. S. Xavier. CMB in a box: Causal structure and the Fourier-Bessel expansion. *Physical Review D.*, 82(4):043510, Aug. 2010. doi: 10.1103/PhysRevD.82. 043510.
- V. Acquaviva, N. Bartolo, S. Matarrese, and A. Riotto. Gauge-invariant second-order perturbations and non-Gaussianity from inflation. *Nuclear Physics B*, 667:119–148, Sept. 2003. doi: 10.1016/ S0550-3213(03)00550-9.
- J. K. Adelman-McCarthy, M. A. Agüeros, S. S. Allam, C. Allende Prieto, K. S. J. Anderson, S. F. Anderson, J. Annis, N. A. Bahcall, C. A. L. Bailer-Jones, I. K. Baldry, J. C. Barentine, B. A. Bassett, A. C. Becker, T. C. Beers, E. F. Bell, A. A. Berlind, M. Bernardi, M. R. Blanton, J. J. Bochanski, W. N. Boroski, J. Brinchmann, J. Brinkmann, R. J. Brunner, T. Budavári, S. Carliles, M. A. Carr, F. J. Castander, D. Cinabro, R. J. Cool, K. R. Covey, I. Csabai, C. E. Cunha, J. R. A. Davenport, B. Dilday, M. Doi, D. J. Eisenstein, M. L. Evans, X. Fan, D. P. Finkbeiner, S. D. Friedman, J. A. Frieman, M. Fukugita, B. T. Gänsicke, E. Gates, B. Gillespie, K. Glazebrook, J. Gray, E. K. Grebel, J. E. Gunn, V. K. Gurbani, P. B. Hall, P. Harding, M. Harvanek, S. L. Hawley, J. Hayes, T. M. Heckman, J. S. Hendry, R. B. Hindsley, C. M. Hirata, C. J. Hogan, D. W. Hogg, J. B. Hyde, S.-i. Ichikawa, Ž. Ivezić, S. Jester, J. A. Johnson, A. M. Jorgensen, M. Jurić, S. M. Kent, R. Kessler, S. J. Kleinman, G. R. Knapp, R. G. Kron, J. Krzesinski, N. Kuropatkin, D. Q. Lamb, H. Lampeitl, S. Lebedeva, Y. S. Lee, R. F. Leger, S. Lépine, M. Lima, H. Lin, D. C. Long, C. P. Loomis, J. Loveday, R. H. Lupton, O. Malanushenko, V. Malanushenko, R. Mandelbaum, B. Margon, J. P. Marriner,

D. Martínez-Delgado, T. Matsubara, P. M. McGehee, T. A. McKay, A. Meiksin, H. L. Morrison, J. A. Munn, R. Nakajima, E. H. Neilsen, Jr., H. J. Newberg, R. C. Nichol, T. Nicinski, M. Nieto-Santisteban, A. Nitta, S. Okamura, R. Owen, H. Oyaizu, N. Padmanabhan, K. Pan, C. Park, J. Peoples, Jr., J. R. Pier, A. C. Pope, N. Purger, M. J. Raddick, P. Re Fiorentin, G. T. Richards, M. W. Richmond, A. G. Riess, H.-W. Rix, C. M. Rockosi, M. Sako, D. J. Schlegel, D. P. Schneider, M. R. Schreiber, A. D. Schwope, U. Seljak, B. Sesar, E. Sheldon, K. Shimasaku, T. Sivarani, J. A. Smith, S. A. Snedden, M. Steinmetz, M. A. Strauss, M. SubbaRao, Y. Suto, A. S. Szalay, I. Szapudi, P. Szkody, M. Tegmark, A. R. Thakar, C. A. Tremonti, D. L. Tucker, A. Uomoto, D. E. Vanden Berk, J. Vandenberg, S. Vidrih, M. S. Vogeley, W. Voges, N. P. Vogt, Y. Wadadekar, D. H. Weinberg, A. A. West, S. D. M. White, B. C. Wilhite, B. Yanny, D. R. Yocum, D. G. York, I. Zehavi, and D. B. Zucker. The Sixth Data Release of the Sloan Digital Sky Survey. *Astrophys. J. Supp.*, 175:297–313, Apr. 2008. doi: 10.1086/524984.

- N. Agarwal, S. Ho, A. D. Myers, H.-J. Seo, A. J. Ross, et al. Characterizing unknown systematics in large scale structure surveys. *Journal of Cosmology and Astroparticle Physics*, 1404:007, 2014a. doi: 10.1088/1475-7516/2014/04/007.
- N. Agarwal, S. Ho, and S. Shandera. Constraining the initial conditions of the Universe using large scale structure. *Journal of Cosmology and Astroparticle Physics*, 1402:038, 2014b. doi: 10.1088/1475-7516/2014/02/038.
- I. Agullo and S. Shandera. Large non-Gaussian halo bias from single field inflation. Journal of Cosmology and Astroparticle Physics, 9:007, Sept. 2012. doi: 10.1088/1475-7516/2012/09/007.
- A. Albrecht and P. J. Steinhardt. Cosmology for grand unified theories with radiatively induced symmetry breaking. *Phys. Rev. Lett.*, 48:1220–1223, Apr 1982. doi: 10.1103/PhysRevLett.48.1220. URL http://link.aps.org/doi/10.1103/PhysRevLett.48.1220.
- R. Allahverdi, R. Brandenberger, F.-Y. Cyr-Racine, and A. Mazumdar. Reheating in Inflationary Cosmology: Theory and Applications. *Annual Review of Nuclear and Particle Science*, 60:27–51, Nov. 2010. doi: 10.1146/annurev.nucl.012809.104511.
- T. J. Allen, B. Grinstein, and M. B. Wise. Non-gaussian density perturbations in inflationary cosmologies. *Physics Letters B*, 197:66–70, Oct. 1987. doi: 10.1016/0370-2693(87)90343-1.
- L. Amendola et al. Cosmology and fundamental physics with the Euclid satellite. *Living Rev.Rel.*, 16: 6, 2013.
- L. Anderson, E. Aubourg, S. Bailey, D. Bizyaev, M. Blanton, A. S. Bolton, J. Brinkmann, J. R. Brownstein, A. Burden, A. J. Cuesta, L. A. N. da Costa, K. S. Dawson, R. de Putter, D. J. Eisenstein, J. E. Gunn, H. Guo, J.-C. Hamilton, P. Harding, S. Ho, K. Honscheid, E. Kazin, D. Kirkby, J.-P. Kneib, A. Labatie, C. Loomis, R. H. Lupton, E. Malanushenko, V. Malanushenko, R. Mandelbaum, M. Manera, C. Maraston, C. K. McBride, K. T. Mehta, O. Mena, F. Montesano, D. Muna, R. C. Nichol, S. E. Nuza, M. D. Olmstead, D. Oravetz, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, J. Parejko, I. Pâris, W. J. Percival, P. Petitjean, F. Prada, B. Reid, N. A. Roe, A. J. Ross, N. P. Ross, L. Samushia, A. G. Sánchez, D. J. Schlegel, D. P. Schneider, C. G. Scóccola, H.-J. Seo, E. S. Sheldon, A. Simmons, R. A. Skibba, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. V. Magaña, L. Verde, C. Wagner, D. A. Wake, B. A. Weaver, D. H. Weinberg, M. White, X. Xu, C. Yèche, I. Zehavi, and G.-B. Zhao. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Release 9 spectroscopic galaxy sample. *Mon. Not. Roy. Astron. Soc.*, 427:3435–3467, Dec. 2012. doi: 10.1111/j.1365-2966.2012.22066.x.
- L. Anderson, E. Aubourg, S. Bailey, F. Beutler, V. Bhardwaj, M. Blanton, A. S. Bolton, J. Brinkmann, J. R. Brownstein, A. Burden, C.-H. Chuang, A. J. Cuesta, K. S. Dawson, D. J. Eisenstein, S. Escoffier, J. E. Gunn, H. Guo, S. Ho, K. Honscheid, C. Howlett, D. Kirkby, R. H. Lupton, M. Manera, C. Maraston, C. K. McBride, O. Mena, F. Montesano, R. C. Nichol, S. E. Nuza, M. D. Olmstead,

N. Padmanabhan, N. Palanque-Delabrouille, J. Parejko, W. J. Percival, P. Petitjean, F. Prada, A. M. Price-Whelan, B. Reid, N. A. Roe, A. J. Ross, N. P. Ross, C. G. Sabiu, S. Saito, L. Samushia, A. G. Sanchez, D. J. Schlegel, D. P. Schneider, C. G. Scoccola, H.-J. Seo, R. A. Skibba, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. Vargas Magana, L. Verde, D. A. Wake, B. A. Weaver, D. H. Weinberg, M. White, X. Xu, C. Yeche, I. Zehavi, and G.-B. Zhao. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations in the Data Release 10 and 11 galaxy samples. *ArXiv e-prints*, Dec. 2013.

- J. Angle, E. Aprile, F. Arneodo, L. Baudis, A. Bernstein, A. Bolozdynya, P. Brusov, L. C. C. Coelho, C. E. Dahl, L. Deviveiros, A. D. Ferella, L. M. P. Fernandes, S. Fiorucci, R. J. Gaitskell, K. L. Giboni, R. Gomez, R. Hasty, L. Kastens, J. Kwong, J. A. M. Lopes, N. Madden, A. Manalaysay, A. Manzur, D. N. McKinsey, M. E. Monzani, K. Ni, U. Oberlack, J. Orboeck, G. Plante, R. Santorelli, J. M. F. Dos Santos, P. Shagin, T. Shutt, P. Sorensen, S. Schulte, C. Winant, and M. Yamashita. First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory. *Physical Review Letters*, 100(2):021303, Jan. 2008. doi: 10.1103/PhysRevLett.100.021303.
- J.-P. Antoine and P. Vandergheynst. Wavelets on the n-sphere and related manifolds. J. Math. Phys, 39(8):3987–4008, 1998.
- J.-P. Antoine and P. Vandergheynst. Wavelets on the 2-sphere: a group theoretical approach. *Applied* and *Computational Harmonic Analysis*, 7:1–30, 1999.
- P. Audet. Directional wavelet analysis on the sphere: Application to gravity and topography of the terrestrial planets. J. Geophys. Res., 116, Jan. 2011. doi: 10.1029/2010JE003710.
- P. Audet and C. L. Johnson. Lithospheric and thermal structure of the Moon from gravity and topography. *AGU Fall Meeting Abstracts*, page D1738, Dec. 2011.
- R. Aurich and S. Lustig. Can one reconstruct the masked cosmic microwave background sky? *Mon. Not. Roy. Astron. Soc.*, 411:124–136, Feb. 2011. doi: 10.1111/j.1365-2966.2010.17667.x.
- D. Babich, P. Creminelli, and M. Zaldarriaga. The shape of non-Gaussianities. *Journal of Cosmology* and Astroparticle Physics, 8:009, Aug. 2004. doi: 10.1088/1475-7516/2004/08/009.
- P. Baldi, G. Kerkyacharian, D. Marinucci, and D. Picard. Asymptotics for spherical needlets. *Annals of Statistics*, 37 No.3:1150–1171, 2009. doi: 10.1214/08-AOS601.
- R. B. Barreiro, M. P. Hobson, A. N. Lasenby, A. J. Banday, K. M. Górski, and G. Hinshaw. Testing the Gaussianity of the COBE DMR data with spherical wavelets. *Mon. Not. Roy. Astron. Soc.*, 318: 475–481, Oct. 2000. doi: 10.1046/j.1365-8711.2000.03772.x.
- N. Bartolo, E. Komatsu, S. Matarrese, and A. Riotto. Non-Gaussianity from inflation: theory and observations. *Physics Reports*, 402:103–266, Nov. 2004. doi: 10.1016/j.physrep.2004.08.022.
- S. Basak and J. Delabrouille. A needlet internal linear combination analysis of WMAP 7-year data: estimation of CMB temperature map and power spectrum. *Mon. Not. Roy. Astron. Soc.*, 419: 1163–1175, Jan. 2012. doi: 10.1111/j.1365-2966.2011.19770.x.
- R. A. Battye and A. Moss. Evidence for Massive Neutrinos from Cosmic Microwave Background and Lensing Observations. *Physical Review Letters*, 112(5):051303, Feb. 2014. doi: 10.1103/ PhysRevLett.112.051303.
- D. Baumann. TASI Lectures on Inflation. ArXiv e-prints, July 2009.
- A. Becker, D. Huterer, and K. Kadota. Scale-Dependent Non-Gaussianity as a Generalization of the Local Model. *Journal of Cosmology and Astroparticle Physics*, 1101:006, 2011. doi: 10.1088/1475-7516/2011/01/006.

- A. Becker, D. Huterer, and K. Kadota. Constraining Scale-Dependent Non-Gaussianity with Future Large-Scale Structure and the CMB. *Journal of Cosmology and Astroparticle Physics*, 1212:034, 2012. doi: 10.1088/1475-7516/2012/12/034.
- C. Bennett, D. Larson, J. Weiland, N. Jarosik, G. Hinshaw, et al. Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. 2012.
- C. L. Bennett, D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, E. Wollack, J. Dunkley, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *ArXiv e-prints*, Dec. 2012.
- J. Beringer, J. F. Arguin, R. M. Barnett, K. Copic, O. Dahl, D. E. Groom, C. J. Lin, J. Lys, H. Murayama, C. G. Wohl, W. M. Yao, P. A. Zyla, C. Amsler, M. Antonelli, D. M. Asner, H. Baer, H. R. Band, T. Basaglia, C. W. Bauer, J. J. Beatty, V. I. Belousov, E. Bergren, G. Bernardi, W. Bertl, S. Bethke, H. Bichsel, O. Biebel, E. Blucher, S. Blusk, G. Brooijmans, O. Buchmueller, R. N. Cahn, M. Carena, A. Ceccucci, D. Chakraborty, M. C. Chen, R. S. Chivukula, G. Cowan, G. D'Ambrosio, T. Damour, D. de Florian, A. de Gouvêa, T. DeGrand, P. de Jong, G. Dissertori, B. Dobrescu, M. Doser, M. Drees, D. A. Edwards, S. Eidelman, J. Erler, V. V. Ezhela, W. Fetscher, B. D. Fields, B. Foster, T. K. Gaisser, L. Garren, H. J. Gerber, G. Gerbier, T. Gherghetta, S. Golwala, M. Goodman, C. Grab, A. V. Gritsan, J. F. Grivaz, M. Grünewald, A. Gurtu, T. Gutsche, H. E. Haber, K. Hagiwara, C. Hagmann, C. Hanhart, S. Hashimoto, K. G. Hayes, M. Heffner, B. Heltsley, J. J. Hernández-Rey, K. Hikasa, A. Höcker, J. Holder, A. Holtkamp, J. Huston, J. D. Jackson, K. F. Johnson, T. Junk, D. Karlen, D. Kirkby, S. R. Klein, E. Klempt, R. V. Kowalewski, F. Krauss, M. Kreps, B. Krusche, Y. V. Kuyanov, Y. Kwon, O. Lahav, J. Laiho, P. Langacker, A. Liddle, Z. Ligeti, T. M. Liss, L. Littenberg, K. S. Lugovsky, S. B. Lugovsky, T. Mannel, A. V. Manohar, W. J. Marciano, A. D. Martin, A. Masoni, J. Matthews, D. Milstead, R. Miquel, K. Mönig, F. Moortgat, K. Nakamura, M. Narain, P. Nason, S. Navas, M. Neubert, P. Nevski, Y. Nir, K. A. Olive, L. Pape, J. Parsons, C. Patrignani, J. A. Peacock, S. T. Petcov, A. Piepke, A. Pomarol, G. Punzi, A. Quadt, S. Raby, G. Raffelt, B. N. Ratcliff, P. Richardson, S. Roesler, S. Rolli, A. Romaniouk, L. J. Rosenberg, J. L. Rosner, C. T. Sachrajda, Y. Sakai, G. P. Salam, S. Sarkar, F. Sauli, O. Schneider, K. Scholberg, D. Scott, W. G. Seligman, M. H. Shaevitz, S. R. Sharpe, M. Silari, T. Sjöstrand, P. Skands, J. G. Smith, G. F. Smoot, S. Spanier, H. Spieler, A. Stahl, T. Stanev, S. L. Stone, T. Sumiyoshi, M. J. Syphers, F. Takahashi, M. Tanabashi, J. Terning, M. Titov, N. P. Tkachenko, N. A. Törnqvist, D. Tovey, G. Valencia, K. van Bibber, G. Venanzoni, M. G. Vincter, P. Vogel, A. Vogt, W. Walkowiak, C. W. Walter, D. R. Ward, T. Watari, G. Weiglein, E. J. Weinberg, L. R. Wiencke, L. Wolfenstein, J. Womersley, C. L. Woody, R. L. Workman, A. Yamamoto, G. P. Zeller, O. V. Zenin, J. Zhang, R. Y. Zhu, G. Harper, V. S. Lugovsky, and P. Schaffner. Review of particle physics. Phys. Rev. D, 86:010001, Jul 2012. doi: 10. 1103/PhysRevD.86.010001. URL http://link.aps.org/doi/10.1103/PhysRevD.86.010001.
- R. Bernabei, P. Belli, F. Cappella, R. Cerulli, C. J. Dai, A. D'Angelo, H. L. He, A. Incicchitti, H. H. Kuang, J. M. Ma, F. Montecchia, F. Nozzoli, D. Prosperi, X. D. Sheng, and Z. P. Ye. First results from DAMA/LIBRA and the combined results with DAMA/NaI. *European Physical Journal C*, 56: 333, Aug. 2008. doi: 10.1140/epjc/s10052-008-0662-y.
- F. Beutler, C. Blake, M. Colless, D. H. Jones, L. Staveley-Smith, L. Campbell, Q. Parker, W. Saunders, and F. Watson. The 6dF Galaxy Survey: baryon acoustic oscillations and the local Hubble constant. *Mon. Not. Roy. Astron. Soc.*, 416:3017–3032, Oct. 2011. doi: 10.1111/j.1365-2966.2011.19250.x.
- F. Beutler, S. Saito, H.-J. Seo, J. Brinkmann, K. S. Dawson, D. J. Eisenstein, A. Font-Ribera, S. Ho, C. K. McBride, F. Montesano, W. J. Percival, A. J. Ross, N. P. Ross, L. Samushia, D. J. Schlegel, A. G. Sánchez, J. L. Tinker, and B. A. Weaver. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Testing gravity with redshift-space distortions using the power spectrum multipoles. *ArXiv e-prints*, Dec. 2013.
- F. Beutler, S. Saito, J. R. Brownstein, C.-H. Chuang, A. J. Cuesta, W. J. Percival, A. J. Ross, N. P. Ross, D. P. Schneider, L. Samushia, A. G. Sánchez, H.-J. Seo, J. L. Tinker, C. Wagner, and B. A.

Weaver. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: Signs of neutrino mass in current cosmological datasets. *ArXiv e-prints*, Mar. 2014.

- BICEP2 Collaboration, P. A. R. Ade, R. W. Aikin, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, J. A. Brevik, I. Buder, E. Bullock, C. D. Dowell, L. Duband, J. P. Filippini, S. Fliescher, S. R. Golwala, M. Halpern, M. Hasselfield, S. R. Hildebrandt, G. C. Hilton, V. V. Hristov, K. D. Irwin, K. S. Karkare, J. P. Kaufman, B. G. Keating, S. A. Kernasovskiy, J. M. Kovac, C. L. Kuo, E. M. Leitch, M. Lueker, P. Mason, C. B. Netterfield, H. T. Nguyen, R. O'Brient, R. W. Ogburn, IV, A. Orlando, C. Pryke, C. D. Reintsema, S. Richter, R. Schwarz, C. D. Sheehy, Z. K. Staniszewski, R. V. Sudiwala, G. P. Teply, J. E. Tolan, A. D. Turner, A. G. Vieregg, C. L. Wong, and K. W. Yoon. BICEP2 I: Detection Of B-mode Polarization at Degree Angular Scales. *ArXiv e-prints*, Mar. 2014.
- C. Blake, S. Brough, M. Colless, W. Couch, S. Croom, et al. The WiggleZ Dark Energy Survey: the selection function and z=0.6 galaxy power spectrum. *Mon. Not. Roy. Astron. Soc.*, 406:803–821, 2010.
- C. Blake, E. Kazin, F. Beutler, T. Davis, D. Parkinson, et al. The WiggleZ Dark Energy Survey: mapping the distance-redshift relation with baryon acoustic oscillations. *Mon. Not. Roy. Astron. Soc.*, 418:1707–1724, 2011. doi: 10.1111/j.1365-2966.2011.19592.x.
- C. Bond, P. Fulda, L. Carbone, K. Kokeyama, and A. Freise. Higher order Laguerre-Gauss mode degeneracy in realistic, high finesse cavities. *Physical Review D.*, 84(10):102002, Nov. 2011.
- J. R. Bond, A. H. Jaffe, and L. Knox. Estimating the power spectrum of the cosmic microwave background. *Physical Review D.*, 57:2117–2137, Feb. 1998. doi: 10.1103/PhysRevD.57.2117.
- J. R. Bond, A. H. Jaffe, and L. Knox. Radical Compression of Cosmic Microwave Background Data. *Astrophys. J.*, 533:19–37, Apr. 2000. doi: 10.1086/308625.
- J. Bovy, J. F. Hennawi, D. W. Hogg, A. D. Myers, J. A. Kirkpatrick, D. J. Schlegel, N. P. Ross, E. S. Sheldon, I. D. McGreer, D. P. Schneider, and B. A. Weaver. Think Outside the Color Box: Probabilistic Target Selection and the SDSS-XDQSO Quasar Targeting Catalog. *Astrophys. J.*, 729 (2):141, 2011a.
- J. Bovy, D. W. Hogg, and S. T. Roweis. Extreme deconvolution: Inferring complete distribution functions from noisy, heterogeneous and incomplete observations. *Annals of Applied Statistics*, 5: 1657–1677, jun 2011b. doi: 10.1214/10-AOAS439.
- J. Bovy, A. D. Myers, J. F. Hennawi, D. W. Hogg, R. G. McMahon, D. Schiminovich, E. S. Sheldon, J. Brinkmann, D. P. Schneider, and B. A. Weaver. Photometric Redshifts and Quasar Probabilities from a Single, Data-driven Generative Model. *Astrophys. J.*, 749:41, Apr. 2012. doi: 10.1088/ 0004-637X/749/1/41.
- M. L. Brown, P. G. Castro, and A. N. Taylor. Cosmic microwave background temperature and polarization pseudo-C_l estimators and covariances. *Mon. Not. Roy. Astron. Soc.*, 360:1262–1280, July 2005. doi: 10.1111/j.1365-2966.2005.09111.x.
- C. Carbone, L. Verde, and S. Matarrese. Non-Gaussian Halo Bias and Future Galaxy Surveys. *Astrophys. J. Lett.*, 684:L1–L4, Sept. 2008. doi: 10.1086/592020.
- S. M. Carroll. The Cosmological Constant. *Living Reviews in Relativity*, 4:1, Feb. 2001. doi: 10.12942/lrr-2001-1.
- L. Cayón, J. L. Sanz, E. Martínez-González, A. J. Banday, F. Argüeso, J. E. Gallegos, K. M. Górski, and G. Hinshaw. Spherical Mexican hat wavelet: an application to detect non-Gaussianity in the COBE-DMR maps. *Mon. Not. Roy. Astron. Soc.*, 326:1243–1248, Oct. 2001. doi: 10.1046/j. 1365-8711.2001.04641.x.
- A. Challinor and A. Lewis. Linear power spectrum of observed source number counts. *Phys. Rev. D*, 84(4):043516, Aug. 2011. doi: 10.1103/PhysRevD.84.043516.

- J. Charléty, G. Nolet, S. Voronin, I. Loris, F. J. Simons, I. Daubechies, and K. Sigloch. Inversion with a sparsity constraint: application to mantle tomography. In *EGU General Assembly Conference Abstracts*, EGU General Assembly Conference Abstracts, 2012.
- A. Chow. Orthogonal and symmetric haar wavelets on the three-dimensional ball. Master's thesis, University of Toronto, Canada, 2010.
- T. Clifton, P. G. Ferreira, A. Padilla, and C. Skordis. Modified gravity and cosmology. *Physics Reports*, 513:1–189, Mar. 2012. doi: 10.1016/j.physrep.2012.01.001.
- D. Clowe, A. Gonzalez, and M. Markevitch. Weak-Lensing Mass Reconstruction of the Interacting Cluster 1E 0657-558: Direct Evidence for the Existence of Dark Matter. *Astrophys. J.*, 604: 596–603, Apr. 2004. doi: 10.1086/381970.
- D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky. A Direct Empirical Proof of the Existence of Dark Matter. *Astrophys. J.*, 648:L109–L113, Sept. 2006. doi: 10.1086/508162.
- A. A. Collister and O. Lahav. ANNz: Estimating Photometric Redshifts Using Artificial Neural Networks. *The Publications of the Astronomical Society of the Pacific*, 116:345–351, Apr. 2004. doi: 10.1086/383254.
- J. M. Conrad, W. C. Louis, and M. H. Shaevitz. The LSND and MiniBooNE Oscillation Searches at High Δm^2 . Ann.Rev.Nucl.Part.Sci., 63:45–67, 2013. doi: 10.1146/annurev-nucl-102711-094957.
- C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman. Bias in low-multipole CMB reconstructions. *ArXiv e-prints*, Mar. 2011.
- P.-S. Corasaniti, T. Giannantonio, and A. Melchiorri. Constraining dark energy with cross-correlated cmb and large scale structure data. *Phys. Rev. D*, 71:123521, Jun 2005. doi: 10.1103/PhysRevD. 71.123521. URL http://link.aps.org/doi/10.1103/PhysRevD.71.123521.
- J.-P. Couderc and W. Zareba. Contribution of the wavelet analysis to the noninvasive electrocardiology. *Annals of Noninvasive Electrocardiology*, 3(1):54–62, 1998. ISSN 1542-474X. doi: 10.1111/j.1542-474X.1998.tb00030.x. URL http://dx.doi.org/10.1111/j.1542-474X.1998. tb00030.x.
- R. T. Cox. Probability, frequency and reasonable expectation. Am.J.Th.Phys, 14:1, 1946.
- M. Crocce, S. Pueblas, and R. Scoccimarro. Transients from initial conditions in cosmological simulations. *Mon. Not. Roy. Astron. Soc.*, 373:369–381, Nov. 2006. doi: 10.1111/j.1365-2966. 2006.11040.x.
- S. M. Croom, G. T. Richards, T. Shanks, B. J. Boyle, R. G. Sharp, J. Bland-Hawthorn, T. Bridges, R. J. Brunner, R. Cannon, D. Carson, K. Chiu, M. Colless, W. Couch, R. de Propris, M. J. Drinkwater, A. Edge, S. Fine, J. Loveday, L. Miller, A. D. Myers, R. C. Nichol, P. Outram, K. Pimbblet, I. Roseboom, N. Ross, D. P. Schneider, A. Smith, C. Stoughton, M. A. Strauss, and D. Wake. The 2dF-SDSS LRG and QSO Survey: the spectroscopic QSO catalogue. *Mon. Not. Roy. Astron. Soc.*, 392:19–44, Jan. 2009. doi: 10.1111/j.1365-2966.2008.14052.x.
- N. Dalal, O. Doré, D. Huterer, and A. Shirokov. Imprints of primordial non-Gaussianities on largescale structure: Scale-dependent bias and abundance of virialized objects. *Physical Review D.*, 77 (12):123514, June 2008. doi: 10.1103/PhysRevD.77.123514.
- S. Das, B. D. Sherwin, P. Aguirre, J. W. Appel, J. R. Bond, C. S. Carvalho, M. J. Devlin, J. Dunkley, R. Dünner, T. Essinger-Hileman, J. W. Fowler, A. Hajian, M. Halpern, M. Hasselfield, A. D. Hincks, R. Hlozek, K. M. Huffenberger, J. P. Hughes, K. D. Irwin, J. Klein, A. Kosowsky, R. H. Lupton, T. A. Marriage, D. Marsden, F. Menanteau, K. Moodley, M. D. Niemack, M. R. Nolta, L. A. Page, L. Parker, E. D. Reese, B. L. Schmitt, N. Sehgal, J. Sievers, D. N. Spergel, S. T. Staggs, D. S. Swetz,

E. R. Switzer, R. Thornton, K. Visnjic, and E. Wollack. Detection of the Power Spectrum of Cosmic Microwave Background Lensing by the Atacama Cosmology Telescope. *Physical Review Letters*, 107(2):021301, July 2011. doi: 10.1103/PhysRevLett.107.021301.

- S. Das, T. Louis, M. R. Nolta, G. E. Addison, E. S. Battistelli, J. Bond, E. Calabrese, D. C. M. J. Devlin, S. Dicker, J. Dunkley, R. Dünner, J. W. Fowler, M. Gralla, A. Hajian, M. Halpern, M. Hasselfield, M. Hilton, A. D. Hincks, R. Hlozek, K. M. Huffenberger, J. P. Hughes, K. D. Irwin, A. Kosowsky, R. H. Lupton, T. A. Marriage, D. Marsden, F. Menanteau, K. Moodley, M. D. Niemack, L. A. Page, B. Partridge, E. D. Reese, B. L. Schmitt, N. Sehgal, B. D. Sherwin, J. L. Sievers, D. N. Spergel, S. T. Staggs, D. S. Swetz, E. R. Switzer, R. Thornton, H. Trac, and E. Wollack. The Atacama Cosmology Telescope: Temperature and Gravitational Lensing Power Spectrum Measurements from Three Seasons of Data. *ArXiv e-prints*, Jan. 2013.
- I. Daubechies. *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1992a. ISBN 0-89871-274-2.
- I. Daubechies. Ten Lectures on Wavelets. Society for Industrial and Applied Mathematics, 1992b. doi: 10.1137/1.9781611970104. URL http://epubs.siam.org/doi/abs/10.1137/ 1.9781611970104.
- M. Davis, J. Huchra, D. W. Latham, and J. Tonry. A survey of galaxy redshifts. II The large scale space distribution. *Astrophys. J.*, 253:423–445, Feb. 1982. doi: 10.1086/159646.
- K. S. Dawson, D. J. Schlegel, C. P. Ahn, S. F. Anderson, É. Aubourg, S. Bailey, R. H. Barkhouser, J. E. Bautista, A. Beifiori, A. A. Berlind, V. Bhardwaj, D. Bizyaev, C. H. Blake, M. R. Blanton, M. Blomqvist, A. S. Bolton, A. Borde, J. Bovy, W. N. Brandt, H. Brewington, J. Brinkmann, P. J. Brown, J. R. Brownstein, K. Bundy, N. G. Busca, W. Carithers, A. R. Carnero, M. A. Carr, Y. Chen, J. Comparat, N. Connolly, F. Cope, R. A. C. Croft, A. J. Cuesta, L. N. da Costa, J. R. A. Davenport, T. Delubac, R. de Putter, S. Dhital, A. Ealet, G. L. Ebelke, D. J. Eisenstein, S. Escoffier, X. Fan, N. Filiz Ak, H. Finley, A. Font-Ribera, R. Génova-Santos, J. E. Gunn, H. Guo, D. Haggard, P. B. Hall, J.-C. Hamilton, B. Harris, D. W. Harris, S. Ho, D. W. Hogg, D. Holder, K. Honscheid, J. Huehnerhoff, B. Jordan, W. P. Jordan, G. Kauffmann, E. A. Kazin, D. Kirkby, M. A. Klaene, J.-P. Kneib, J.-M. Le Goff, K.-G. Lee, D. C. Long, C. P. Loomis, B. Lundgren, R. H. Lupton, M. A. G. Maia, M. Makler, E. Malanushenko, V. Malanushenko, R. Mandelbaum, M. Manera, C. Maraston, D. Margala, K. L. Masters, C. K. McBride, P. McDonald, I. D. McGreer, R. G. McMahon, O. Mena, J. Miralda-Escudé, A. D. Montero-Dorta, F. Montesano, D. Muna, A. D. Myers, T. Naugle, R. C. Nichol, P. Noterdaeme, S. E. Nuza, M. D. Olmstead, A. Oravetz, D. J. Oravetz, R. Owen, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, J. K. Parejko, I. Pâris, W. J. Percival, I. Pérez-Fournon, I. Pérez-Ràfols, P. Petitjean, R. Pfaffenberger, J. Pforr, M. M. Pieri, F. Prada, A. M. Price-Whelan, M. J. Raddick, R. Rebolo, J. Rich, G. T. Richards, C. M. Rockosi, N. A. Roe, A. J. Ross, N. P. Ross, G. Rossi, J. A. Rubiño-Martin, L. Samushia, A. G. Sánchez, C. Sayres, S. J. Schmidt, D. P. Schneider, C. G. Scóccola, H.-J. Seo, A. Shelden, E. Sheldon, Y. Shen, Y. Shu, A. Slosar, S. A. Smee, S. A. Snedden, F. Stauffer, O. Steele, M. A. Strauss, A. Streblyanska, N. Suzuki, M. E. C. Swanson, T. Tal, M. Tanaka, D. Thomas, J. L. Tinker, R. Tojeiro, C. A. Tremonti, M. Vargas Magaña, L. Verde, M. Viel, D. A. Wake, M. Watson, B. A. Weaver, D. H. Weinberg, B. J. Weiner, A. A. West, M. White, W. M. Wood-Vasey, C. Yeche, I. Zehavi, G.-B. Zhao, and Z. Zheng. The Baryon Oscillation Spectroscopic Survey of SDSS-III. Astrophys. J., 145:10, Jan. 2013. doi: 10.1088/0004-6256/145/1/10.
- R. de Putter, E. V. Linder, and A. Mishra. Inflationary Freedom and Cosmological Neutrino Constraints. *ArXiv e-prints*, Jan. 2014.
- E. Deriaz, J.-L. Starck, and S. Pires. Wavelet Helmholtz decomposition for weak lensing mass map reconstruction. *Astron. & Astrophys.*, 540:A34, Apr. 2012. doi: 10.1051/0004-6361/201117129.
- V. Desjacques and U. Seljak. Primordial non-Gaussianity from the large-scale structure. *Classical and Quantum Gravity*, 27(12):124011, June 2010a. doi: 10.1088/0264-9381/27/12/124011.

- V. Desjacques and U. Seljak. Signature of primordial non-Gaussianity of ϕ^3 type in the mass function and bias of dark matter haloes. *Phys. Rev. D*, 81(2):023006, Jan. 2010b. doi: 10.1103/PhysRevD. 81.023006.
- M. Dias, R. H. Ribeiro, and D. Seery. Scale-dependent bias from multiple-field inflation. *Phys. Rev. D*, 87(10):107301, May 2013. doi: 10.1103/PhysRevD.87.107301.
- R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson. Cosmic Black-Body Radiation. *Astrophys. J.*, 142:414–419, July 1965. doi: 10.1086/148306.
- J. P. Dietrich, N. Werner, D. Clowe, A. Finoguenov, T. Kitching, L. Miller, and A. Simionescu. A filament of dark matter between two clusters of galaxies. *Nature*, 487:202–204, July 2012. doi: 10.1038/nature11224.
- S. Dodelson. Modern cosmology. 2003.
- A. G. Doroshkevich, P. D. Naselsky, O. V. Verkhodanov, D. I. Novikov, V. I. Turchaninov, I. D. Novikov, P. R. Christensen, and L.-Y. Chiang. Gauss-Legendre Sky Pixelization (glesp) for CMB Maps. *International Journal of Modern Physics D*, 14:275–290, 2005. doi: 10.1142/S0218271805006183.
- A. G. Doroshkevich, O. V. Verkhodanov, P. D. Naselsky, J. Kim, D. I. Novikov, V. I. Turchaninov, I. D. Novikov, L.-Y. Chiang, and M. Hansen. The Gauss-Legendre Sky Pixelization for the CMB Polarization Glesp-Pol Errors due to Pixelization of the CMB Sky. *International Journal of Modern Physics D*, 20:1053–1078, 2011. doi: 10.1142/S0218271811019219.
- J. R. Driscoll and D. M. J. Healy. Computing Fourier transforms and convolutions on the sphere. *Advances in Applied Mathematics*, 15:202–250, 1994.
- C. F. Dunkl. A Laguerre polynomial orthogonality and the hydrogen atom. *Analysis & Appl.*, 1: 177–178, 2003.
- J. Dunkley, R. Hlozek, J. Sievers, V. Acquaviva, P. A. R. Ade, P. Aguirre, M. Amiri, J. W. Appel, L. F. Barrientos, E. S. Battistelli, J. R. Bond, B. Brown, B. Burger, J. Chervenak, S. Das, M. J. Devlin, S. R. Dicker, W. Bertrand Doriese, R. Dünner, T. Essinger-Hileman, R. P. Fisher, J. W. Fowler, A. Hajian, M. Halpern, M. Hasselfield, C. Hernández-Monteagudo, G. C. Hilton, M. Hilton, A. D. Hincks, K. M. Huffenberger, D. H. Hughes, J. P. Hughes, L. Infante, K. D. Irwin, J. B. Juin, M. Kaul, J. Klein, A. Kosowsky, J. M. Lau, M. Limon, Y.-T. Lin, R. H. Lupton, T. A. Marriage, D. Marsden, P. Mauskopf, F. Menanteau, K. Moodley, H. Moseley, C. B. Netterfield, M. D. Niemack, M. R. Nolta, L. A. Page, L. Parker, B. Partridge, B. Reid, N. Sehgal, B. Sherwin, D. N. Spergel, S. T. Staggs, D. S. Swetz, E. R. Switzer, R. Thornton, H. Trac, C. Tucker, R. Warne, E. Wollack, and Y. Zhao. The Atacama Cosmology Telescope: Cosmological Parameters from the 2008 Power Spectrum. *Astrophys. J.*, 739:52, Sept. 2011. doi: 10.1088/0004-637X/739/1/52.
- C. Dvorkin, M. Wyman, D. H. Rudd, and W. Hu. Neutrinos help reconcile Planck measurements with both Early and Local Universe. 2014.
- G. Efstathiou. A maximum likelihood analysis of the low cosmic microwave background multipoles from the Wilkinson Microwave Anisotropy Probe. *Mon. Not. Roy. Astron. Soc.*, 348:885–896, Mar. 2004a. doi: 10.1111/j.1365-2966.2004.07409.x.
- G. Efstathiou. Myths and truths concerning estimation of power spectra: the case for a hybrid estimator. *Mon. Not. Roy. Astron. Soc.*, 349:603–626, Apr. 2004b. doi: 10.1111/j.1365-2966. 2004.07530.x.
- G. Efstathiou. Hybrid estimation of cosmic microwave background polarization power spectra. *Mon. Not. Roy. Astron. Soc.*, 370:343–362, July 2006. doi: 10.1111/j.1365-2966.2006.10486.x.
- G. Efstathiou. H0 Revisited. 2013.

- D. J. Eisenstein, D. H. Weinberg, E. Agol, H. Aihara, C. Allende Prieto, S. F. Anderson, J. A. Arns, É. Aubourg, S. Bailey, E. Balbinot, and et al. SDSS-III: Massive Spectroscopic Surveys of the Distant Universe, the Milky Way, and Extra-Solar Planetary Systems. *Astrophys. J.*, 142:72, Sept. 2011. doi: 10.1088/0004-6256/142/3/72.
- H. K. Eriksen, C. Dickinson, C. R. Lawrence, C. Baccigalupi, A. J. Banday, K. M. Górski, F. K. Hansen, P. B. Lilje, E. Pierpaoli, M. D. Seiffert, K. M. Smith, and K. Vanderlinde. Cosmic Microwave Background Component Separation by Parameter Estimation. *Astrophys. J.*, 641:665–682, Apr. 2006. doi: 10.1086/500499.
- H. K. Eriksen, G. Huey, R. Saha, F. K. Hansen, J. Dick, A. J. Banday, K. M. Górski, P. Jain, J. B. Jewell, L. Knox, D. L. Larson, I. J. O'Dwyer, T. Souradeep, and B. D. Wandelt. A Reanalysis of the 3 Year Wilkinson Microwave Anisotropy Probe Temperature Power Spectrum and Likelihood. *Astrophys. J.*, 656:641–652, Feb. 2007. doi: 10.1086/509911.
- H. K. Eriksen, J. B. Jewell, C. Dickinson, A. J. Banday, K. M. Górski, and C. R. Lawrence. Joint Bayesian Component Separation and CMB Power Spectrum Estimation. *Astrophys. J.*, 676:10–32, Mar. 2008. doi: 10.1086/525277.
- T. Falk, R. Rangarajan, and M. Srednicki. The angular dependence of the three-point correlation function of the cosmic microwave background radiation as predicted by inflationary cosmologies. *Astrophys. J. Lett.*, 403:L1–L3, Jan. 1993. doi: 10.1086/186707.
- X.-H. Fan, M. A. Strauss, R. H. Becker, R. L. White, J. E. Gunn, et al. Constraining the evolution of the ionizing background and the epoch of reionization with z 6 quasars. 2. a sample of 19 quasars. *Astron.J.*, 132:117–136, 2006. doi: 10.1086/504836.
- N. Fanidakis, A. V. Macciò, C. M. Baugh, C. G. Lacey, and C. S. Frenk. The most luminous quasars do not live in the most massive dark matter haloes at any redshift. *Mon. Not. Roy. Astron. Soc.*, 436:315–326, Nov. 2013. doi: 10.1093/mnras/stt1567.
- G. Faÿ, F. Guilloux, M. Betoule, J.-F. Cardoso, J. Delabrouille, and M. Le Jeune. CMB power spectrum estimation using wavelets. *Physical Review D.*, 78(8):083013, Oct. 2008. doi: 10.1103/PhysRevD. 78.083013.
- S. M. Feeney, M. C. Johnson, D. J. Mortlock, and H. V. Peiris. First observational tests of eternal inflation: Analysis methods and WMAP 7-year results. *Physical Review D.*, 84(4):043507, Aug. 2011a. doi: 10.1103/PhysRevD.84.043507.
- S. M. Feeney, H. V. Peiris, and A. Pontzen. Avoiding bias in reconstructing the largest observable scales from partial-sky data. *Physical Review D.*, 84(10):103002, Nov. 2011b. doi: 10.1103/ PhysRevD.84.103002.
- S. M. Feeney, M. C. Johnson, J. D. McEwen, D. J. Mortlock, and H. V. Peiris. Hierarchical Bayesian detection algorithm for early-universe relics in the cosmic microwave background. *Phys. Rev. D*, 88(4):043012, Aug. 2013. doi: 10.1103/PhysRevD.88.043012.
- S. M. Feeney, H. V. Peiris, and L. Verde. Is there evidence for additional neutrino species from cosmology? *Journal of Cosmology and Astroparticle Physics*, 1304:036, 2013. doi: 10.1088/ 1475-7516/2013/04/036.
- M. Fengler, D. Michel, and V. Michel. Harmonic spline-wavelets on the 3 dimensional ball and their application to the reconstruction of the earth's density distribution from gravitational data at arbitrarily shaped satellite orbits. *ZAMM Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 86(11):856–873, 2006. ISSN 1521-4001. doi: 10.1002/zamm.200510277.
- D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. *Publications of the Astronomical Society of the Pacific*, 125:306–312, Mar. 2013. doi: 10.1086/670067.

- W. L. Freedman, B. F. Madore, B. K. Gibson, L. Ferrarese, D. D. Kelson, S. Sakai, J. R. Mould, R. C. Kennicutt, Jr., H. C. Ford, J. A. Graham, J. P. Huchra, S. M. G. Hughes, G. D. Illingworth, L. M. Macri, and P. B. Stetson. Final Results from the Hubble Space Telescope Key Project to Measure the Hubble Constant. *Astrophys. J.*, 553:47–72, May 2001. doi: 10.1086/320638.
- A. Gangui and J. Martin. Cosmic microwave background bispectrum and slow-roll inflation. Mon. Not. Roy. Astron. Soc., 313:323–330, Apr. 2000. doi: 10.1046/j.1365-8711.2000.03210.x.
- A. Gangui, F. Lucchin, S. Matarrese, and S. Mollerach. The three-point correlation function of the cosmic microwave background in inflationary models. *Astrophys. J.*, 430:447–457, Aug. 1994. doi: 10.1086/174421.
- E. GÂólich and C. Markett. A convolution structure for Laguerre series. *Indagationes Mathematicae* (*Proceedings*), 85(2):161 – 171, 1982. ISSN 1385-7258. doi: 10.1016/S1385-7258(82)80008-5.
- D. Geller and D. Marinucci. Mixed Needlets. ArXiv e-prints, June 2010.
- D. Geller, F. K. Hansen, D. Marinucci, G. Kerkyacharian, and D. Picard. Spin needlets for cosmic microwave background polarization data analysis. *Physical Review D.*, 78(12):123533, Dec. 2008. doi: 10.1103/PhysRevD.78.123533.
- T. Giannantonio and W. J. Percival. Using correlations between cosmic microwave background lensing and large-scale structure to measure primordial non-Gaussianity. *Mon. Not. Roy. Astron. Soc.*, 441:L16–L20, June 2014. doi: 10.1093/mnrasl/slu036.
- T. Giannantonio, R. G. Crittenden, B. Nichol, R. Scranton, G. Richards, A. Myers, R. Brunner, A. Gray, A. Connolly, and D. Schneider. High redshift detection of the integrated sachs-wolfe effect. *Physical Review D.*, 74(6):063520, September 2006. URL http://eprints.port.ac.uk/6682/.
- T. Giannantonio, R. Scranton, R. G. Crittenden, R. C. Nichol, S. P. Boughn, A. D. Myers, and G. T. Richards. Combined analysis of the integrated Sachs-Wolfe effect and cosmological implications. *Physical Review D.*, 77(12):123520, June 2008. doi: 10.1103/PhysRevD.77.123520.
- T. Giannantonio, R. Crittenden, R. Nichol, and A. J. Ross. The significance of the integrated sachs-wolfe effect revisited. *Mon. Not. Roy. Astron. Soc.*, 426(3):2581-2599, 2012. doi: 10. 1111/j.1365-2966.2012.21896.x. URL http://mnras.oxfordjournals.org/content/426/3/ 2581.abstract.
- T. Giannantonio, C. Porciani, J. Carron, A. Amara, and A. Pillepich. Constraining primordial non-Gaussianity with future galaxy surveys. *Mon. Not. Roy. Astron. Soc.*, 422:2854–2877, June 2012. doi: 10.1111/j.1365-2966.2012.20604.x.
- T. Giannantonio, A. J. Ross, W. J. Percival, R. Crittenden, D. Bacher, M. Kilbinger, R. Nichol, and J. Weller. Improved primordial non-Gaussianity constraints from measurements of galaxy clustering and the integrated Sachs-Wolfe effect. *Phys. Rev. D*, 89(2):023511, Jan. 2014. doi: 10.1103/PhysRevD.89.023511.
- E. Giusarma, R. de Putter, S. Ho, and O. Mena. Constraints on neutrino masses from Planck and Galaxy clustering data. *Physical Review D.*, 88(6):063515, Sept. 2013. doi: 10.1103/PhysRevD. 88.063515.
- E. Giusarma, E. Di Valentino, M. Lattanzi, A. Melchiorri, and O. Mena. Relic Neutrinos, thermal axions and cosmology in early 2014. 2014.
- M. Gonzalez-Garcia, M. Maltoni, J. Salvado, and T. Schwetz. Global fit to three neutrino mixing: critical look at present precision. *JHEP*, 1212:123, 2012. doi: 10.1007/JHEP12(2012)123.
- K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *Astrophys. J.*, 622:759–771, Apr. 2005. doi: 10.1086/427976.

- M. M. Graça and M. Esmeralda Sousa-Dias. A unified framework for the computation of polynomial quadrature weights and errors. *Arxiv preprint 1203.4795*, Mar. 2012.
- P. Graff, F. Feroz, M. P. Hobson, and A. Lasenby. SKYNET: an efficient and robust neural network training tool for machine learning in astronomy. *Mon. Not. Roy. Astron. Soc.*, 441:1741–1759, June 2014. doi: 10.1093/mnras/stu642.
- H. F. Gruetjen and E. P. S. Shellard. Towards efficient and optimal analysis of CMB anisotropies on a masked sky. *ArXiv e-prints*, Dec. 2012.
- J. E. Gunn, M. Carr, C. Rockosi, M. Sekiguchi, K. Berry, B. Elms, E. de Haas, Ž. Ivezić, G. Knapp, R. Lupton, G. Pauls, R. Simcoe, R. Hirsch, D. Sanford, S. Wang, D. York, F. Harris, J. Annis, L. Bartozek, W. Boroski, J. Bakken, M. Haldeman, S. Kent, S. Holm, D. Holmgren, D. Petravick, A. Prosapio, R. Rechenmacher, M. Doi, M. Fukugita, K. Shimasaku, N. Okada, C. Hull, W. Siegmund, E. Mannery, M. Blouke, D. Heidtman, D. Schneider, R. Lucinio, and J. Brinkman. The Sloan Digital Sky Survey Photometric Camera. *Astrophys. J.*, 116:3040–3081, Dec. 1998. doi: 10.1086/300645.
- J. E. Gunn, W. A. Siegmund, E. J. Mannery, R. E. Owen, C. L. Hull, R. F. Leger, L. N. Carey, G. R. Knapp, D. G. York, W. N. Boroski, S. M. Kent, R. H. Lupton, C. M. Rockosi, M. L. Evans, P. Waddell, J. E. Anderson, J. Annis, J. C. Barentine, L. M. Bartoszek, S. Bastian, S. B. Bracker, H. J. Brewington, C. I. Briegel, J. Brinkmann, Y. J. Brown, M. A. Carr, P. C. Czarapata, C. C. Drennan, T. Dombeck, G. R. Federwitz, B. A. Gillespie, C. Gonzales, S. U. Hansen, M. Harvanek, J. Hayes, W. Jordan, E. Kinney, M. Klaene, S. J. Kleinman, R. G. Kron, J. Kresinski, G. Lee, S. Limmongkol, C. W. Lindenmeyer, D. C. Long, C. L. Loomis, P. M. McGehee, P. M. Mantsch, E. H. Neilsen, Jr., R. M. Neswold, P. R. Newman, A. Nitta, J. Peoples, Jr., J. R. Pier, P. S. Prieto, A. Prosapio, C. Rivetta, D. P. Schneider, S. Snedden, and S.-i. Wang. The 2.5 m Telescope of the Sloan Digital Sky Survey. Astrophys. J., 131:2332–2359, Apr. 2006. doi: 10.1086/500975.
- A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356, Jan. 1981.
- A. H. Guth and E. J. Weinberg. Cosmological consequences of a first-order phase transition in the SU₅ grand unified model. *Phys. Rev. D*, 23:876–885, Feb. 1981. doi: 10.1103/PhysRevD.23.876.
- A. Hajian, N. Battaglia, D. N. Spergel, J. R. Bond, C. Pfrommer, and J. L. Sievers. Measuring the thermal Sunyaev-Zel'dovich effect through the cross correlation of Planck and WMAP maps with ROSAT galaxy cluster catalogs. *Journal of Cosmology and Astroparticle Physics*, 11:064, Nov. 2013. doi: 10.1088/1475-7516/2013/11/064.
- J. Hamann and J. Hasenkamp. A new life for sterile neutrinos: resolving inconsistencies using hot dark matter. *Journal of Cosmology and Astroparticle Physics*, 10:044, Oct. 2013. doi: 10.1088/1475-7516/2013/10/044.
- N. Hamaus, U. Seljak, and V. Desjacques. Optimal constraints on local primordial non-Gaussianity from the two-point statistics of large-scale structure. *Phys. Rev. D*, 84(8):083509, Oct. 2011. doi: 10.1103/PhysRevD.84.083509.
- A. J. S. Hamilton and M. Tegmark. A scheme to deal accurately and efficiently with complex angular masks in galaxy surveys. *Mon. Not. Roy. Astron. Soc.*, 349:115–128, Mar. 2004. doi: 10.1111/j.1365-2966.2004.07490.x.
- S. Hamimeche and A. Lewis. Properties and use of CMB power spectrum likelihoods. *Physical Review D.*, 79(8):083012, Apr. 2009. doi: 10.1103/PhysRevD.79.083012.
- D. Healy, Jr., D. Rockmore, P. J. Kostelec, and S. S. B. Moore. Ffts for the 2-sphere improvements and variations. *The Journal of Fourier Analysis and Applications*, 9:341–385, 1996.

- G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright. Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *ArXiv e-prints*, Dec. 2012.
- G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *Astrophys. J. Supp.*, 208:19, Oct. 2013. doi: 10.1088/0067-0049/208/2/19.
- E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen. MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets. *Astrophys. J.*, 567:2–17, Mar. 2002. doi: 10.1086/338126.
- R. Hlozek, J. Dunkley, G. Addison, J. W. Appel, J. R. Bond, C. Sofia Carvalho, S. Das, M. J. Devlin, R. Dünner, T. Essinger-Hileman, J. W. Fowler, P. Gallardo, A. Hajian, M. Halpern, M. Hasselfield, M. Hilton, A. D. Hincks, J. P. Hughes, K. D. Irwin, J. Klein, A. Kosowsky, T. A. Marriage, D. Marsden, F. Menanteau, K. Moodley, M. D. Niemack, M. R. Nolta, L. A. Page, L. Parker, B. Partridge, F. Rojas, N. Sehgal, B. Sherwin, J. Sievers, D. N. Spergel, S. T. Staggs, D. S. Swetz, E. R. Switzer, R. Thornton, and E. Wollack. The Atacama Cosmology Telescope: A Measurement of the Primordial Power Spectrum. *Astrophys. J.*, 749:90, Apr. 2012. doi: 10.1088/0004-637X/749/1/90.
- S. Ho, C. Hirata, N. Padmanabhan, U. Seljak, and N. Bahcall. Correlation of CMB with large-scale structure. I. Integrated Sachs-Wolfe tomography and cosmological implications. *Phys. Rev. D*, 78 (4):043519, Aug. 2008. doi: 10.1103/PhysRevD.78.043519.
- S. Ho, A. Cuesta, H.-J. Seo, R. de Putter, A. J. Ross, M. White, N. Padmanabhan, S. Saito, D. J. Schlegel, E. Schlafly, U. Seljak, C. Hernández-Monteagudo, A. G. Sánchez, W. J. Percival, M. Blanton, R. Skibba, D. Schneider, B. Reid, O. Mena, M. Viel, D. J. Eisenstein, F. Prada, B. A. Weaver, N. Bahcall, D. Bizyaev, H. Brewinton, J. Brinkman, L. Nicolaci da Costa, J. R. Gott, E. Malanushenko, V. Malanushenko, B. Nichol, D. Oravetz, K. Pan, N. Palanque-Delabrouille, N. P. Ross, A. Simmons, F. de Simoni, S. Snedden, and C. Yeche. Clustering of Sloan Digital Sky Survey III Photometric Luminous Galaxies: The Measurement, Systematics, and Cosmological Implications. *Astrophys. J.*, 761:14, Dec. 2012. doi: 10.1088/0004-637X/761/1/14.
- S. Ho, N. Agarwal, A. D. Myers, R. Lyons, A. Disbrow, H.-J. Seo, A. Ross, C. Hirata, N. Padmanabhan, R. O'Connell, E. Huff, D. Schlegel, A. Slosar, D. Weinberg, M. Strauss, N. P. Ross, D. P. Schneider, N. Bahcall, J. Brinkmann, N. Palanque-Delabrouille, and C. Yèche. Sloan Digital Sky Survey III Photometric Quasar Clustering: Probing the Initial Conditions of the Universe using the Largest Volume. *ArXiv e-prints*, Nov. 2013.
- J. Hu, R. Cai, Z. Guo, and B. Hu. Cosmological parameter estimation from CMB and X-ray clusters after Planck. *ArXiv e-prints*, 2014.
- D. Huterer, L. Knox, and R. C. Nichol. The Angular Power Spectrum of Edinburgh/Durham Southern Galaxy Catalogue Galaxies. *Astrophys. J.*, 555:547–557, July 2001. doi: 10.1086/323328.
- D. Huterer, C. E. Cunha, and W. Fang. Calibration errors unleashed: effects on cosmological parameters and requirements for large-scale structure surveys. *Mon. Not. Roy. Astron. Soc.*, May 2013. doi: 10.1093/mnras/stt653.
- A. H. Jaffe, J. R. Bond, P. G. Ferreira, and L. E. Knox. CMB likelihood functions for beginners and experts. In L. Maiani, F. Melchiorri, and N. Vittorio, editors, *3K cosmology*, volume 476 of *American Institute of Physics Conference Series*, pages 249–365, May 1999. doi: 10.1063/1.59329.

- Y. Kanjin. On algebras with convolution structures for Laguerre polynomials. *Transactions of the American Mathematical Society*, 295(2):pp. 783–794, 1986. ISSN 00029947.
- D. Karagiannis, T. Shanks, and N. P. Ross. Search for primordial non-Gaussianity in the quasars of SDSS-III BOSS DR9. *ArXiv e-prints*, Oct. 2013.
- R. Keisler, C. L. Reichardt, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzapfel, S. Hoover, Z. Hou, J. D. Hrubes, M. Joy, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, J. J. Mohr, T. E. Montroy, T. Natoli, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, L. Shaw, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, K. Story, A. van Engelen, K. Vanderlinde, J. D. Vieira, R. Williamson, and O. Zahn. A Measurement of the Damping Tail of the Cosmic Microwave Background Power Spectrum with the South Pole Telescope. *Astrophys. J.*, 743:28, Dec. 2011. doi: 10.1088/0004-637X/743/1/28.
- M. Kilbinger, L. Fu, C. Heymans, F. Simpson, J. Benjamin, et al. CFHTLenS: Combined probe cosmological model comparison using 2D weak gravitational lensing. *Mon. Not. Roy. Astron. Soc.*, 430(3):2200–2220, 2013. doi: 10.1093/mnras/stt041.
- L. Knox, J. R. Bond, and A. H. Jaffe. CMB Power Spectrum Estimation. In A. V. Olinto, J. A. Frieman, & D. N. Schramm, editor, *Eighteenth Texas Symposium on Relativistic Astrophysics*, page 282, 1998.
- E. Komatsu and D. N. Spergel. Acoustic signatures in the primary microwave background bispectrum. *Phys. Rev. D*, 63(6):063002, Mar. 2001. doi: 10.1103/PhysRevD.63.063002.
- E. Komatsu, N. Afshordi, N. Bartolo, D. Baumann, J. R. Bond, E. I. Buchbinder, C. T. Byrnes, X. Chen, D. J. H. Chung, A. Cooray, P. Creminelli, N. Dalal, O. Dore, R. Easther, A. V. Frolov, J. Khoury, W. H. Kinney, L. Kofman, K. Koyama, L. Leblond, J.-L. Lehners, J. E. Lidsey, M. Liguori, E. A. Lim, A. Linde, D. H. Lyth, J. Maldacena, S. Matarrese, L. McAllister, P. McDonald, S. Mukohyama, B. Ovrut, H. V. Peiris, A. Riotto, Y. Rodrigues, M. Sasaki, R. Scoccimarro, D. Seery, A. Sefusatti, K. M. Smith, A. A. Starobinsky, P. J. Steinhardt, F. Takahashi, M. Tegmark, A. J. Tolley, L. Verde, B. D. Wandelt, D. Wands, S. Weinberg, M. Wyman, A. P. S. Yadav, and M. Zaldarriaga. Non-Gaussianity as a Probe of the Physics of the Primordial Universe and the Astrophysics of the Low Redshift Universe. In *astro2010: The Astronomy and Astrophysics Decadal Survey*, volume 2010 of *Astronomy*, page 158, 2009.
- K. Kuijken. Shears from shapelets. Astron. & Astrophys., 456:827–838, Sept. 2006. doi: 10.1051/ 0004-6361:20054794.
- A. Labatie, J. L. Starck, and M. Lachièze-Rey. Detecting Baryon Acoustic Oscillations. *Astron. J.*, 746:172, Feb. 2012. doi: 10.1088/0004-637X/746/2/172.
- X. Lan and D. Marinucci. The needlets bispectrum. *Electronic Journal of Statistics*, 2:332–367, 2008. doi: 10.1214/08-EJS197.
- F. Lanusse, A. Rassat, and J.-L. Starck. Spherical 3D isotropic wavelets. *Astron. & Astrophys.*, 540: A92, Apr. 2012. doi: 10.1051/0004-6361/201118568.
- B. Leistedt and J. D. McEwen. Exact wavelets on the ball. IEEE Trans. Sig. Proc., 60(12), 2012.
- B. Leistedt and H. V. Peiris. Exploiting the full potential of photometric quasar surveys: Optimal power spectra through blind mitigation of systematics. *ArXiv e-prints*, 2014.
- B. Leistedt, A. Rassat, A. Réfrégier, and J.-L. Starck. 3DEX: a code for fast spherical Fourier-Bessel decomposition of 3D surveys. *Astron. & Astrophys.*, 540:A60, Apr. 2012. doi: 10.1051/0004-6361/ 201118463.

- B. Leistedt, H. V. Peiris, D. J. Mortlock, A. Benoit-Lévy, and A. Pontzen. Estimating the large-scale angular power spectrum in the presence of systematics: a case study of Sloan Digital Sky Survey quasars. *Mon. Not. Roy. Astron. Soc.*, 435:1857–1873, Nov. 2013. doi: 10.1093/mnras/stt1359.
- B. Leistedt, H. V. Peiris, and N. Roth. Constraints on primordial non-Gaussianity from 800,000 photometric quasars. *ArXiv e-prints*, 2014.
- D. Lemoine. The discrete Bessel transform algorithm. *The Journal of Chemical Physics*, 101(5): 3936–3944, 1994. doi: 10.1063/1.468428.
- J. Lesgourgues and S. Pastor. Neutrino mass from Cosmology. ArXiv e-prints, Dec. 2012.
- C. Lessig. Orthogonal and symmetric haar wavelets on the sphere. Master's thesis, University of Toronto, Canada, 2007.
- C. Lessig and E. Fiume. Soho: Orthogonal and symmetric haar wavelets on the sphere. *ACM Trans. Graph.*, 27(1):4:1–4:11, Mar. 2008. ISSN 0730-0301. doi: 10.1145/1330511.1330515.
- A. Lewis and S. Bridle. Cosmological parameters from CMB and other data: a Monte- Carlo approach. *Phys. Rev.*, D66:103511, 2002.
- A. Lidz, E. J. Baxter, P. Adshead, and S. Dodelson. Primordial non-Gaussianity and reionization. *Phys. Rev. D*, 88(2):023534, July 2013. doi: 10.1103/PhysRevD.88.023534.
- A. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6):389 – 393, 1982. ISSN 0370-2693. doi: 10.1016/0370-2693(82)91219-9. URL http://www.sciencedirect.com/ science/article/pii/0370269382912199.
- I. Loris, F. J. Simons, I. Daubechies, G. Nolet, M. Fornasier, P. Vetter, S. Judd, S. Voronin, C. Vonesch, and J. Charléty. A new approach to global seismic tomography based on regularization by sparsity in a novel 3D spherical wavelet basis. In *EGU General Assembly Conference Abstracts*, volume 12 of *EGU General Assembly Conference Abstracts*, page 6033, May 2010.
- M. LoVerde, A. Miller, S. Shandera, and L. Verde. Effects of scale-dependent non-Gaussianity on cosmological structures. *Astrophys. J. Lett.*, 4:14, Apr. 2008. doi: 10.1088/1475-7516/2008/04/ 014.
- LSST Dark Energy Science Collaboration. Large Synoptic Survey Telescope: Dark Energy Science Collaboration. *ArXiv e-prints*, Nov. 2012.
- J. Maldacena. Non-gaussian features of primordial fluctuations in single field inflationary models. *Journal of High Energy Physics*, 5:013, May 2003. doi: 10.1088/1126-6708/2003/05/013.
- S. Mallat. A Wavelet Tour of Signal Processing. Wavelet Tour of Signal Processing. Elsevier Science, 1999. ISBN 9780080520834. URL http://books.google.co.uk/books?id=hbV0fWQNtB8C.
- M. Manera, R. Scoccimarro, W. J. Percival, L. Samushia, C. K. McBride, A. J. Ross, R. K. Sheth, M. White, B. A. Reid, A. G. Sánchez, R. de Putter, X. Xu, A. A. Berlind, J. Brinkmann, C. Maraston, B. Nichol, F. Montesano, N. Padmanabhan, R. A. Skibba, R. Tojeiro, and B. A. Weaver. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: a large sample of mock galaxy catalogues. *Mon. Not. Roy. Astron. Soc.*, 428:1036–1054, Jan. 2013. doi: 10.1093/mnras/sts084.
- Y. Mao, A. D'Aloisio, J. Zhang, and P. R. Shapiro. Primordial non-Gaussianity estimation using 21 cm tomography from the epoch of reionization. *Phys. Rev. D*, 88(8):081303, Oct. 2013. doi: 10.1103/PhysRevD.88.081303.
- L. Marian, S. Hilbert, R. E. Smith, P. Schneider, and V. Desjacques. Measuring Primordial Nongaussianity Through Weak-lensing Peak Counts. *Astrophys. J. Lett.*, 728:L13, Feb. 2011. doi: 10.1088/2041-8205/728/1/L13.

- D. Marinucci, D. Pietrobon, A. Balbi, P. Baldi, P. Cabella, G. Kerkyacharian, P. Natoli, D. Picard, and N. Vittorio. Spherical needlets for cosmic microwave background data analysis. *Mon. Not. Roy. Astron. Soc.*, 383:539–545, Jan. 2008. doi: 10.1111/j.1365-2966.2007.12550.x.
- C. Markett. A new proof of Watson's product formula for Laguerre polynomials via a Cauchy problem associated with a singular differential operator. *SIAM Journal on Mathematical Analysis*, 17(4):1010–1032, 1986. doi: 10.1137/0517072.
- J. Martin. Everything you always wanted to know about the cosmological constant problem (but were afraid to ask). *Comptes Rendus Physique*, 13:566–665, July 2012. doi: 10.1016/j.crhy.2012. 04.008.
- R. Massey, J. Rhodes, R. Ellis, N. Scoville, A. Leauthaud, A. Finoguenov, P. Capak, D. Bacon, H. Aussel, J.-P. Kneib, A. Koekemoer, H. McCracken, B. Mobasher, S. Pires, A. Refregier, S. Sasaki, J.-L. Starck, Y. Taniguchi, A. Taylor, and J. Taylor. Dark matter maps reveal cosmic scaffolding. *Nature*, 445:286–290, Jan. 2007. doi: 10.1038/nature05497.
- S. Matarrese and L. Verde. The Effect of Primordial Non-Gaussianity on Halo Bias. *Astrophys. J. Lett.*, 677:L77–L80, Apr. 2008. doi: 10.1086/587840.
- D. J. Matthews and J. A. Newman. Reconstructing Redshift Distributions with Cross-Correlations: Tests and an Optimized Recipe. *Astrophys.J.*, 721:456–468, 2010. doi: 10.1088/0004-637X/721/1/456.
- J. D. McEwen and Y. Wiaux. A novel sampling theorem on the sphere. *IEEE Trans. Sig. Proc.*, 59 (12):5876–5887, 2011. doi: 10.1109/TSP.2011.2166394.
- J. D. McEwen, M. P. Hobson, and A. N. Lasenby. A directional continuous wavelet transform on the sphere. *Arxiv preprint astro-ph/0609159*, Sept. 2006.
- J. D. McEwen, M. P. Hobson, A. N. Lasenby, and D. J. Mortlock. A high-significance detection of non-Gaussianity in the WMAP 3-year data using directional spherical wavelets. *Mon. Not. Roy. Astron. Soc.*, 371(1):L50–L54, 2006. doi: 10.1111/j.1745-3933.2006.00206.x.
- J. D. McEwen, P. Vielva, M. P. Hobson, E. Martínez-González, and A. N. Lasenby. Detection of the ISW effect and corresponding dark energy constraints made with directional spherical wavelets. *Mon. Not. Roy. Astron. Soc.*, 376(3):1211–1226, 2007a. doi: 10.1111/j.1365-2966.2007.11505.x.
- J. D. McEwen, P. Vielva, Y. Wiaux, R. B. Barreiro, L. Cayón, M. P. Hobson, A. N. Lasenby, E. Martínez-González, and J. L. Sanz. Cosmological applications of a wavelet analysis on the sphere. J. Fourier Anal. and Appl., 13(4):495–510, 2007b. doi: 10.1007/s00041-006-6918-8.
- J. D. McEwen, Y. Wiaux, M. P. Hobson, P. Vandergheynst, and A. N. Lasenby. Probing dark energy with steerable wavelets through correlation of WMAP and NVSS local morphological measures. *Mon. Not. Roy. Astron. Soc.*, 384(4):1289–1300, 2008. doi: 10.1111/j.1365-2966.2007.12776.x.
- J. D. McEwen, P. Vandergheynst, and Y. Wiaux. On the computation of directional scale-discretized wavelet transforms on the sphere. In *Wavelets and Sparsity XIV, SPIE international symposium on optics and photonics, invited contribution*, 2013.
- M. McQuinn and M. White. On using angular cross-correlations to determine source redshift distributions. *Mon. Not. Roy. Astron. Soc.*, 433:2857, 2013. doi: 10.1093/mnras/stt914.
- V. Michel. Wavelets on the 3 dimensional ball. PAMM, 5(1):775–776, 2005. ISSN 1617-7061. doi: 10.1002/pamm.200510362.
- A. D. Myers, R. J. Brunner, G. T. Richards, R. C. Nichol, D. P. Schneider, D. E. Vanden Berk, R. Scranton, A. G. Gray, and J. Brinkmann. First Measurement of the Clustering Evolution of Photometrically Classified Quasars. *Astrophys. J.*, 638:622–634, Feb. 2006. doi: 10.1086/499093.

- A. D. Myers, R. J. Brunner, R. C. Nichol, G. T. Richards, D. P. Schneider, and N. A. Bahcall. Clustering Analyses of 300,000 Photometrically Classified Quasars. I. Luminosity and Redshift Evolution in Quasar Bias. *Astrophys. J.*, 658:85–98, Mar. 2007a. doi: 10.1086/511519.
- A. D. Myers, R. J. Brunner, G. T. Richards, R. C. Nichol, D. P. Schneider, and N. A. Bahcall. Clustering Analyses of 300,000 Photometrically Classified Quasars. II. The Excess on Very Small Scales. *Astrophys. J.*, 658:99–106, Mar. 2007b. doi: 10.1086/511520.
- F. J. Narcowich, P. Petrushev, and J. D. Ward. Localized tight frames on spheres. SIAM J. Math. Anal., 38(2):574–594, 2006.
- J. H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. *Bulletin of the Astronomical Institutes of the Netherlands*, 6:249, Aug. 1932.
- P. J. Outram, F. Hoyle, T. Shanks, S. M. Croom, B. J. Boyle, L. Miller, R. J. Smith, and A. D. Myers. The 2dF QSO Redshift Survey - XI. The QSO power spectrum. *Mon. Not. Roy. Astron. Soc.*, 342: 483–495, June 2003. doi: 10.1046/j.1365-8711.2003.06527.x.
- N. Padmanabhan, C. M. Hirata, U. Seljak, D. J. Schlegel, J. Brinkmann, and D. P. Schneider. Correlating the CMB with luminous red galaxies: The integrated Sachs-Wolfe effect. *Physical Review D.*, 72(4):043525, Aug. 2005. doi: 10.1103/PhysRevD.72.043525.
- N. Padmanabhan, X. Xu, D. J. Eisenstein, R. Scalzo, A. J. Cuesta, et al. A 2Oscillations I : Methods and Application to the Sloan Digital Sky Survey. *Mon. Not. Roy. Astron. Soc.*, 427:2132–2145, Dec. 2012. doi: 10.1111/j.1365-2966.2012.21888.x.
- I. Pâris, P. Petitjean, É. Aubourg, S. Bailey, N. P. Ross, A. D. Myers, M. A. Strauss, S. F. Anderson, E. Arnau, J. Bautista, D. Bizyaev, A. S. Bolton, J. Bovy, W. N. Brandt, H. Brewington, J. R. Browstein, N. Busca, D. Capellupo, W. Carithers, R. A. C. Croft, K. Dawson, T. Delubac, G. Ebelke, D. J. Eisenstein, P. Engelke, X. Fan, N. Filiz Ak, H. Finley, A. Font-Ribera, J. Ge, R. R. Gibson, P. B. Hall, F. Hamann, J. F. Hennawi, S. Ho, D. W. Hogg, Ž. Ivezić, L. Jiang, A. E. Kimball, D. Kirkby, J. A. Kirkpatrick, K.-G. Lee, J.-M. Le Goff, B. Lundgren, C. L. MacLeod, E. Malanushenko, V. Malanushenko, C. Maraston, I. D. McGreer, R. G. McMahon, J. Miralda-Escudé, D. Muna, P. Noterdaeme, D. Oravetz, N. Palanque-Delabrouille, K. Pan, I. Perez-Fournon, M. M. Pieri, G. T. Richards, E. Rollinde, E. S. Sheldon, D. J. Schlegel, D. P. Schneider, A. Slosar, A. Shelden, Y. Shen, A. Simmons, S. Snedden, N. Suzuki, J. Tinker, M. Viel, B. A. Weaver, D. H. Weinberg, M. White, W. M. Wood-Vasey, and C. Yèche. The Sloan Digital Sky Survey quasar catalog: ninth data release. *Astron. & Astrophys.*, 548:A66, Dec. 2012. doi: 10.1051/0004-6361/201220142.
- C.-G. Park, J.-c. Hwang, and H. Noh. Axion as a cold dark matter candidate: Low-mass case. *Phys. Rev. D*, 86(8):083535, Oct. 2012. doi: 10.1103/PhysRevD.86.083535.
- D. Parkinson, S. Riemer-Sorensen, C. Blake, G. B. Poole, T. M. Davis, et al. The WiggleZ Dark Energy Survey: Final data release and cosmological results. *Phys.Rev.*, D86:103518, 2012. doi: 10.1103/PhysRevD.86.103518.
- J. E. G. Peek and G. J. Graves. A Correction to the Standard Galactic Reddening Map: Passive Galaxies as Standard Crayons. *Astrophys. J.*, 719:415–424, Aug. 2010. doi: 10.1088/0004-637X/719/1/415.
- A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophys. J.*, 142:419–421, July 1965. doi: 10.1086/148307.
- S. Perlmutter, G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and T. S. C. Project. Measurements of Ω and Λ from 42 High-Redshift Supernovae. *Astrophys. J.*, 517:565–586, June 1999. doi: 10.1086/307221.

- D. Pietrobon, A. Amblard, A. Balbi, P. Cabella, A. Cooray, and D. Marinucci. Needlet detection of features in the WMAP CMB sky and the impact on anisotropies and hemispherical asymmetries. *Physical Review D.*, 78(10):103504, Nov. 2008. doi: 10.1103/PhysRevD.78.103504.
- D. Pietrobon, A. Balbi, P. Cabella, and K. M. Gorski. NeedATool: A Needlet Analysis Tool for Cosmological Data Processing. *Astrophys. J.*, 723:1–9, Nov. 2010. doi: 10.1088/0004-637X/723/ 1/1.
- A. Pillepich, C. Porciani, and T. H. Reiprich. The X-ray cluster survey with eRosita: forecasts for cosmology, cluster physics and primordial non-Gaussianity. *Mon. Not. Roy. Astron. Soc.*, 422: 44–69, May 2012. doi: 10.1111/j.1365-2966.2012.20443.x.
- Planck Collaboration. Planck 2013 results. XXIII. Isotropy and statistics of the CMB. *ArXiv e-prints*, Mar. 2013a.
- Planck Collaboration. Planck 2013 results. XXV. Searches for cosmic strings and other topological defects. *ArXiv e-prints*, Mar. 2013b.
- Planck Collaboration. Planck 2013 results. XXVI. Background geometry and topology of the Universe. *ArXiv e-prints*, Mar. 2013c.
- Planck Collaboration. Planck 2013 results. XVII. Gravitational lensing by large-scale structure. *ArXiv e-prints*, Mar. 2013d.
- Planck Collaboration. Planck intermediate results. XVII. Emission of dust in the diffuse interstellar medium from the far-infrared to microwave frequencies. *ArXiv e-prints*, Dec. 2013e.
- Planck Collaboration. Planck 2013 results. XII. Component separation. ArXiv e-prints, Mar. 2013f.
- Planck Collaboration. Planck 2013 results. XVI. Cosmological parameters. ArXiv e-prints, Mar. 2013g.
- Planck Collaboration. Planck 2013 Results. XXIV. Constraints on primordial non-Gaussianity. *ArXiv e-prints*, Mar. 2013h.
- Planck Collaboration. Planck 2013 results. XV. CMB power spectra and likelihood. *ArXiv e-prints*, Mar. 2013i.
- Planck Collaboration. Planck 2013 results. XX. Cosmology from Sunyaev-Zeldovich cluster counts. *ArXiv e-prints*, Mar. 2013j.
- Planck Collaboration. Planck 2013 results. XVI. Cosmological parameters. ArXiv e-prints, Mar. 2013k.
- Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, and et al. Planck 2013 results. XXII. Constraints on inflation. *ArXiv e-prints*, Mar. 2013.
- H. Pollard. Representation of an analytic function by a Laguerre series. *The Annals of Mathematics*, 48(2):pp. 358–365, 1947. ISSN 0003486X.
- A. Pontzen and H. V. Peiris. The cut-sky cosmic microwave background is not anomalous. *Physical Review D.*, 81(10):103008, May 2010. doi: 10.1103/PhysRevD.81.103008.
- C. Porciani and P. Norberg. Luminosity- and redshift-dependent quasar clustering. *Mon. Not. Roy. Astron. Soc.*, 371:1824–1834, Oct. 2006. doi: 10.1111/j.1365-2966.2006.10813.x.
- A. R. Pullen and C. M. Hirata. Systematic Effects in Large-Scale Angular Power Spectra of Photometric Quasars and Implications for Constraining Primordial Non-Gaussianity. *Publications of the Astronomical Society of the Pacific*, 125:705–718, June 2013. doi: 10.1086/671189.

- A. Rassat and A. Refregier. 3D spherical analysis of baryon acoustic oscillations. Astron. & Astrophys., 540:A115, Apr. 2012. doi: 10.1051/0004-6361/201118638.
- A. Refregier and D. Bacon. Shapelets II. A method for weak lensing measurements. *Mon. Not. Roy. Astron. Soc.*, 338:48–56, Jan. 2003. doi: 10.1046/j.1365-8711.2003.05902.x.
- D. Regan, M. Gosenca, and D. Seery. Constraining the WMAP9 bispectrum and trispectrum with needlets. 2013.
- C. L. Reichardt, L. Shaw, O. Zahn, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzapfel, S. Hoover, Z. Hou, J. D. Hrubes, M. Joy, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, J. J. Mohr, T. E. Montroy, T. Natoli, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, K. Story, A. van Engelen, K. Vanderlinde, J. D. Vieira, and R. Williamson. A Measurement of Secondary Cosmic Microwave Background Anisotropies with Two Years of South Pole Telescope Observations. *Astrophys. J.*, 755:70, Aug. 2012. doi: 10.1088/0004-637X/755/1/70.
- B. A. Reid, W. J. Percival, D. J. Eisenstein, L. Verde, D. N. Spergel, R. A. Skibba, N. A. Bahcall, T. Budavari, J. A. Frieman, M. Fukugita, J. R. Gott, J. E. Gunn, Ž. Ivezić, G. R. Knapp, R. G. Kron, R. H. Lupton, T. A. McKay, A. Meiksin, R. C. Nichol, A. C. Pope, D. J. Schlegel, D. P. Schneider, C. Stoughton, M. A. Strauss, A. S. Szalay, M. Tegmark, M. S. Vogeley, D. H. Weinberg, D. G. York, and I. Zehavi. Cosmological constraints from the clustering of the Sloan Digital Sky Survey DR7 luminous red galaxies. *Mon. Not. Roy. Astron. Soc.*, 404:60–85, May 2010. doi: 10.1111/j.1365-2966.2010.16276.x.
- M. Reinecke. Libpsht algorithms for efficient spherical harmonic transforms. *Astron. & Astrophys.*, 526:A108, Feb. 2011. doi: 10.1051/0004-6361/201015906.
- M. Reinecke and D. S. Seljebotn. Libsharp spherical harmonic transforms revisited. Astron. & Astrophys., 554:A112, June 2013. doi: 10.1051/0004-6361/201321494.
- G. T. Richards, M. A. Strauss, X. Fan, P. B. Hall, S. Jester, D. P. Schneider, D. E. Vanden Berk, C. Stoughton, S. F. Anderson, R. J. Brunner, J. Gray, J. E. Gunn, Ž. Ivezić, M. K. Kirkland, G. R. Knapp, J. Loveday, A. Meiksin, A. Pope, A. S. Szalay, A. R. Thakar, B. Yanny, D. G. York, J. C. Barentine, H. J. Brewington, J. Brinkmann, M. Fukugita, M. Harvanek, S. M. Kent, S. J. Kleinman, J. Krzesiński, D. C. Long, R. H. Lupton, T. Nash, E. H. Neilsen, Jr., A. Nitta, D. J. Schlegel, and S. A. Snedden. The Sloan Digital Sky Survey Quasar Survey: Quasar Luminosity Function from Data Release 3. *Astrophys. J.*, 131:2766–2787, June 2006. doi: 10.1086/503559.
- G. T. Richards, A. D. Myers, A. G. Gray, R. N. Riegel, R. C. Nichol, R. J. Brunner, A. S. Szalay, D. P. Schneider, and S. F. Anderson. Efficient Photometric Selection of Quasars from the Sloan Digital Sky Survey: II. ~1,000,000 Quasars from Data Release Six. Astrophys. J. Supp., 180(1):67, 2009.
- S. Riemer-Sorensen, D. Parkinson, and T. M. Davis. Combining Planck with Large Scale Structure gives strong neutrino mass constraint. 2013.
- A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *Astrophys. J.*, 116:1009–1038, Sept. 1998. doi: 10.1086/300499.
- A. G. Riess, L.-G. Strolger, J. Tonry, S. Casertano, H. C. Ferguson, B. Mobasher, P. Challis, A. V. Filippenko, S. Jha, W. Li, R. Chornock, R. P. Kirshner, B. Leibundgut, M. Dickinson, M. Livio, M. Giavalisco, C. C. Steidel, T. Benítez, and Z. Tsvetanov. Type Ia Supernova Discoveries at z ¿ 1 from the Hubble Space Telescope: Evidence for Past Deceleration and Constraints on Dark Energy Evolution. *Astrophys. J.*, 607:665–687, June 2004. doi: 10.1086/383612.

- A. G. Riess, L. Macri, S. Casertano, H. Lampeitl, H. C. Ferguson, A. V. Filippenko, S. W. Jha, W. Li, and R. Chornock. A 3% Solution: Determination of the Hubble Constant with the Hubble Space Telescope and Wide Field Camera 3. *Astrophys. J.*, 730:119, Apr. 2011a. doi: 10.1088/0004-637X/730/2/119.
- A. G. Riess, L. Macri, S. Casertano, H. Lampeitl, H. C. Ferguson, A. V. Filippenko, S. W. Jha, W. Li, and R. Chornock. A 3% Solution: Determination of the Hubble Constant with the Hubble Space Telescope and Wide Field Camera 3. *Astrophys. J.*, 730:119, Apr. 2011b. doi: 10.1088/0004-637X/730/2/119.
- T. Risbo. Fourier transform summation of Legendre series and D-functions. *Journal of Geodesy*, 70: 383–396, July 1996. doi: 10.1007/BF01090814.
- J. Ritsema, A. Deuss, H. J. van Heijst, and J. H. Woodhouse. S40RTS: a degree-40 shear-velocity model for the mantle from new Rayleigh wave dispersion, teleseismic traveltime and normalmode splitting function measurements. *Geophysical Journal International*, 184(3):1223–1236, 2011. ISSN 1365-246X. doi: 10.1111/j.1365-246X.2010.04884.x.
- A. J. Ross, S. Ho, A. J. Cuesta, R. Tojeiro, W. J. Percival, D. Wake, K. L. Masters, R. C. Nichol, A. D. Myers, F. de Simoni, H. J. Seo, C. Hernández-Monteagudo, R. Crittenden, M. Blanton, J. Brinkmann, L. A. N. da Costa, H. Guo, E. Kazin, M. A. G. Maia, C. Maraston, N. Padmanabhan, F. Prada, B. Ramos, A. Sanchez, E. F. Schlafly, D. J. Schlegel, D. P. Schneider, R. Skibba, D. Thomas, B. A. Weaver, M. White, and I. Zehavi. Ameliorating systematic uncertainties in the angular clustering of galaxies: a study using the SDSS-III. *Mon. Not. Roy. Astron. Soc.*, 417:1350–1373, Oct. 2011. doi: 10.1111/j.1365-2966.2011.19351.x.
- A. J. Ross, W. J. Percival, A. G. Sánchez, L. Samushia, S. Ho, E. Kazin, M. Manera, B. Reid, M. White, R. Tojeiro, C. K. McBride, X. Xu, D. A. Wake, M. A. Strauss, F. Montesano, M. E. C. Swanson, S. Bailey, A. S. Bolton, A. M. Dorta, D. J. Eisenstein, H. Guo, J.-C. Hamilton, R. C. Nichol, N. Padmanabhan, F. Prada, D. J. Schlegel, M. V. Magaña, I. Zehavi, M. Blanton, D. Bizyaev, H. Brewington, A. J. Cuesta, E. Malanushenko, V. Malanushenko, D. Oravetz, J. Parejko, K. Pan, D. P. Schneider, A. Shelden, A. Simmons, S. Snedden, and G.-b. Zhao. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: analysis of potential systematics. *Mon. Not. Roy. Astron. Soc.*, 424:564–590, July 2012a. doi: 10.1111/j.1365-2966.2012.21235.x.
- A. J. Ross, W. J. Percival, A. Carnero, G.-b. Zhao, M. Manera, A. Raccanelli, E. Aubourg, D. Bizyaev, H. Brewington, J. Brinkmann, J. R. Brownstein, A. J. Cuesta, L. A. N. da Costa, D. J. Eisenstein, G. Ebelke, H. Guo, J.-C. Hamilton, M. V. Magaña, E. Malanushenko, V. Malanushenko, C. Maraston, F. Montesano, R. C. Nichol, D. Oravetz, K. Pan, F. Prada, A. G. Sánchez, L. Samushia, D. J. Schlegel, D. P. Schneider, H.-J. Seo, A. Sheldon, A. Simmons, S. Snedden, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, and I. Zehavi. The clustering of galaxies in the SDSS-III DR9 Baryon Oscillation Spectroscopic Survey: constraints on primordial non-Gaussianity. *Mon. Not. Roy. Astron. Soc.*, 428:1116–1127, Jan. 2013. doi: 10.1093/mnras/sts094.
- N. P. Ross, Y. Shen, M. A. Strauss, D. E. Vanden Berk, A. J. Connolly, G. T. Richards, D. P. Schneider, D. H. Weinberg, P. B. Hall, N. A. Bahcall, and R. J. Brunner. Clustering of Low-redshift (z i = 2.2) Quasars from the Sloan Digital Sky Survey. *Astrophys. J.*, 697:1634–1655, June 2009. doi: 10.1088/0004-637X/697/2/1634.
- N. P. Ross, A. D. Myers, E. S. Sheldon, C. Yèche, M. A. Strauss, J. Bovy, J. A. Kirkpatrick, G. T. Richards, É. Aubourg, M. R. Blanton, W. N. Brandt, W. C. Carithers, R. A. C. Croft, R. da Silva, K. Dawson, D. J. Eisenstein, J. F. Hennawi, S. Ho, D. W. Hogg, K.-G. Lee, B. Lundgren, R. G. McMahon, J. Miralda-Escudé, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, M. M. Pieri, J. Rich, N. A. Roe, D. Schiminovich, D. J. Schlegel, D. P. Schneider, A. Slosar, N. Suzuki, J. L. Tinker, D. H. Weinberg, A. Weyant, M. White, and W. M. Wood-Vasey. The SDSS-III Baryon Oscillation Spectroscopic Survey: Quasar Target Selection for Data Release Nine. *Astrophys. J. Supp.*, 199:3, Mar. 2012b. doi: 10.1088/0067-0049/199/1/3.

- N. P. Ross, A. D. Myers, E. S. Sheldon, C. Yèche, M. A. Strauss, J. Bovy, J. A. Kirkpatrick, G. T. Richards, É. Aubourg, M. R. Blanton, W. N. Brandt, W. C. Carithers, R. A. C. Croft, R. da Silva, K. Dawson, D. J. Eisenstein, J. F. Hennawi, S. Ho, D. W. Hogg, K.-G. Lee, B. Lundgren, R. G. McMahon, J. Miralda-Escudé, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, M. M. Pieri, J. Rich, N. A. Roe, D. Schiminovich, D. J. Schlegel, D. P. Schneider, A. Slosar, N. Suzuki, J. L. Tinker, D. H. Weinberg, A. Weyant, M. White, and W. M. Wood-Vasey. The SDSS-III Baryon Oscillation Spectroscopic Survey: Quasar Target Selection for Data Release Nine. *Astrophys. J. Supp.*, 199:3, Mar. 2012c. doi: 10.1088/0067-0049/199/1/3.
- N. Roth and C. Porciani. Can we really measure f_{NL} from the galaxy power spectrum? *Mon. Not. Roy. Astron. Soc.*, 425:L81–L85, 2012. doi: 10.1111/j.1745-3933.2012.01305.x.
- E. Rozo, E. S. Rykoff, J. G. Bartlett, and A. E. Evrard. Cluster Cosmology at a Crossroads: Neutrino Masses. *ArXiv e-prints*, Feb. 2013.
- E. Rozo, J. G. Bartlett, A. E. Evrard, and E. S. Rykoff. Closing the loop: a self-consistent model of optical, X-ray and Sunyaev-Zel'dovich scaling relations for clusters of Galaxies. *Mon. Not. Roy. Astron. Soc.*, 438:78–96, Feb. 2014. doi: 10.1093/mnras/stt2161.
- V. C. Rubin and W. K. Ford, Jr. A Comparison of Dynamical Models of the Andromeda Nebula and the Galaxy. In W. Becker and G. I. Kontopoulos, editors, *The Spiral Structure of our Galaxy*, volume 38 of *IAU Symposium*, page 61, 1970.
- D. S. Salopek and J. R. Bond. Nonlinear evolution of long-wavelength metric fluctuations in inflationary models. *Phys. Rev. D*, 42:3936–3962, Dec. 1990. doi: 10.1103/PhysRevD.42.3936.
- L. Samushia, B. A. Reid, M. White, W. J. Percival, A. J. Cuesta, L. Lombriser, M. Manera, R. C. Nichol, D. P. Schneider, D. Bizyaev, H. Brewington, E. Malanushenko, V. Malanushenko, D. Oravetz, K. Pan, A. Simmons, A. Shelden, S. Snedden, J. L. Tinker, B. A. Weaver, D. G. York, and G.-B. Zhao. The clustering of galaxies in the SDSS-III DR9 Baryon Oscillation Spectroscopic Survey: testing deviations from Λ and general relativity using anisotropic clustering of galaxies. *Mon. Not. Roy. Astron. Soc.*, 429:1514–1528, Feb. 2013. doi: 10.1093/mnras/sts443.
- A. G. Sánchez, C. G. Scóccola, A. J. Ross, W. Percival, M. Manera, F. Montesano, X. Mazzalay, A. J. Cuesta, D. J. Eisenstein, E. Kazin, C. K. McBride, K. Mehta, A. D. Montero-Dorta, N. Padmanabhan, F. Prada, J. A. Rubiño-Martín, R. Tojeiro, X. Xu, M. V. Magaña, E. Aubourg, N. A. Bahcall, S. Bailey, D. Bizyaev, A. S. Bolton, H. Brewington, J. Brinkmann, J. R. Brownstein, J. R. Gott, J. C. Hamilton, S. Ho, K. Honscheid, A. Labatie, E. Malanushenko, V. Malanushenko, C. Maraston, D. Muna, R. C. Nichol, D. Oravetz, K. Pan, N. P. Ross, N. A. Roe, B. A. Reid, D. J. Schlegel, A. Shelden, D. P. Schneider, A. Simmons, R. Skibba, S. Snedden, D. Thomas, J. Tinker, D. A. Wake, B. A. Weaver, D. H. Weinberg, M. White, I. Zehavi, and G. Zhao. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological implications of the large-scale two-point correlation function. *Mon. Not. Roy. Astron. Soc.*, 425:415–437, Sept. 2012. doi: 10.1111/j.1365-2966.2012.21502.x.
- A. Sandage. Current Problems in the Extragalactic Distance Scale. Astrophys. J., 127:513, May 1958. doi: 10.1086/146483.
- D. J. Schlegel, D. P. Finkbeiner, and M. Davis. Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds. *Astrophys. J.*, 500:525, June 1998. doi: 10.1086/305772.
- J. Schmitt, J. L. Starck, J. M. Casandjian, J. Fadili, and I. Grenier. Poisson denoising on the sphere: application to the Fermi gamma ray space telescope. *Astron. & Astrophys.*, 517:A26, July 2010. doi: 10.1051/0004-6361/200913822.
- D. P. Schneider, G. T. Richards, P. B. Hall, M. A. Strauss, S. F. Anderson, T. A. Boroson, N. P. Ross, Y. Shen, W. N. Brandt, X. Fan, N. Inada, S. Jester, G. R. Knapp, C. M. Krawczyk, A. R.

Thakar, D. E. Vanden Berk, W. Voges, B. Yanny, D. G. York, N. A. Bahcall, D. Bizyaev, M. R. Blanton, H. Brewington, J. Brinkmann, D. Eisenstein, J. A. Frieman, M. Fukugita, J. Gray, J. E. Gunn, P. Hibon, Ž. Ivezić, S. M. Kent, R. G. Kron, M. G. Lee, R. H. Lupton, E. Malanushenko, V. Malanushenko, D. Oravetz, K. Pan, J. R. Pier, T. N. Price, III, D. H. Saxe, D. J. Schlegel, A. Simmons, S. A. Snedden, M. U. SubbaRao, A. S. Szalay, and D. H. Weinberg. The Sloan Digital Sky Survey Quasar Catalog. V. Seventh Data Release. *Astrophys. J.*, 139:2360, June 2010. doi: 10.1088/0004-6256/139/6/2360.

- P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. *Computer Graphics Proceedings (SIGGRAPH 95)*, pages 161–172, 1995.
- R. Scoccimarro. Transients from Initial Conditions: A Perturbative Analysis. *ArXiv Astrophysics e-prints*, Nov. 1997.
- T. Sekiguchi and N. Sugiyama. Optimal constraint on g nl from cmb. *Journal of Cosmology and Astroparticle Physics*, 2013(09):002, 2013.
- W. Serber, N. Bahcall, B. Ménard, and G. Richards. The Small-Scale Environment of Quasars. *Astrophys. J.*, 643:68–74, May 2006. doi: 10.1086/501443.
- Y. Shen, M. A. Strauss, M. Oguri, J. F. Hennawi, X. Fan, G. T. Richards, P. B. Hall, J. E. Gunn, D. P. Schneider, A. S. Szalay, A. R. Thakar, D. E. Vanden Berk, S. F. Anderson, N. A. Bahcall, A. J. Connolly, and G. R. Knapp. Clustering of High-Redshift (z *i* = 2.9) Quasars from the Sloan Digital Sky Survey. *Astrophys. J.*, 133:2222–2241, May 2007. doi: 10.1086/513517.
- Y. Shen, M. A. Strauss, N. P. Ross, P. B. Hall, Y.-T. Lin, G. T. Richards, D. P. Schneider, D. H. Weinberg, A. J. Connolly, X. Fan, J. F. Hennawi, F. Shankar, D. E. Vanden Berk, N. A. Bahcall, and R. J. Brunner. Quasar Clustering from SDSS DR5: Dependences on Physical Properties. *Astrophys. J.*, 697:1656–1673, June 2009. doi: 10.1088/0004-637X/697/2/1656.
- B. D. Sherwin, S. Das, A. Hajian, G. Addison, J. R. Bond, D. Crichton, M. J. Devlin, J. Dunkley, M. B. Gralla, M. Halpern, J. C. Hill, A. D. Hincks, J. P. Hughes, K. Huffenberger, R. Hlozek, A. Kosowsky, T. Louis, T. A. Marriage, D. Marsden, F. Menanteau, K. Moodley, M. D. Niemack, L. A. Page, E. D. Reese, N. Sehgal, J. Sievers, C. Sifón, D. N. Spergel, S. T. Staggs, E. R. Switzer, and E. Wollack. The Atacama Cosmology Telescope: Cross-correlation of cosmic microwave background lensing and quasars. *Physical Review D.*, 86(8):083006, Oct. 2012. doi: 10.1103/PhysRevD.86.083006.
- A. E. Siegman. Hermite-Gaussian functions of complex argument as optical-beam eigenfunctions. J. *Opt. Soc. Am.*, 63(9):1093–1094, Sep 1973. doi: 10.1364/JOSA.63.001093.
- J. L. Sievers, R. A. Hlozek, M. R. Nolta, V. Acquaviva, G. E. Addison, P. A. R. Ade, P. Aguirre, M. Amiri, J. W. Appel, L. F. Barrientos, E. S. Battistelli, N. Battaglia, J. R. Bond, B. Brown, B. Burger, E. Calabrese, J. Chervenak, D. Crichton, S. Das, M. J. Devlin, S. R. Dicker, W. Bertrand Doriese, J. Dunkley, R. Dünner, T. Essinger-Hileman, D. Faber, R. P. Fisher, J. W. Fowler, P. Gallardo, M. S. Gordon, M. B. Gralla, A. Hajian, M. Halpern, M. Hasselfield, C. Hernández-Monteagudo, J. C. Hill, G. C. Hilton, M. Hilton, A. D. Hincks, D. Holtz, K. M. Huffenberger, D. H. Hughes, J. P. Hughes, L. Infante, K. D. Irwin, D. R. Jacobson, B. Johnstone, J. Baptiste Juin, M. Kaul, J. Klein, A. Kosowsky, J. M. Lau, M. Limon, Y.-T. Lin, T. Louis, R. H. Lupton, T. A. Marriage, D. Marsden, K. Martocci, P. Mauskopf, M. McLaren, F. Menanteau, K. Moodley, H. Moseley, C. B. Netterfield, M. D. Niemack, L. A. Page, W. A. Page, L. Parker, B. Partridge, R. Plimpton, H. Quintana, E. D. Reese, B. Reid, F. Rojas, N. Sehgal, B. D. Sherwin, B. L. Schmitt, D. N. Spergel, S. T. Staggs, O. Stryzak, D. S. Swetz, E. R. Switzer, R. Thornton, H. Trac, C. Tucker, M. Uehara, K. Visnjic, R. Warne, G. Wilson, E. Wollack, Y. Zhao, and C. Zunckel. The Atacama Cosmology Telescope: cosmological parameters from three seasons of data. *Journal of Cosmology and Astroparticle Physics*, 10:060, Oct. 2013. doi: 10.1088/1475-7516/2013/10/060.
- F. J. Simons, I. Loris, E. Brevdo, and I. C. Daubechies. Wavelets and wavelet-like transforms on the sphere and their application to geophysical data inversion. In M. Papadakis, D. Van de Ville, and

V. K. Goyal, editors, *Wavelets and Sparsity XIV*, volume 81380, page 81380X. SPIE, SPIE, 2011. doi: 10.1117/12.892285.

- F. J. Simons, I. Loris, G. Nolet, I. C. Daubechies, S. Voronin, J. S. Judd, P. A. Vetter, J. Charléty, and C. Vonesch. Solving or resolving global tomographic models with spherical wavelets, and the scale and sparsity of seismic heterogeneity. *Geophysical Journal International*, 187:969–988, Nov. 2011. doi: 10.1111/j.1365-246X.2011.05190.x.
- A. Slosar, U. Seljak, and A. Makarov. Exact likelihood evaluations and foreground marginalization in low resolution WMAP data. *Physical Review D.*, 69(12):123003, June 2004. doi: 10.1103/ PhysRevD.69.123003.
- A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan. Constraints on local primordial non-Gaussianity from large scale structure. *Journal of Cosmology and Astroparticle Physics*, 8:031, Aug. 2008. doi: 10.1088/1475-7516/2008/08/031.
- J. Smidt, A. Amblard, C. T. Byrnes, A. Cooray, A. Heavens, and D. Munshi. Cmb contraints on primordial non-gaussianity from the bispectrum and trispectrum and a new consistency test of single-field inflation. *Phys. Rev. D*, 81:123007, Jun 2010. doi: 10.1103/PhysRevD.81.123007.
- K. M. Smith, S. Ferraro, and M. LoVerde. Halo clustering and g_{NL} -type primordial non-Gaussianity. *Journal of Cosmology and Astroparticle Physics*, 1203:032, 2012. doi: 10.1088/1475-7516/2012/03/032.
- G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, G. Hinshaw, P. D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S. S. Meyer, S. H. Moseley, T. Murdock, L. Rokke, R. F. Silverberg, L. Tenorio, R. Weiss, and D. T. Wilkinson. Structure in the COBE differential microwave radiometer first-year maps. *Astrophysical Journal Letters*, 396:L1–L5, Sept. 1992. doi: 10.1086/186504.
- D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. *The Astrophysical Journal Supplement Series*, 148: 175–194, Sept. 2003. doi: 10.1086/377226.
- D. N. Spergel, R. Bean, O. Doré, M. R. Nolta, C. L. Bennett, J. Dunkley, G. Hinshaw, N. Jarosik, E. Komatsu, L. Page, H. V. Peiris, L. Verde, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications for Cosmology. *Astrophys. J. Supp.*, 170:377–408, June 2007. doi: 10.1086/513700.
- J.-L. Starck, Y. Moudden, P. Abrial, and M. Nguyen. Wavelets, ridgelets and curvelets on the sphere. *Astron. & Astrophys.*, 446:1191–1204, Feb. 2006a. doi: 10.1051/0004-6361:20053246.
- J.-L. Starck, S. Pires, and A. Réfrégier. Weak lensing mass reconstruction using wavelets. *Astron. & Astrophys.*, 451:1139–1150, June 2006b. doi: 10.1051/0004-6361:20052997.
- J.-L. Starck, Y. Moudden, and J. Bobin. Polarized wavelets and curvelets on the sphere. *Astron. & Astrophys.*, 497:931–943, Apr. 2009. doi: 10.1051/0004-6361/200811343.
- N. E. Strand, R. J. Brunner, and A. D. Myers. AGN Environments in the Sloan Digital Sky Survey.
 I. Dependence on Type, Redshift, and Luminosity. *Astrophys. J.*, 688:180–189, Nov. 2008. doi: 10.1086/592099.
- M. E. C. Swanson, M. Tegmark, A. J. S. Hamilton, and J. C. Hill. Methods for rapidly processing angular masks of next-generation galaxy surveys. *Mon. Not. Roy. Astron. Soc.*, 387:1391–1402, July 2008. doi: 10.1111/j.1365-2966.2008.13296.x.

- W. Sweldens. The lifting scheme: A custom-design construction of biorthogonal wavelets. *Appl. Comput. Harmon. Anal.*, 3(2):186–200, 1996.
- W. Sweldens. The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.*, 29(2):511–546, 1997.
- A. N. Taylor, S. Dye, T. J. Broadhurst, N. Benítez, and E. van Kampen. Gravitational Lens Magnification and the Mass of Abell 1689. Astrophys. J., 501:539–553, July 1998. doi: 10.1086/305827.
- M. Tegmark. How to measure CMB power spectra without losing information. *Phys. Rev. D*, 55: 5895–5907, May 1997. doi: 10.1103/PhysRevD.55.5895.
- M. Tegmark and A. de Oliveira-Costa. How to measure CMB polarization power spectra without losing information. *Phys. Rev. D*, 64(6):063001, Sept. 2001. doi: 10.1103/PhysRevD.64.063001.
- M. Tegmark, A. N. Taylor, and A. F. Heavens. Karhunen-Loeve Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets? *Astrophys. J.*, 480:22, May 1997. doi: 10.1086/303939.
- M. Tegmark, A. J. S. Hamilton, M. A. Strauss, M. S. Vogeley, and A. S. Szalay. Measuring the Galaxy Power Spectrum with Future Redshift Surveys. *Astrophys. J.*, 499:555, May 1998. doi: 10.1086/305663.
- M. Tegmark, S. Dodelson, D. J. Eisenstein, V. Narayanan, R. Scoccimarro, R. Scranton, M. A. Strauss, A. Connolly, J. A. Frieman, J. E. Gunn, L. Hui, B. Jain, D. Johnston, S. Kent, J. Loveday, R. C. Nichol, L. O'Connell, R. K. Sheth, A. Stebbins, A. S. Szalay, I. Szapudi, M. S. Vogeley, I. Zehavi, J. Annis, N. A. Bahcall, J. Brinkmann, I. Csabai, M. Doi, M. Fukugita, G. Hennessy, Ž. Ivezíc, G. R. Knapp, D. Q. Lamb, B. C. Lee, R. H. Lupton, T. A. McKay, P. Kunszt, J. A. Munn, J. Peoples, J. R. Pier, M. Richmond, C. Rockosi, D. Schlegel, C. Stoughton, D. L. Tucker, B. Yanny, and D. G. York. The Angular Power Spectrum of Galaxies from Early Sloan Digital Sky Survey Data. *Astrophys. J.*, 571:191–205, May 2002a. doi: 10.1086/339894.
- M. Tegmark, A. J. S. Hamilton, and Y. Xu. The power spectrum of galaxies in the 2dF 100k redshift survey. *Mon. Not. Roy. Astron. Soc.*, 335:887–908, Oct. 2002b. doi: 10.1046/j.1365-8711.2002. 05622.x.
- R. Teyssier, S. Pires, S. Prunet, D. Aubert, C. Pichon, A. Amara, K. Benabed, S. Colombi, A. Refregier, and J.-L. Starck. Full-sky weak-lensing simulation with 70 billion particles. *Astron. & Astrophys.*, 497:335–341, Apr. 2009. doi: 10.1051/0004-6361/200810657.
- S. A. Thomas, F. B. Abdalla, and O. Lahav. The Angular Power Spectra of Photometric SDSS LRGs. *Mon. Not. Roy. Astron. Soc.*, Nov. 2010.
- S. A. Thomas, F. B. Abdalla, and O. Lahav. Upper Bound of 0.28eV on the Neutrino Masses from the Largest Photometric Redshift Survey. *Phys. Rev. Lett.*, 105:031301, 2010. doi: 10.1103/ PhysRevLett.105.031301.
- S. A. Thomas, F. B. Abdalla, and O. Lahav. Excess Clustering on Large Scales in the MegaZ DR7 Photometric Redshift Survey. *Phys. Rev. Lett.*, 106(24):241301, June 2011. doi: 10.1103/ PhysRevLett.106.241301.
- J. L. Tinker, B. E. Robertson, A. V. Kravtsov, A. Klypin, M. S. Warren, G. Yepes, and S. Gottlöber. The Large-scale Bias of Dark Matter Halos: Numerical Calibration and Model Tests. *Astrophys. J.*, 724:878–886, Dec. 2010. doi: 10.1088/0004-637X/724/2/878.
- S. Trapani and J. Navaza. Calculation of spherical harmonics and Wigner *d* functions by FFT. Applications to fast rotational matching in molecular replacement and implementation into *AMoRe*. *Acta Crystallographica Section A*, 62(4):262–269, 2006. doi: 10.1107/S0108767306017478.

- A. van Engelen, R. Keisler, O. Zahn, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzapfel, S. Hoover, Z. Hou, J. D. Hrubes, M. Joy, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, J. J. Mohr, T. E. Montroy, T. Natoli, S. Padin, T. Plagge, C. Pryke, C. L. Reichardt, J. E. Ruhl, J. T. Sayre, K. K. Schaffer, L. Shaw, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, K. Story, K. Vanderlinde, J. D. Vieira, and R. Williamson. A Measurement of Gravitational Lensing of the Microwave Background Using South Pole Telescope Data. *Astrophys. J.*, 756:142, Sept. 2012. doi: 10.1088/0004-637X/756/2/142.
- D. E. Vanden Berk, G. T. Richards, A. Bauer, M. A. Strauss, D. P. Schneider, T. M. Heckman, D. G. York, P. B. Hall, X. Fan, G. R. Knapp, S. F. Anderson, J. Annis, N. A. Bahcall, M. Bernardi, J. W. Briggs, J. Brinkmann, R. Brunner, S. Burles, L. Carey, F. J. Castander, A. J. Connolly, J. H. Crocker, I. Csabai, M. Doi, D. Finkbeiner, S. Friedman, J. A. Frieman, M. Fukugita, J. E. Gunn, G. S. Hennessy, Ž. Ivezić, S. Kent, P. Z. Kunszt, D. Q. Lamb, R. F. Leger, D. C. Long, J. Loveday, R. H. Lupton, A. Meiksin, A. Merelli, J. A. Munn, H. J. Newberg, M. Newcomb, R. C. Nichol, R. Owen, J. R. Pier, A. Pope, C. M. Rockosi, D. J. Schlegel, W. A. Siegmund, S. Smee, Y. Snir, C. Stoughton, C. Stubbs, M. SubbaRao, A. S. Szalay, G. P. Szokoly, C. Tremonti, A. Uomoto, P. Waddell, B. Yanny, and W. Zheng. Composite Quasar Spectra from the Sloan Digital Sky Survey. *The Astrophysical Journal*, 122:549–564, Aug. 2001. doi: 10.1086/321167.
- L. Verde, L. Wang, A. F. Heavens, and M. Kamionkowski. Large-scale structure, the cosmic microwave background and primordial non-Gaussianity. *Mon. Not. Roy. Astron. Soc.*, 313:141–147, Mar. 2000. doi: 10.1046/j.1365-8711.2000.03191.x.
- L. Verde, S. M. Feeney, D. J. Mortlock, and H. V. Peiris. (Lack of) Cosmological evidence for dark radiation after Planck. *Journal of Cosmology and Astroparticle Physics*, 1309:013, 2013. doi: 10.1088/1475-7516/2013/09/013.
- L. Verde, P. Protopapas, and R. Jimenez. Planck and the local Universe: Quantifying the tension. *Physics of the Dark Universe*, 2:166–175, Sept. 2013. doi: 10.1016/j.dark.2013.09.002.
- P. Vielva, E. Martínez-González, R. B. Barreiro, J. L. Sanz, and L. Cayón. Detection of Non-Gaussianity in the Wilkinson Microwave Anisotropy Probe First-Year Data Using Spherical Wavelets. *Astrophys. J.*, 609:22–34, July 2004. doi: 10.1086/421007.
- P. Vielva, E. Martínez-González, and M. Tucci. Cross-correlation of the cosmic microwave background and radio galaxies in real, harmonic and wavelet spaces: detection of the integrated Sachs-Wolfe effect and dark energy constraints. *Mon. Not. Roy. Astron. Soc.*, 365:891–901, Jan. 2006a. doi: 10.1111/j.1365-2966.2005.09764.x.
- P. Vielva, Y. Wiaux, E. Martínez-González, and P. Vandergheynst. Steerable wavelet analysis of CMB structures alignment. *New Astron. Rev.*, 50:880–888, Dec. 2006b. doi: 10.1016/j.newar.2006.09. 007.
- P. Vielva, Y. Wiaux, E. Martínez-González, and P. Vandergheynst. Alignment and signed-intensity anomalies in Wilkinson Microwave Anisotropy Probe data. *Mon. Not. Roy. Astron. Soc.*, 381: 932–942, Nov. 2007. doi: 10.1111/j.1365-2966.2007.12290.x.
- A. Vikhlinin, A. V. Kravtsov, R. A. Burenin, H. Ebeling, W. R. Forman, A. Hornstrup, C. Jones, S. S. Murray, D. Nagai, H. Quintana, and A. Voevodkin. Chandra Cluster Cosmology Project III: Cosmological Parameter Constraints. *Astrophys. J.*, 692:1060–1074, Feb. 2009. doi: 10.1088/ 0004-637X/692/2/1060.
- M. S. Vogeley and A. S. Szalay. Eigenmode Analysis of Galaxy Redshift Surveys. I. Theory and Methods. *Astrophys. J.*, 465:34, July 1996. doi: 10.1086/177399.

- A. von der Linden, A. Mantz, S. W. Allen, D. E. Applegate, P. L. Kelly, R. G. Morris, A. Wright, M. T. Allen, P. R. Burchat, D. L. Burke, D. Donovan, and H. Ebeling. Robust Weak-lensing Mass Calibration of Planck Galaxy Clusters. *ArXiv e-prints*, Feb. 2014.
- C. Wagner and L. Verde. N-body simulations with generic non-Gaussian initial conditions II: halo bias. *Journal of Cosmology and Astroparticle Physics*, 3:002, Mar. 2012. doi: 10.1088/1475-7516/ 2012/03/002.
- B. D. Wandelt, E. Hivon, and K. M. Górski. Cosmic microwave background anisotropy power spectrum statistics for high precision cosmology. *Physical Review D.*, 64(8):083003, Oct. 2001. doi: 10.1103/PhysRevD.64.083003.
- B. D. Wandelt, D. L. Larson, and A. Lakshminarayanan. Global, exact cosmic microwave background data analysis using Gibbs sampling. *Phys.Rev.*, D70:083511, 2004. doi: 10.1103/PhysRevD.70. 083511.
- L. Wang and M. Kamionkowski. Cosmic microwave background bispectrum and inflation. *Phys. Rev. D*, 61(6):063504, Mar. 2000. doi: 10.1103/PhysRevD.61.063504.
- G. Watson. A treatise on the theory of Bessel functions. Cambridge University Press, 1995. ISBN 9780521483919.
- G. Watson. Numerical Recipes: The Art of Scientific Computing, Third Edition. Cambridge University Press, 2007. ISBN 9780521880688.
- E. J. Weniger. On the analyticity of Laguerre series. *Journal of Physics A: Mathematical and Theoretical*, 41(42):425207, 2008.
- M. White, A. D. Myers, N. P. Ross, D. J. Schlegel, J. F. Hennawi, Y. Shen, I. McGreer, M. A. Strauss, A. S. Bolton, J. Bovy, X. Fan, J. Miralda-Escude, N. Palanque-Delabrouille, I. Paris, P. Petitjean, D. P. Schneider, M. Viel, D. H. Weinberg, C. Yeche, I. Zehavi, K. Pan, S. Snedden, D. Bizyaev, H. Brewington, J. Brinkmann, V. Malanushenko, E. Malanushenko, D. Oravetz, A. Simmons, A. Sheldon, and B. A. Weaver. The clustering of intermediate-redshift quasars as measured by the Baryon Oscillation Spectroscopic Survey. *Mon. Not. Roy. Astron. Soc.*, 424:933–950, Aug. 2012. doi: 10.1111/j.1365-2966.2012.21251.x.
- Y. Wiaux, L. Jacques, and P. Vandergheynst. Correspondence principle between spherical and Euclidean wavelets. *Astrophys. J.*, 632:15–28, 2005.
- Y. Wiaux, L. Jacques, P. Vielva, and P. Vandergheynst. Fast directional correlation on the sphere with steerable filters. *Astrophys. J.*, 652:820–832, 2006. doi: 10.1086/507692.
- Y. Wiaux, P. Vielva, E. Martínez-González, and P. Vandergheynst. Global Universe Anisotropy Probed by the Alignment of Structures in the Cosmic Microwave Background. *Physical Review Letters*, 96 (15):151303, Apr. 2006. doi: 10.1103/PhysRevLett.96.151303.
- Y. Wiaux, J. D. McEwen, and P. Vielva. Complex data processing: fast wavelet analysis on the sphere. J. Fourier Anal. and Appl., 13(4):477–493, 2007. doi: 10.1007/s00041-006-6917-9.
- Y. Wiaux, J. D. McEwen, P. Vandergheynst, and O. Blanc. Exact reconstruction with directional wavelets on the sphere. *Mon. Not. Roy. Astron. Soc.*, 388(2):770–788, 2008. doi: 10.1111/j. 1365-2966.2008.13448.x.
- Y. Wiaux, P. Vielva, R. B. Barreiro, E. Martínez-González, and P. Vandergheynst. Non-Gaussianity analysis on local morphological measures of WMAP data. *Mon. Not. Roy. Astron. Soc.*, 385: 939–947, Apr. 2008. doi: 10.1111/j.1365-2966.2008.12901.x.
- M. Wyman, D. H. Rudd, R. A. Vanderveld, and W. Hu. $\lambda u \in \mathbb{R}$ Neutrinos help reconcile Planck with the Local Universe. *ArXiv e-prints*, July 2013.

- J.-Q. Xia, M. Viel, C. Baccigalupi, and S. Matarrese. The high redshift Integrated Sachs-Wolfe effect. *Journal of Cosmology and Astroparticle Physics*, 9:003, Sept. 2009. doi: 10.1088/1475-7516/2009/09/003.
- J.-Q. Xia, A. Bonaldi, C. Baccigalupi, G. De Zotti, S. Matarrese, L. Verde, and M. Viel. Constraining primordial non-Gaussianity with high-redshift probes. *Journal of Cosmology and Astroparticle Physics*, 8:013, Aug. 2010. doi: 10.1088/1475-7516/2010/08/013.
- J.-Q. Xia, C. Baccigalupi, S. Matarrese, L. Verde, and M. Viel. Constraints on primordial non-Gaussianity from large scale structure probes. *Journal of Cosmology and Astroparticle Physics*, 8: 033, Aug. 2011. doi: 10.1088/1475-7516/2011/08/033.
- B. Yeo, W. Ou, and P. Golland. On the construction of invertible filter banks on the 2-sphere. *Image Processing, IEEE Transactions on*, 17(3):283–300, march 2008. ISSN 1057-7149. doi: 10.1109/TIP.2007.915550.
- J.-F. Zhang, Y.-H. Li, and X. Zhang. Sterile neutrinos help reconcile the observational results of primordial gravitational waves from Planck and BICEP2. 2014.
- F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.