A clear view of the primordial Universe

Keir Kwame Rogers

Submitted for the degree of Doctor of Philosophy

Department of Physics and Astronomy

University College London, University of London

October 2017

I, *Keir Kwame Rogers*, 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 temperature anisotropies in the cosmic microwave background (CMB) and measurements of the large-scale structure of matter have established the standard Lambda cold dark matter model of cosmology. Precise measurements of new observables will test extensions to the standard cosmological model, *e. g.*, a non-zero tensor-to-scalar ratio of primordial perturbations, a running of the spectral index of the primordial power spectrum (both tests of cosmic inflation), or new components like massive neutrinos and warm dark matter (WDM). Two of the most promising observables to test these extensions in upcoming surveys are polarisation anisotropies in the CMB and correlations in the Lyman-alpha forest. Accurate cosmological parameter estimation, however, is only achievable through careful consideration of instrumental and astrophysical systematic effects, either by removing contamination in data or modelling its effect during statistical inference. I present new approaches to controlling contaminants to CMB temperature and polarisation and the Lyman-alpha forest.

The primary contamination to the CMB is foreground Galactic radiation, *e. g.*, synchrotron and thermal dust emission. I demonstrate the use of directional wavelets in more accurately reconstructing CMB temperature maps in the presence of these foregrounds, using *Planck* simulations and data. The complexity of polarised Galactic emissions limits constraints on inflation and neutrinos using CMB polarisation. I show how spin directional wavelets can allow additional morphological information to improve cosmic and foreground component separation.

The Lyman-alpha forest probes the primordial power spectrum and the suppression of small-scale clustering by neutrinos or WDM. However, estimation of the shape of the power spectrum is biased by broadened absorption lines formed by high density systems of neutral hydrogen. I present models of their effect, built from Illustris cosmological hydrodynamical simulations. Being functions of

absorber column density provides the flexibility to model residual contamination, after the largest absorbers have been removed from data.

Acknowledgements

I owe a debt of gratitude to my excellent supervisor, Hiranya Peiris, for the faith she has shown in me, sharing her expertise, effort and humour. I thank her for the opportunity to do what I've always wanted to and for always finding time to help me. I must also thank my second supervisor, Andrew Pontzen, for his insight, precision and advice, and for also always being on hand to assist me. In the department, I thank my older "brother," Boris Leistedt, whose committed help kept me going, especially at the beginning of my degree; latterly, I thank Andreu Font-Ribera for explaining the forest and making me pizza. Outside the department, I thank Jason McEwen and Simeon Bird for their instrumental help in our respective projects. I thank all the other friends I have made at UCL and beyond, especially the Italians willing to talk to me about Jeremy Corbyn.

I owe a debt of gratitude to my excellent father, Andrew Mbiu, for single-handedly raising me while sharing with me a desire to understand all things in the Universe. Without his care, guidance, education and love, this endeavour would never have started, let alone be nearing its end.

Contents

List of Figures

List of Tables

1 Introduction

1.1 The cosmological principle

This thesis is a study of the origin, evolution and contents of the Universe, by designing methods to accurately test the principles and models proposed to explain these phenomena with different observables; *i. e.*, a study of physical cosmology. In this Chapter, I will introduce the established knowledge about these principles, models and observations, review some of the tools that are used and indicate the directions of current research, in order to motivate and provide background information to the work that is presented in Chapters 2 to 5.

A guiding notion is the cosmological principle, an assumption that will motivate some of the decisions made below (*e. g.*, see § 1.2). It states that when viewed on sufficiently large scales, the properties of the Universe are the same for all observers. It follows that on these large scales the Universe is *homogeneous*, *i. e.*, that the same observational evidence is available at every location in the Universe. It also follows that the Universe, on these large scales, is *isotropic*, *i. e.*, that the same observational evidence is available by looking in any direction in the Universe. It further follows that the observable Universe (which I will otherwise usually refer to as simply the Universe) is a fair sample of the whole and that the same physical laws apply throughout. There is a body of evidence that supports the cosmological principle (*e. g.*, see § 1.6.1 and 1.6.2), but this remains an active area of research.

1.2 The general theory of relativity

The general theory of relativity (GR) is currently the most complete theory of gravity[1]. It describes gravity as a geometric property of spacetime. It relates the curvature of spacetime to the energy and momentum of matter and radiation (*i. e.*, the contents of the Universe); this relation is expressed by the Einstein field equations (EFE), a set of ten partial differential equations which can be expressed in tensor form:

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R + \Lambda g_{\mu\nu} = 8\pi G T_{\mu\nu}, \tag{1.1}$$

where I have included the cosmological constant Λ, in order to be consistent with the standard cosmological model (see § 1.3). The left-hand side expresses the curvature of spacetime in terms of the metric $g_{\mu\nu}$, which characterises all the geometric and causal structure of spacetime; the Ricci curvature tensor $R_{\mu\nu}$ is a function of the metric and its derivatives and the Ricci scalar R is the contraction of the Ricci tensor with the metric. The right-hand side expresses the energy and momentum of the contents of the Universe through the energy-momentum tensor $T_{\mu\nu}$, which characterises the density and flux of energy and momentum in spacetime (G is Newton's gravitational constant and, throughout this chapter, I am using units where the speed of light in vacuum $c = 1$). μ and ν index over the (four) coordinates of spacetime. Solutions of the EFE consist of spacetime metrics.

The Universe is observed to be expanding; more specifically, the radiation from distant objects (*e. g.*, galaxies) is redshifted (the wavelength is increased) in all directions and the further away they are, the greater is the redshift. This observation and the cosmological principle (see § 1.1) support the use of the Friedmann-Lemaître-Robertson-Walker (FLRW) metric, an (exact) solution of the EFE for a homogeneous, isotropic, expanding (or indeed contracting) universe:

$$g_{\mu\nu} \mathrm{d}x^\mu \mathrm{d}x^\nu = -\mathrm{d}t^2 + a^2(t) \left(\frac{\mathrm{d}r^2}{1 - kr^2} + r^2 (\mathrm{d}\theta^2 + \sin^2 \mathrm{d}\phi^2) \right), \tag{1.2}$$

where the expansion is fully determined by the homogeneous and isotropic scale factor $a(t)$, which is a function of physical time t only and by convention has a value today $a_0 = 1$. The spatial part of the metric is given in reduced-circumference polar coordinates (r, θ, ϕ) for a three-dimensional space of uniform curvature, such that x^μ indexes (t, r, θ, ϕ). k is a constant representing the (global) curvature of space; it is taken to be in the set $\{-1, 0, 1\}$, corresponding to negative, flat or positive curvatures respectively.

[1] There are numerous good textbooks about GR. For a good review of GR in the context of cosmology, see Dodelson (2003); I will here follow his conventions.

The physical distance travelled by a photon $\mathrm{d}t = a\,\mathrm{d}x$, where x is a comoving spatial coordinate (as used in Eq. (1.2)) and the expansion of the Universe has been "scaled-out." It follows that the comoving particle horizon, the maximum comoving distance that a photon could have travelled in a time interval $t - t_\mathrm{i}$ is

$$\chi(t) \equiv \int_{t_\mathrm{i}}^{t} \frac{\mathrm{d}t'}{a(t')} \equiv \eta, \tag{1.3}$$

where we have additionally defined the conformal time η. $\chi(t)$ defines a causal horizon, beyond which particles have not been causally connected since t_i.

The laws of motion in GR are essentially determined by the geodesic equation; by solving the geodesic equation for a photon $g_{\mu\nu}\mathrm{d}x^\mu \mathrm{d}x^\nu = 0$ with the FLRW metric given in Eq. (1.2), it follows that the physical momentum and energy of photons $\propto \frac{1}{a}$. Since the energy of a photon is inversely proportional to its wavelength λ, I can define the (cosmological) redshift due to the expansion of the Universe in terms of the observed (λ_obsv) and emitted (λ_emit) wavelengths and then relate these to the scale factors at the times of observation (a_obsv) and emission (a_emit):

$$1 + z \equiv \frac{\lambda_\mathrm{obsv}}{\lambda_\mathrm{emit}} = \frac{a_\mathrm{obsv}}{a_\mathrm{emit}}. \tag{1.4}$$

Since by convention $a_0 = 1$ (as mentioned above), the redshift observed today of distant objects $= \frac{1}{a} - 1$.

Ignoring any comoving motion ($\dot{x} = 0$, *i. e.*, no "peculiar" velocity), the relative velocity due only to the metric expansion of the Universe $v = \dot{a}x = Hd$ (where $\dot{}$ indicates the derivative with respect to time). Here, in the second equality, I have expressed the Hubble law and related this relative velocity to the physical distance d between objects by defining the Hubble parameter

$$H \equiv \frac{1}{a}\frac{\mathrm{d}a}{\mathrm{d}t}. \tag{1.5}$$

In general, the Hubble parameter is a function of time and so determinations of its value through cosmic history reveal the changing rate of expansion of the Universe. The Hubble parameter at the present epoch is usually denoted the Hubble constant H_0 and is often expressed in terms of the dimensionless parameter h as $H_0 = 100\,h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

I now return to the EFE in Eq. (1.1); the FLRW metric was given as a solution to this equation for the observed (homogeneous and isotropic) expansion of the Universe in Eq. (1.2). The remaining piece is the energy-momentum tensor $T_{\mu\nu}$, for which I will consider an isotropic fluid $\mathrm{diag}(-\rho, \mathcal{P}, \mathcal{P}, \mathcal{P})$, where ρ is its energy density and $\mathcal{P}$ is its pressue, to represent the contents of the

Universe. By plugging in $g_{\mu\nu}$ and $T_{\mu\nu}$ into Eq. (1.1), I can derive the Friedmann equations, which more usefully relate the geometry and the contents of the Universe, by connecting the (background) evolution of the scale factor to the energy-momentum of the Universe's contents:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{k}{a^2} + \frac{\Lambda}{3}; \tag{1.6}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(3\mathcal{P} + \rho\right) + \frac{\Lambda}{3}, \tag{1.7}$$

where $\ddot{}$ indicates the second derivative with respect to time.

The conservation of energy-momentum in GR is expressed as $\nabla_\mu T^\mu_\nu = 0$, where ∇_μ is the covariant derivative. For the metric and energy-momentum tensors specified above, this simplifies to

$$\dot{\rho} + 3\frac{\dot{a}}{a}\left(\rho + \mathcal{P}\right) = 0. \tag{1.8}$$

The isotropic fluid considered above is additionally "perfect" (*i. e.*, completely characterised by its rest-frame energy density and pressure) and so has an equation of state $\mathcal{P} = w\rho$. By plugging this into Eq. (1.8), we find the solution $\rho \propto a^{-3(1+w)}$, describing the evolution of the energy contents of the Universe as it expands. Further, by combining this solution with the Friedmann equations (Eqs. (1.6) and (1.7)), we derive equations for the evolution of the scale factor when a given fluid (as specified by an equation of state parameter w) dominates the total energy density of the Universe:

$$a(t) = \begin{cases} t^{\frac{2}{3(1+w)}}, & \text{if } w \neq -1, \\ e^{Ht}, & \text{if } w = -1. \end{cases} \tag{1.9}$$

I will discuss the different components of the Universe, their equations of state and their consequent effects on the evolution of the Universe in § 1.3.1.

I finally comment that effective energy densities can be associated with the last two terms in Eq. (1.6); the cosmological constant Λ implies a constant energy density (*i. e.*, $w = -1$). An explicit curvature term can be formed by manipulating Eq. (1.6) further by defining a critical density $\rho_{\rm crit} = 3H^2/8\pi G$:

$$\Omega(a) - 1 = \frac{k^2}{H^2 a^2}, \tag{1.10}$$

where the total energy density is expressed in units of the critical density $\Omega(a) = \rho/\rho_{\rm crit}$. The Universe is observed to be consistent with a flat global curvature ($k = 0$; see § 1.3 and 1.6.1 for more details); *i. e.*, the total energy density is very close to the critical density. However, this suggests the

"flatness problem" (which I will return to in § 1.4 to motivate the model of cosmic inflation), which is a cosmological fine-tuning problem. If the Universe deviated very slightly from perfect flatness, then by Eq. (1.10), it can be seen that the total energy density would diverge quickly away from the critical density; conversely, considering that the total energy density is observed to be very close to the critical density, then it must have been even closer to the critical density in the past, suggesting that the total energy density of the Universe has been "fine-tuned" to this value. Proposed solutions to this problem include an epoch of primordial cosmic inflation (§ 1.4).

1.3 The standard Lambda cold dark matter model of cosmology

The current concordance cosmology is the Lambda cold dark matter (ΛCDM) model, which is a parameterisation of the Big Bang cosmological model. The Big Bang model explains that the Universe began in a small, dense and high-temperature state and it has since expanded and cooled, symmetries have been broken and lower-energy states persist. This expansion is modelled by the FLRW metric (Eq. (1.2)) and the laws of GR summarised in § 1.2. The Big Bang model was initially proposed to explain the expansion of the Universe observed through the redshifts of distant galaxies (as mentioned in § 1.2; other models were also proposed, most notably the Steady State theory; Bondi and Gold, 1948; Hoyle, 1948), but the key evidence came with the observation of the cosmic microwave background (CMB; Penzias and Wilson, 1965, § 1.6.1; "relic radiation" from the early Universe) which the Big Bang model had predicted.

The baseline ΛCDM parameterisation of the Big Bang model has six free parameters (and a number of fixed parameters) that determine the components (§ 1.3.1), origin (§ 1.3.2) and evolution (§ 1.3.3) of the Universe[2]. The values of these parameters (or more specifically, their probability distributions) are estimated from a number of cosmological observations (§ 1.6), most powerfully the CMB (§ 1.6.1). The Universe is observed to be flat and so the total energy density is fixed ($\Omega = 1$). This energy density is today distributed between three main components (ordinary "baryonic" matter, cold dark matter (CDM) and "dark energy;" these components will be explained in § 1.3.1); their varying proportions are parameterised by Ω_b and Ω_c[3], the energy densities today of baryons and CDM respectively (as a fraction of the total). The initial conditions of the Universe are modelled as purely adiabatic scalar perturbations to the background with a (dimensionless)

[2] I explicitly follow the parameterisation given in Planck Collaboration et al. (2014d, 2016e).
[3] The combinations $\Omega_\mathrm{b}h^2$ and $\Omega_c h^2$ are actually usually sampled (*e. g.*, Planck Collaboration et al., 2014d, 2016e).

curvature power spectrum, parameterised by an amplitude A_s and a power-law scale index (or "tilt") n_s (more details about the primordial power spectrum will be given in § 1.3.2). (The origin of these "initial" conditions is discussed in § 1.4.) The evolution of the Universe from these initial conditions passed through two major epochs: recombination when originally ionised hydrogen was made electrically neutral by cooling; and, later, reionisation when hydrogen ionised again by the formation of the first stars. The remaining two free parameters essentially relate to these phase transitions: $\theta_\star$ is the angular size of the sound horizon[4] at recombination[5]; and τ is the Thomson scattering optical depth due to reionisation. (See § 1.3.3 for more details.) This six-parameter model remains the baseline cosmology, but parameter extensions (as well as more radical departures) are considered (but not preferred by observations); I will discuss some extensions to the model in § 1.4 and 1.5.

1.3.1 Components

There are a number of components to the energy contents of the Universe; here, I will outline the main categories, specify their equations of state and therefore how their energy densities evolve. Furthermore, since these energy densities scale very differently, it follows that their relative contributions to the total energy density vary and the Universe has passed through different epochs, where one component has dominated and determined the expansion rate.

- *Relativistic components*: these are photons and neutrinos (in the early Universe). (Indeed, even matter behaves relativistically in the very early Universe when the kinetic energy is much greater than the mass energy.) These relativistic components have an equation of state $w = \frac{1}{3}$ and so their energy density $\rho_\gamma \propto a^{-4}$. These components dominate the total energy density in the very early Universe and in this radiation-dominated epoch, $a \propto t^{\frac{1}{2}}$. The decoupling of photons from the primordial plasma (*i. e.*, the CMB) will be discussed in § 1.3.3; neutrinos decoupled earlier when the rate of their (weak force) interactions with electrons and positrons fell below the rate of expansion of the Universe (at a temperature $\sim 1\,\mathrm{MeV}$). This cosmic neutrino background is not (yet) directly observed. In the baseline $\Lambda\mathrm{CDM}$ model, the sum of the masses of the different species of neutrino $\sum m_\nu$ is fixed to the minimum value given by oscillation experiments (0.06 eV; Esteban et al., 2017; Forero et al., 2012, 2014; Gonzalez-Garcia et al., 2014). The effective number of (neutrino-like) relativistic degrees of freedom N_eff is fixed to the canonical value 3.046, which assumes three species (corrected

[4]The comoving horizon of a sound wave.
[5]An approximation to this parameter is used in Planck Collaboration et al. (2014d, 2016e).

slightly due to heating by electron-positron annihilation during decoupling). I will discuss in § 1.5 extended cosmological models where these parameters are allowed to vary and the implications for neutrino physics and cosmology.

- *Baryonic matter*: this is the "ordinary" matter that forms part of the Standard Model of particle physics and is "visible" since it interacts by the electromagnetic (EM) force (it also interacts gravitationally and by the weak and strong forces). The fraction of the total energy density today composed of baryons $\Omega_{\rm b}$ is inferred to be 4.9% (Planck Collaboration et al., 2016e). These baryons primarily form hydrogen, helium and other light elements (hydrogen, helium and trace amounts of lithium were formed in the early Universe by Big Bang nucleosynthesis (BBN; see § 1.3.3 for more details)). The mass fraction in helium $Y_{\rm P}$ is fixed in Planck Collaboration et al. (2014d, 2016e) to a standard value from BBN calculations (24.77%; Pisanti et al., 2008).

- *Cold dark matter*: this is non-baryonic, pressureless and non-relativistic matter that does not form part of the Standard Model of particle physics. It does not interact (or at least it does very weakly) by the EM force; it has not been directly or indirectly detected, but its presence is inferred by a number of cosmological and astronomical observations. These include the CMB (Planck Collaboration et al., 2014d, 2016e); galaxy velocity curves (Rubin and Ford, 1970); galaxy velocity dispersions; X-ray emission from galaxy clusters; gravitational lensing of galaxies and galaxy clusters (Taylor et al., 1998); and galaxy clustering (Davis et al., 1982; Oort, 1932; Zwicky, 1933). However, the particle nature of CDM is unknown and remains an active area of research (I will discuss in § 1.5 how cosmological observations can help distinguish between alternative models of dark matter). The leading candidates for CDM can broadly be separated into three categories. Weakly interacting massive particles (WIMPs) are proposed particles that interact by gravity, but by any other forces weakly; candidate particles come, *e. g.*, from supersymmetry theory (see Jungman et al. 1996 for an early, comprehensive review). Low-mass axions are proposed particles that have a non-zero coupling to photons (Peccei and Quinn, 1977a,b). The third category is massive astrophysical compact halo objects (MACHOs), of which there is renewed interest in primordial black holes (Carr and Hawking, 1974) due to some consistency with gravitational wave observations (Bird et al., 2016). The fraction of the total energy density today composed of CDM $\Omega_{\rm c}$ is inferred to be 26.5% (Planck Collaboration et al., 2016e) and so it constitutes 84% of the total matter energy density. The evolution of CDM in cosmic history plays a central role in the formation of large-scale structure

(see § 1.3.3). Since CDM is pressureless, it has an equation of state $w = 0$. It follows that the energy density of non-relativistic matter $\rho_{\rm m} \propto a^{-3}$. It follows that as the Universe continued to expand, the relative proportions of the energy densities of matter to radiation increased, such that the Universe entered a matter-dominated epoch, where $a \propto t^{\frac{2}{3}}$.

- *Dark energy*: this is modelled in ΛCDM cosmology by the cosmological constant Λ (introduced in Eq. (1.1)). It is included in the inferred energy contents of the Universe to explain the observed late-time accelerated expansion of the Universe (Perlmutter et al., 1999; Riess et al., 1998, 2004). This is because its energy density ρ_Λ is constant (as noted in § 1.2; *i. e.*, its equation of state $w = -1$) and so at late times, it becomes the dominant component of the total energy contents of the Universe and in this Λ-dominated epoch, $a \propto e^{Ht}$. Indeed, any equation of state for dark energy $w < -\frac{1}{3}$ would generate late-time accelerated expansion, but current observations are consistent (to percent precision) with the simplest model of the cosmological constant and so in the baseline cosmology w is fixed to -1, with no time-evolution in its equation of state. However, when Λ is considered as vacuum energy, there is a disagreement with theoretical predictions from quantum field theory for the value of Λ of up to 120 orders of magnitude – the so-called "cosmological constant problem." Alternatives to the ΛCDM model aim to avoid this problem by proposing dynamical models of dark energy or alternative theories of gravity; modern reviews of the theoretical challenges can be found in Carroll (2001); Clifton et al. (2012); Martin (2012). The fraction of the total energy density today composed of dark energy Ω_Λ is inferred to be 68% (Planck Collaboration et al., 2016e).

1.3.2 The primordial Universe

Having discussed the different components that constitute the energy contents of the Universe and how they determine the background evolution of the Universe (§ 1.3.1), I will now discuss the initial conditions of the Universe, or more specifically of the perturbations to the background. The majority of the above adheres rigidly to the cosmological principle and concerns an entirely homogeneous and isotropic Universe. However, there are manifestly some small-scale inhomogeneities and anisotropies, *i. e.*, structure has formed in the Universe (this does not contradict the cosmological principle, which only concerns very large scales). In this section, I will present the simple phenomenological parameterisation for the initial perturbations, which went on to grow into the structure that is observed today; however, the origin of these primordial perturbations will only be explored in § 1.4, when I discuss the theory of cosmic inflation. The baseline ΛCDM model

considers only scalar comoving curvature perturbations $\mathcal{R}$, which are purely adiabatic (*i. e.*, have a constant matter-to-radiation ratio everywhere, as opposed to isocurvature perturbations). Tensor perturbations will be discussed in § 1.4 and in the baseline model, the ratio of tensor primordial power to curvature power (or "tensor-to-scalar ratio") r is fixed to zero. These perturbations are Gaussian-distributed and so all the information is characterised by a (dimensionless) primordial power spectrum[6], which is a function of scale k:

$$P_{\mathcal{R}}(k) = A_{\mathrm{s}} \left(\frac{k}{k_0} \right)^{n_{\mathrm{s}}-1} , \tag{1.11}$$

where A_{s} is the amplitude, n_{s} is the scalar spectral index and k_0 is a pivot scale (set to $0.05\,\mathrm{Mpc}^{-1}$ in Planck Collaboration et al. 2016e). The values and 68% credible intervals of A_{s} and n_{s} inferred in Planck Collaboration et al. (2016e) are respectively $(2.198^{+0.076}_{-0.085}) \times 10^{-9}$ and 0.965 ± 0.0062; the latter implies a nearly scale-invariant primordial power spectrum. I will discuss extensions of the standard cosmological model that generalise the form of the primordial power spectrum in § 1.4.

1.3.3 Evolution

The Big Bang cosmological model assumes that the Universe begins in a very hot, small[7] and dense state (itself called the "Big Bang") and that as it expanded, it cooled. This allowed the generation of lower energy states, forming the particles of the Standard Model of particle physics (and probably physics beyond the Standard Model, *e. g.*, CDM). As noted in § 1.3.1, when the temperature cooled to about 1 MeV (or about 1 s after the Big Bang), neutrinos decoupled from the other components, forming the (as yet unobserved) cosmic neutrino background. When the primordial plasma cooled further to about 100 keV (or about 10 s after the Big Bang), the processes of Big Bang nucleosynthesis (BBN) began, where protons and neutrons combined to form the nuclei of the light elements (hydrogen, helium and small amounts of lithium).

From this time until the Universe had cooled to about 1 eV (~ 3000 K, or $\sim 380,000$ years after the Big Bang, or a redshift $z \sim 1100$), there existed a plasma of atomic nuclei and electrons (*i. e.*, baryons) and photons, (to a first approximation) tightly-coupled through Coulomb scattering between nuclei and electrons and Compton scattering between electrons and photons. This baryon-photon plasma interacted gravitationally with the CDM. The mean free path of photons in this

[6] I will not discuss here possible primordial non-Gaussianities, but the search for these remains an active area of research; current observations are consistent with zero non-Gaussianity (Planck Collaboration et al., 2016c).

[7] *I. e.*, the physical distances between any points are small; this says nothing about the comoving volume of the entire Universe beyond what is observable, which may even be infinite.

plasma was very small. Acoustic oscillations were set up in this plasma because of the balance between the tendency of gravity to collapse structure (initially seeded as the primordial perturbations discussed in § 1.3.2) and the pressure of the relativistic photons erasing structure. The scales of these oscillations were related to the size of the sound horizon in the fluid. I will discuss this and other features of the plasma when I discuss the primary observable, the CMB, in § 1.6.1. At the end of this epoch, the Universe had cooled sufficiently to allow neutral hydrogen to form: this event is called recombination[8]. After recombination, the photons decoupled from the baryons and began to free-stream, *i. e.*, their mean free path became very large: it is these photons that we observe today as the CMB. It follows that the CMB is a powerful observable, allowing us to probe the physics of the early Universe. The angular size of the sound horizon at recombination $\theta_\star$ is one of the free parameters in the $\Lambda\mathrm{CDM}$ model, constrained by observations of the CMB. The (nearly spherical) surface at which the photons we observe of the CMB last scattered is called the "last scattering surface." The CMB is also (linearly) polarised because of Compton scattering of the quadrupole of the photon anisotropies at the last scattering surface. There are also late-time effects on the CMB, which I will discuss in § 1.6.1.

I now consider how the initial perturbations introduced in § 1.3.2 evolve up to recombination (and beyond). First, we must make a distinction between super-horizon and sub-horizon modes, *i. e.*, perturbations of a scale respectively larger and smaller than the Hubble radius. As time increases, the Hubble radius increases in size and larger-scale modes cross the horizon and enter into causal contact. Since at any given moment, the largest-scale modes observable have just crossed the horizon, they have evolved in causal contact for the least time. Additionally, the evolution is also dependent on the dominant components in the Universe. As explained in § 1.3.1, the Universe passes through radiation-, matter- and Λ-dominated epochs. The Λ-dominated epoch occurs at late times ($z \lesssim 0.5$), but matter-radiation equality occurs before recombination ($z_{\mathrm{eq}} = 3400$ as inferred by Planck Collaboration et al. 2016c). The evolution of the perturbations is fully determined by solving the (nine first-order differential) Einstein-Boltzmann equations, but I will summarise some of the key dependencies. In particular, in the radiation era, sub-horizon perturbations to the matter distribution only grow logarithmically with time, while super-horizon modes $\sim \eta^2$. However, after matter-radiation equality, the matter perturbations grow linearly with the scale factor. The way in which the spectrum of temperature anisotropies in the CMB (*i. e.*, the angular power spectrum C_ℓ^{TT}) depends on the primordial power spectrum of scalar perturbations (*i. e.*, the evolution of

[8]This is a misnomer, since hydrogen had never previously been neutral; helium recombination is a two-step process that occurred earlier, since its ionisation energy is greater than that of hydrogen.

perturbations noted above) is encapsulated by

$$C_\ell^{\rm TT} = \int_0^\infty \frac{{\rm d}k}{k} (\Delta^{\rm s}_{\ell,{\rm T}})^2 P_{\mathcal{R}}(k), \tag{1.12}$$

where $\Delta^{\rm s}_{\ell,{\rm T}}$ is the scalar temperature CMB transfer function. This transfer function also accounts for late-time effects which I will discuss in § 1.6.1. There may also be a sub-dominant contribution of primordial tensor perturbations to the temperature angular power spectrum, which adds in quadrature. Equivalent equations work for anisotropies in the polarisation of the CMB.

After recombination and photon decoupling, the Universe entered the so-called "dark ages," where most of the baryons resided in neutral hydrogen. The Universe was "dark," apart from 21 cm emission from the spin-flip transition of neutral hydrogen. About 150 million years after the Big Bang, the first stars began to form, themselves beginning to form the first heavy elements. These stars formed the first galaxies, built within halos of dark matter. The ultraviolet background radiation from these sources had sufficient energy to ionise the neutral hydrogen, beginning the epoch of reionisation, which probably proceeded in patches that eventually overlapped. Planck Collaboration et al. (2014d, 2016e) model hydrogen reionisation in their ΛCDM parameterisation as having occurred quite sharply (and simultaneously to the first helium reionisation) with a mid-point redshift $z_{\rm re}$, related to a Thomson scattering optical depth τ, and a width parameter $\Delta z_{\rm re} = 0.5$ [9]. Planck Collaboration et al. (2016g) inferred $\tau = 0.055 \pm 0.009$, giving $z_{\rm re} \sim 8$. The optical depth to reionisation is the most poorly constrained ΛCDM parameter (to only $\sim 15\%$ precision[10] in Planck Collaboration et al. 2016g) and this remains one of the most poorly understood epochs of the ΛCDM model and an active area of research. Once the Universe was fully ionised, the largest structures began to form, with the first proto-clusters forming at $z \sim 6$. After this, larger and larger clusters and structures built up forming the "cosmic web" of clusters, filaments, walls and voids we observe locally today ("hierarchical structure formation;" White and Rees, 1978). The final chapter to the evolution of the Universe is the Λ-dominated epoch, which began at $z \sim 0.5$, and an accelerated phase of expansion. In this epoch, gravitational potentials decay inversely with scale factor and the largest structures are frozen.

[9] The second helium reionisation is fixed at $z = 3.5$, consistent with observations of the Lyman-alpha forest (Becker et al., 2011, see § 1.6.3).

[10] 1σ credible interval on the marginalised 1D posterior distribution.

1.4 Cosmic inflation

An extension of the standard $\Lambda\mathrm{CDM}$ cosmological model is the theory of cosmic inflation. Indeed, inflation is required to resolve a number of problems that otherwise arise in the $\Lambda\mathrm{CDM}$ model, such that it is often considered a part of the model itself, but there is not yet any conclusive evidence for it and there remain numerous viable models of inflation. In general, inflation consists of a period of accelerated expansion at the start of the Universe (before the radiation-dominated epoch). I will now discuss the problems in the $\Lambda\mathrm{CDM}$ model and the properties of inflation necessary to resolve them.

- *Horizon problem*: due to observations of the CMB, the Universe is inferred to be very homogeneous on scales larger than the particle horizon at the time of decoupling, *i. e.*, outside causal contact, assuming the standard expansion history of the Universe without inflation (as set out in § 1.3). Inflation resolves this by driving a period of expansion at the beginning of the Universe where the comoving Hubble radius decreased. The expression for the comoving particle horizon in Eq. (1.3) can be re-expressed in terms of the comoving Hubble radius $(aH)^{-1}$: $\chi(a) = \int_0^a \mathrm{d}\ln a'(a'H(a'))^{-1}$. *I. e.*, the comoving Hubble radius is the distance that particles can travel in one expansion time (approximately the time in which the scale factor doubles); or, it is a measure of the scale on which particles can communicate at any moment in time (as opposed to the particle horizon, which is the scale on which particles have ever been able to communicate). By decreasing the comoving Hubble radius, inflation allows particles separated by a given scale to be in causal contact at early times and to then fall out of contact. After the end of inflation, the standard expansion resumes and the comoving Hubble radius begins to increase. The horizon problem is resolved because the scales in the CMB that otherwise appeared to have never been in causal contact were at the start of inflation.

- *Flatness problem*: this was discussed in § 1.2 and is understood by noting that the observed flatness of the Universe $\Omega = 1$ is an unstable point in Eq. (1.10) for an expanding Universe. This is resolved by inflation because a decreasing comoving Hubble radius in Eq. (1.10) drives Ω towards unity at the end of inflation.

- *Origin of structure*: this was noted in § 1.3.2, in that no mechanism was provided to produce the primordial perturbations that later grow into the structure we observe today. As will be expanded below, the primordial inflationary expansion is usually driven in the simplest models by a scalar field. Inflation postulates that quantum fluctuations in this field were expanded

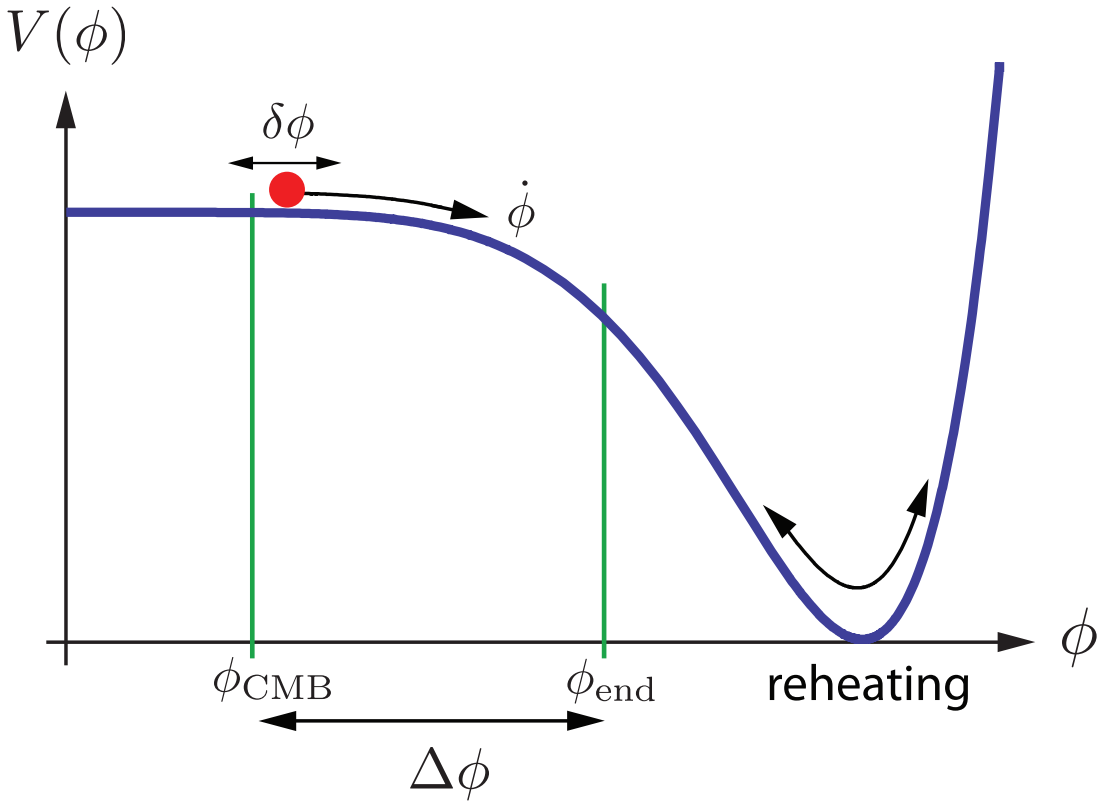

Figure 1.1: An example potential for the inflation field ϕ. Acceleration happens when the potential energy of the field $V(\phi)$ dominates over its kinetic energy $\frac{1}{2}\dot{\phi}^2$. Inflation ends at $\phi_{\rm end}$, when the kinetic energy grows to be comparable to the potential energy. Quantum fluctuations in the field $\delta\phi$ are the source of primordial perturbations and are ultimately the origin of structure in the Universe. At reheating, the energy of the inflaton is converted into radiation. This figure is taken from Baumann (2009).

to macroscopic scales and these are the source of the primordial density perturbations (Starobinsky, 1982).

Above, I motivated why primordial inflationary expansion needs a decreasing comoving Hubble radius. It follows that the expansion is accelerating [$\frac{\rm d}{{\rm d}t}\left(a\frac{\dot{a}}{a}\right) = \ddot{a} > 0$]. I explained in § 1.2 that accelerated expansion occurs for any fluid with an equation of state $w < -\frac{1}{3}$. Another feature of inflation must be that it is finite (at least for our part of the Universe), so that the standard expansion history that we otherwise observe can begin at its end. The fine-tuning problems above are solved if in this time, the scale factor increases by a factor $\gtrsim e^{60}$ (60 "*e*-folds"). At the end of inflation, there must also be a mechanism to couple to the particles of the Standard Model of particle physics ("reheating"). Inflation is usually implemented by a scalar field, the inflaton ϕ. The

initial implementation ("old inflation") drove the accelerated expansion by holding the field in a false vacuum and inflation ended by quantum tunnelling to the true vacuum (Guth, 1981). This model was not viable as it does not allow the Universe to stop inflating and reheat homogeneously (Guth and Weinberg, 1983). The simplest modern implementation ("new inflation") is to have the inflaton roll slowly down a shallow potential (Albrecht and Steinhardt, 1982; Linde, 1982, an example is shown in Fig. 1.1). The Universe will expand at an accelerated rate as long as the potential energy of the field is much greater than its kinetic energy ($V \gg \dot{\phi}^2$), stopping only when the kinetic energy becomes comparable to the potential energy. Once the inflaton reaches the minimum of the potential, it oscillates coherently, coupling to the Standard Model particles and reheating the Universe (Mukhanov and Chibisov, 1981).

There exist numerous viable models of inflation, many more complicated than the simple scheme explained above (see Planck Collaboration et al. 2016d for constraints on many of these different models). A generic prediction of most models of inflation is the generation of Gaussian-distributed primordial perturbations with a nearly scale-invariant spectrum; this has indeed been observed (see § 1.3.2). However, in order to further distinguish between inflationary models, it will be necessary to probe the primordial scalar power spectrum more precisely and to accurately measure new observables like the primordial tensor power spectrum. A standard extension to the baseline ΛCDM parameterisation of the primordial scalar power spectrum is to add an extra scale-dependent term to the index in Eq. (1.11) $\frac{1}{2}\frac{\mathrm{d}n_\mathrm{s}}{\mathrm{d}\ln k}\ln\left(\frac{k}{k_0}\right)$, where $\frac{\mathrm{d}n_\mathrm{s}}{\mathrm{d}\ln k}$ is the "running of the spectral index" (even more complicated forms can be considered, *e. g.*, a "running of the running"). The inflationary epoch would also have generated primordial tensor perturbations (*i. e.*, a stochastic background of primordial gravitational waves). The power spectrum of these modes can be phenomenologically modelled in the same way as the scalar modes (see Eq. (1.11)), with a tensor amplitude A_t and tensor spectral index n_t. The ratio of the tensor and scalar amplitudes is the tensor-to-scalar ratio r. There is a non-zero transfer function for tensor modes for the CMB temperature T and polarisation E and B [11] angular power spectra (although for the T and E modes, the tensor contribution is very sub-dominant to the scalar part). However, there is no transfer of scalar perturbations to the CMB B- mode angular power spectrum and tensor modes are the only primordial source of B modes (Kamionkowski et al., 1997; Seljak and Zaldarriaga, 1997). This makes the measurement of B modes a high priority for future surveys in order to find compelling evidence for inflation and to further distinguish between models (contamination of this signal by lensing and foreground modes will be discussed in § 1.6.1).

[11] The distinction between polarisation E and B modes will be made in § 1.6.1.

1.5 Additional components

Some extensions to the standard cosmological model postulate adding additional components to the energy contents of the Universe. The first I will discuss is massive neutrinos; this is not strictly an additional component – as explained in § 1.3.1, they are included in the baseline model, but varying their parameters (the sum of the three species' masses $\sum m_\nu$ and the number of relativistic degrees of freedom $N_{\rm eff}$) is a non-standard extension. Cosmological observables are sensitive to the neutrino sector partly because of the suppression of small-scale clustering due to massive neutrinos (see § 1.6). Neutrinos are a form of "hot dark matter" and can free-stream out of matter overdensities. This allows *e. g.*, an upper limit to be put on $\sum m_\nu$. Indeed, the value of the best upper limit (0.12 eV[12]; Palanque-Delabrouille et al., 2015) is beginning to approach the lower limits given by neutrino oscillation experiments (Esteban et al., 2017; Forero et al., 2012, 2014; Gonzalez-Garcia et al., 2014), which are different depending on the mass hierarchy of the three neutrino mass eigenstates, *i. e.*, the relative masses of the different species. In the normal hierarchy (implying $\sum m_\nu \gtrsim 0.06\,\mathrm{eV}$), there are two light and one heavy neutrino; whereas in the inverted hierarchy (which has a higher lower limit $\sum m_\nu \gtrsim 0.1\,\mathrm{eV}$), there are two heavy and one light neutrino (Capozzi et al., 2016). The ability to identify the mass hierarchy of the neutrinos will have implications for particle physics, which currently does not have a thorough understanding of why neutrinos are massive. There are two leading mechanisms: the Dirac mechanism (Mohapatra and Senjanović, 1980), requiring a sterile neutrino (also see below); and the Majorana mechanism (Majorana, 1937), where the neutrino is its own antiparticle.

Second, although CDM is part of the standard cosmological model, not all extended models that allow part of the dark matter to be another species are ruled out by current observations. One example of this is warm dark matter (WDM); an example of a candidate particle is the sterile neutrino (Adhikari et al., 2017). WDM would suppress small-scale clustering in a similar way as neutrinos by free-streaming. Cosmological observations of small-scale clustering allow lower limits to be placed on the mass of the WDM particle (see § 1.6; Iršič et al., 2017b). Interest remains in WDM models because they have been suggested to resolve some of the "small-scale crisis" in CDM (*e. g.*, Weinberg et al., 2015): discrepancies between observations and cosmological CDM simulations on small scales (although many of these problems have now been resolved without WDM; *e. g.*, Brooks and Zolotov, 2014).

[12] 95 % confidence limit.

1.6 Cosmological observables

1.6.1 The cosmic microwave background

As discussed in § 1.3.3, the CMB is the background EM radiation, which decoupled shortly after recombination; owing to cosmological redshift, its spectrum today peaks at microwave frequencies. It was first observed by Penzias and Wilson (1965) and identified by Dicke et al. (1965), providing key evidence for the Big Bang cosmological model. The COBE satellite (Boggess et al., 1992; Mather et al., 1990) measured its near-perfect blackbody spectrum (FIRAS instrument; Mather et al., 1994), implying that the CMB was once in thermal equilibrium. It is observed to have a near-uniform temperature of 2.7 K, but COBE also made the first measurement of the small temperature anisotropies of order 1 in 100,000 (DMR instrument; Smoot et al., 1992). The large-scale isotropy of the CMB is key evidence to support the cosmological principle (see § 1.1). Precise measurement of these small anisotropies underpins our determination of cosmological parameters. As explained in § 1.3.2, these anisotropies are sourced by primordial perturbations that themselves are understood to have originated in the early inflationary expansion (§ 1.4). We also learn about the contents and evolution of the Universe as these leave their imprint on the CMB both before and after decoupling. It follows that there have been numerous experiments to make ever more precise measurements; some of the early experiments include *Saskatoon* (Netterfield et al., 1997), BOOMERanG (de Bernardis et al., 2000) and MAXIMA (Hanany et al., 2000). There have been two further generations of satellite experiments: WMAP (Bennett et al., 2003a) and *Planck* (Planck Collaboration et al., 2011); these have been complemented by high-resolution ground-based experiments, ACT (Das et al., 2011a) and SPT (Keisler et al., 2011)[13].

As mentioned in § 1.3.3, the CMB is linearly polarised at the last scattering surface when the quadrupole of the photon anisotropies Compton scatters off the electron distribution. Linear polarisation can be decomposed into Stokes Q and U parameters (relative to some coordinate basis) and these are what is directly measured; angular power spectra can be constructed from their anisotropies in an equivalent way as the temperature anisotropies. However, these do not allow rotationally-invariant measures of angular power spectra and so it is useful to decompose into scalar E and and pseudo-scalar B modes (Kamionkowski et al., 1997; Seljak and Zaldarriaga, 1997, see § 3.4.2 for more details about E-B decomposition). CMB (E-mode) polarisation was first detected by the DASI experiment (Kovac et al., 2002). CMB polarisation anisotropies provide even more statistical power than the temperature modes in constraining cosmological parameters (Galli

[13] See `https://lambda.gsfc.nasa.gov/product/expt` for a comprehensive list of CMB experiments.

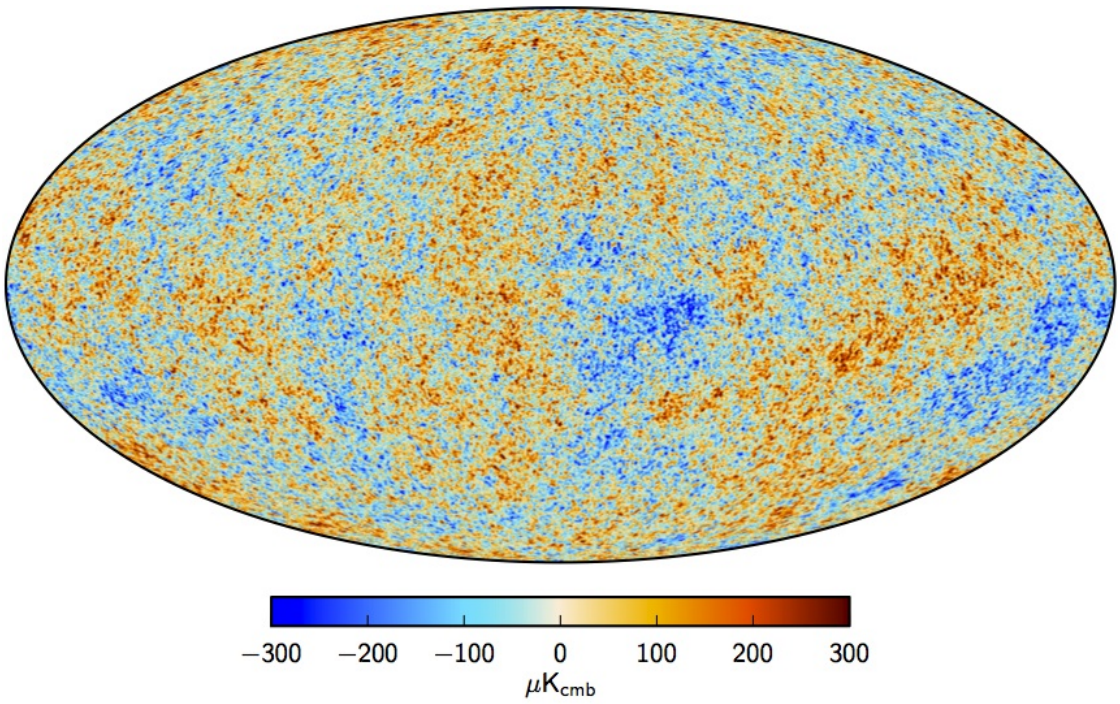

Figure 1.2: A map of the CMB temperature anisotropies as jointly inferred from data from the *Planck* and WMAP satellites and 408 MHz observations (Haslam et al., 1982). This figure is taken from Planck Collaboration et al. (2016a).

et al., 2014) and so there is a continuing effort for precise measurements. As discussed in § 1.4, primordial gravitational waves are a potential primordial source of CMB B modes and detection would provide key evidence to support the theory of inflation (this conclusion is complicated by cosmic defect models which can also source B modes; see *e. g.*, Magueijo and Brandenberger 2000 for a review). As I will discuss below, this is (further) complicated by the presence of dominant secondary sources (from weak gravitational lensing and foreground Galactic polarised microwave emission); the POLARBEAR experiment has made the first detection of lensing B modes (The Polarbear Collaboration: P. A. R. Ade et al., 2014). There are numerous current and upcoming experiments in pursuit of primordial and lensing B-mode detection (see *e. g.*, Errard et al. 2015 for a summary and forecast on their cosmological constraining power). The current best limit on the tensor-to-scalar ratio from CMB B modes is from the BICEP2 and Keck Array experiments (combined with other datasets: $r < 0.07$ [14]; BICEP2 Collaboration et al., 2016).

Figure 1.2 shows the highest-resolution full-sky map to date of the CMB temperature anisotropies. Although most of the information can be extracted in the angular power spectrum (since the modes

[14] 95 % confidence limit.

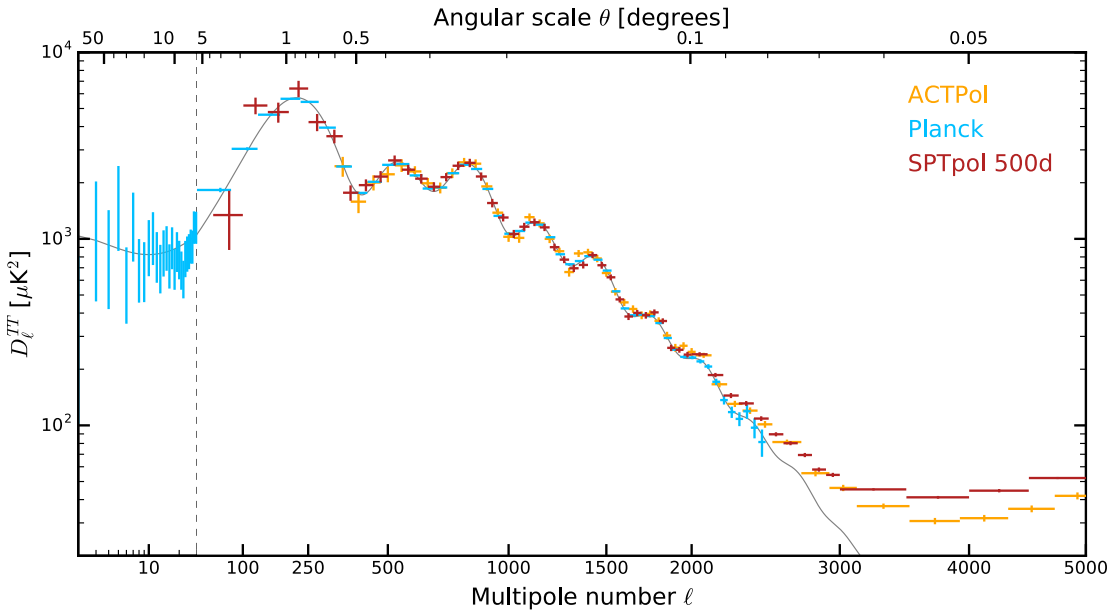

Figure 1.3: The CMB temperature angular power spectrum ($D_\ell^{\rm TT} = \ell(\ell+1)C_\ell^{\rm TT}$). The data-points are measurements from the ACTPol (Louis et al., 2017), *Planck* (Planck Collaboration et al., 2016f) and SPTpol (Henning et al., 2017) experiments. The grey line is the best-fit ΛCDM model as inferred by *Planck*. This figure is taken from Henning et al. (2017).

are almost exactly Gaussian distributed), maps still have a vital use. This includes searching for the signatures of *e. g.*, anisotropic backgrounds and primordial non-Gaussianity; secondary contributions like the integrated Sachs-Wolfe (ISW) effect (see below), reconstruction of the lensing potential and the Sunyaev-Zel'dovich (SZ) effect (Sunyaev and Zeldovich, 1970); and to be used in cross-correlation with other non-CMB maps. Maps are also made of the polarisation Q, U, E and B modes, with equivalent uses. Since at any single frequency, a map of the CMB is contaminated by foreground emission (see below), it is necessary to appropriately combine maps as measured at different frequencies in order to separate the cosmological and foreground signals. Ways to improve this "component separation" in both temperature and polarisation are the focus of Chapters 2 and 3.

Figure 1.3 shows the best-fit ΛCDM model CMB temperature angular power spectrum as inferred from observations by the *Planck* satellite. There is an excellent agreement between data and model for a wide range of angular scales. The acoustic oscillations in the primordial baryon-photon plasma (from which the CMB photons decoupled) discussed in § 1.3.3 leave a clear imprint on the angular power spectrum, in the periodic peaks and troughs. Modes that have undergone a half-integer number of oscillations at recombination are maximally rarefied or compressed at photon decoupling. It follows that structure on these scales, which are the harmonics of the sound horizon

at recombination, is enhanced relative to intermediate scales – forming the acoustic peaks and troughs. Another striking feature of the angular power spectrum is the damping tail on small scales. This arises because the finite speed sound in the plasma means that photons and baryons are not perfectly coupled and the photons have a small non-zero mean free path. This smooths out structure smaller than a characteristic diffusion length-scale (Silk, 1968). There is a slight increase in power on the very largest scales due to the ISW effect (Rees and Sciama, 1968), where the energy of CMB photons changes because they travel through evolving potentials (which decay at late times due to dark energy). The overall scale-dependence of the angular power spectrum indicates the scale-dependence of the primordial power spectrum, especially on large scales, where the modes were super-horizon at recombination and so less evolved than those that had been in causal contact. The positions of the acoustic peaks constrain a combination of the curvature and dark energy density fractions, due to their geometric effect on the angular diameter distance to the last scattering surface. The ratios of the odd and even numbered peaks constrain the baryon energy density fraction, since an increase in the baryon fraction increases this ratio (and vice versa). This is because an increased baryon fraction causes the baryon-photon plasma to compress further than it rarefies. The amplitudes of the small-scale peaks constrain the dark matter energy density fraction, since an increased fraction (at fixed baryon fraction so that the total matter fraction increases) reduces the effect of radiation driving, where in radiation-domination, gravitational potentials decay, increasing plasma oscillation amplitudes. Measurement of the damping tail constrains the fraction of relativistic matter (in particular, neutrinos). An important late-time effect is weak gravitational lensing of the CMB by matter between the last scattering surface and us. The arcminute deflections of CMB photons tend to smooth the amplitudes of the small-scale peaks and troughs.

The E-mode angular power spectrum (first detected by DASI; Kovac et al. 2002) and the TE cross-spectrum (first detected by WMAP; Kogut et al. 2003) have also been measured to high precision up to very small scales (BICEP2 and Keck Array Collaborations et al., 2015; Henning et al., 2017; Louis et al., 2017; Planck Collaboration et al., 2016f). Figure 1.4 illustrates the challenge of primordial B-mode detection. It is already an intrinsically low-amplitude signal, many orders of magnitude fainter than the T and E modes, which therefore requires very sensitive instrumentation to be observed. In addition, on arcminute scales, it is superposed with modes generated by weak gravitational lensing of E modes. These modes contain important information about the matter distribution at high redshift (particularly at $z \sim 2$) and will place constraints on neutrino properties. As mentioned above, this signal has now been detected (Ade et al., 2015; Keisler et al., 2015; The Polarbear Collaboration: P. A. R. Ade et al., 2014). As will be discussed below, both these signals are

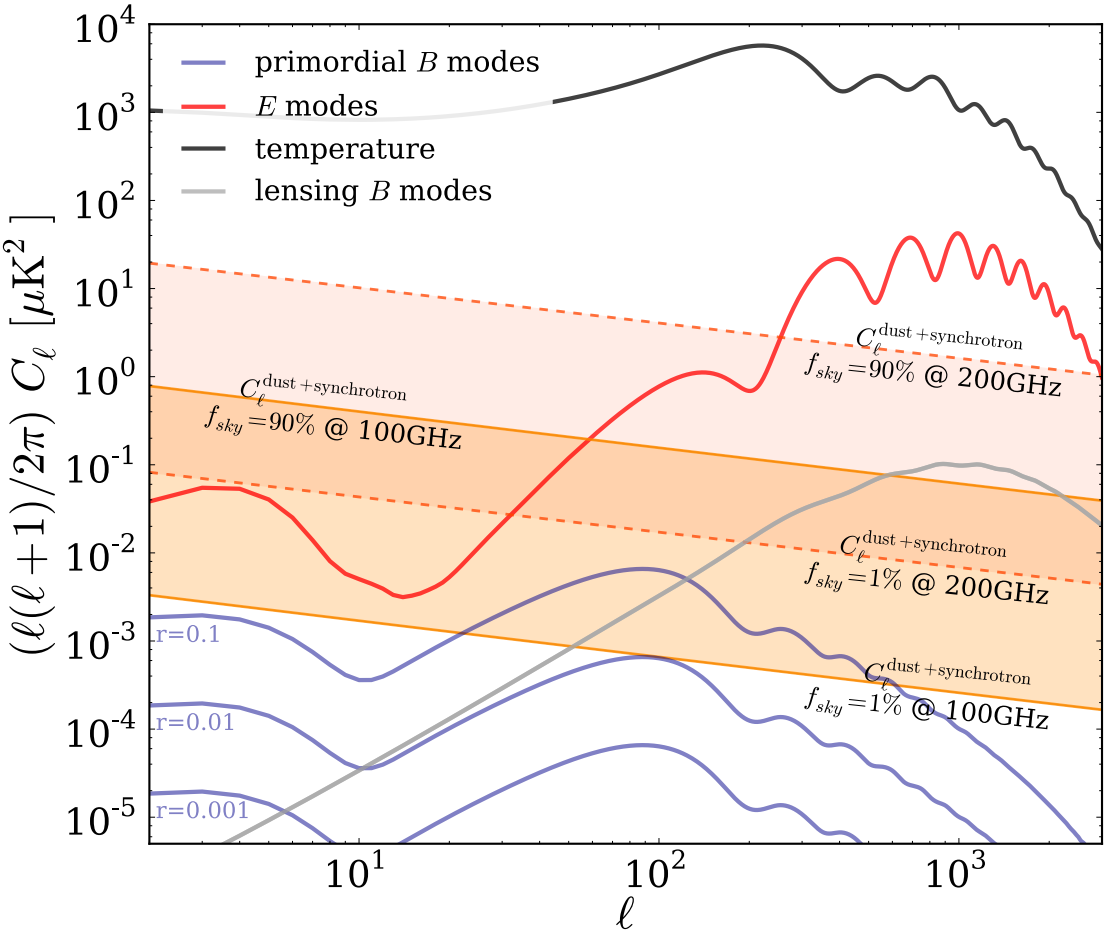

Figure 1.4: The CMB temperature TT and polarisation EE and BB angular power spectra as predicted by the ΛCDM model. The BB angular power spectrum is split into its possible contributions. The contribution from Galactic polarised dust and synchrotron emission is shown in the coloured bands, for different fractions of the full sky and as seen at different frequencies. The contribution from weak lensing modes is the grey line, peaking on arcminute scales. The contribution from primordial modes is shown for different values of the tensor-to-scalar ratio r. This figure is adapted from Errard et al. (2015).

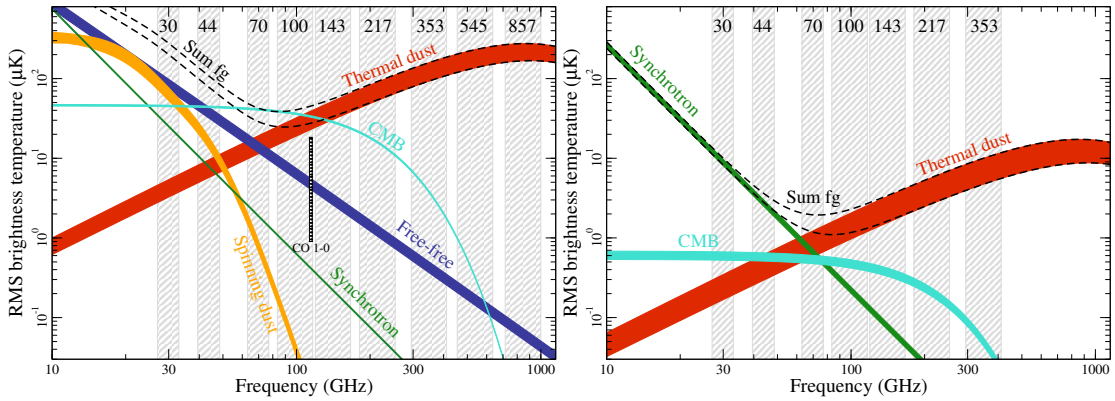

Figure 1.5: RMS brightness temperature as a function of frequency for various astronomical foreground components and the CMB for (*left*) temperature and (*right*) polarisation. The grey bands in the background show the frequency response of each *Planck* channel. This figure is taken from Planck Collaboration et al. (2015a).

contaminated by the polarised emission from Galactic dust and synchrotron. Even at the frequency where the foreground signal is smallest and in the cleanest 1 % of the sky, the primordial signal is very contaminated. The development and testing of methods to overcome this obstacle is the focus of Chapter 3.

Figure 1.5 shows the RMS brightness temperature as a function of frequency for astrophysical microwave-frequency foregrounds and compares them to the CMB. It can be seen that for temperature anisotropies, there is only a small window of frequencies (around 80 GHz) where the CMB is of the same order of magnitude as the sum of contaminants. Meanwhile, for polarisation, foregrounds are dominant at all frequencies.

A significant source of contaminating radiation to the CMB temperature anisotropies derives from warm (about 10 - 100 K) interstellar dust grains in the Milky Way (see review in *e. g.*, Dunkley et al. 2009; red bands in Fig. 1.5). These emit thermal radiation and dominate the foreground signal at high frequencies. This radiation comes from almost the entire sky, though it is particularly concentrated in the Galactic plane, as is most foreground signal (Planck Collaboration et al., 2016b). Free-free emission or *bremsstrahlung* (dark blue band in Fig. 1.5) comes from electron-electron scattering events in warm (10^4 K) ionised gas in the interstellar medium (ISM). This is a dominant foreground at low frequencies. Synchrotron photons are emitted by electrons accelerated in the Galactic magnetic field (green bands in Fig. 1.5). Synchrotron emission, though still concentrated along the Galactic equator, extends far into the "cirrus" part of the Milky Way at higher Galactic latitudes. Spinning dust emission (Draine and Lazarian, 1998, yellow band in Fig. 1.5) is the

preferred explanation for observed anomalous microwave emission (Kogut et al., 1996). It would arise from the electric dipole of rapidly spinning (10 - 60 GHz), small (nm) dust grains, most probably polycyclic aromatic hydrocarbons. There are also various line emissions at the microwave frequencies of interest, in particular from rotational transitions of carbon monoxide (CO; the CO $J = 1 \rightarrow 0$ emission is the black and white band in Fig. 1.5).

Polarised foregrounds are brighter (with respect to the CMB; see Fig. 1.5) and even more complex. The thermal emission from dust grains is polarised due to the alignment of aspherical grains in the Galactic magnetic field. The grains experience radiative torque by photon interactions (Planck Collaboration et al., 2015e). When spinning, the grains align their longest axis perpendicular to the local magnetic field. The grains then emit and absorb radiation most efficiently along their long axis forming a polarisation which traces the structure of the magnetic field in the Milky Way. What is more, there are multiple populations of dust grains: a mixture of silicates and carbonaceous grains with different temperatures, emissivities and polarisation efficiencies lying in different magnetic field alignments. This forms a very complicated foreground signal, again concentrated along the Galactic equator. Polarised thermal dust emission dominates over the CMB at high frequencies. Synchrotron emission is also polarised and again traces interstellar magnetic fields. Polarised synchrotron emission dominates over the CMB at low frequencies. Other polarised foregrounds include free-free emission (polarised by Thomson scattering at the edges of HII regions) and so-called anomalous emission (electric or magnetic dipole emission from rapidly rotating dust grains).

Another set of foregrounds is compact or point sources of both Galactic and extragalactic origin. These normally undergo special handling in CMB analysis by either subtracting flux or masking appropriately. Planck Collaboration et al. (2015f) gives a catalogue of compact sources as identified in *Planck* data.

As mentioned above, the CMB is gravitationally lensed by intervening matter between the last scattering surface and the observer (see Lewis and Challinor 2006 for a good review). The photon deflections are of order arcminute scales; CMB lensing is the dominant source of power in the TT angular power spectrum for $\ell \gtrsim 3000$ (in the absence of other secondary effects). On larger scales, lensing smoothes the acoustic peaks and troughs of the TT power spectrum at the percent level (Hu, 2000). It similarly smoothes the TE and EE angular power spectra. As mentioned above, the lensing of CMB E modes is a significant source of B modes on arcminute scales (see Fig. 1.4), a contamination to the possible primordial signal. Lensing of the CMB also induces non-Gaussian signatures (Amblard et al., 2004; Jain et al., 2000), which can be measured by higher-order point

functions. The deflection angle of CMB photons can be expressed to first order as the gradient of a "lensing potential," itself an integrated form of the gravitational potential along the line-of-sight. This lensing potential can be reconstructed using statistical properties of the unlensed CMB (Hu, 2001a,b; Hu and Okamoto, 2002; Zaldarriaga and Seljak, 1999). On large scales, the gravitational potential is Gaussian and hence so is the lensing potential; all the information can be captured by a lensing potential power spectrum, which peaks on degree scales. Furthermore, the CMB can then be "de-lensed;" this is of particular interest for removing contamination to primordial B modes (Hu, 2002).

Lensing is not simply a contamination to the unlensed CMB; the reconstruction of the lensing potential is a powerful cosmological probe in itself. The kernel of the integration of the gravitational potential that forms the lensing potential peaks at higher redshifts ($z \sim 2$) than the equivalent for galaxy weak lensing ($z \lesssim 0.5$; see § 1.6.2) and so the large-scale (integrated) matter distribution can be probed at high redshift. Moreover, *e. g.*, the lensing potential power spectrum is more sensitive to neutrino masses than the unlensed CMB anisotropies (Kaplinghat et al., 2003) and the cross-correlation to the TT angular power spectrum is sensitive to the ISW effect (Seljak and Zaldarriaga, 1999). CMB lensing was first detected by cross-correlation between the CMB as observed by WMAP and counts of radio galaxies (Smith et al., 2007). It was first detected in the CMB alone by ACT (Das et al., 2011b).

1.6.2 The large-scale structure of matter

Cosmological observation extends beyond the CMB radiation anisotropies presented in § 1.6.1; we also learn about the Universe by measuring the matter inhomogeneities. The CMB is a powerful source of information especially about the early Universe and is less affected by late-time evolution. The large-scale structure of matter informs us about this late-time evolution (*e. g.*, the effect of dark energy), as well as the initial conditions from which the matter perturbations grew. Moreover, there is intrinsically more statistical power in large-scale structure, since the number of CMB modes is limited by the two-dimensional surface from which they came (there are $\sim 10^6$ accessible temperature modes). There is more information to be accessed through CMB polarisation and the CMB is partly sensitive to the late-time matter distribution through *e. g.*, CMB lensing. Nonetheless, current and future surveys of tracers of the matter distribution will access $\gtrsim 10^9$ modes throughout the three-dimensional volume from us to high redshifts.

Like the angular power spectra for the CMB, most of the information in the matter fluctuations

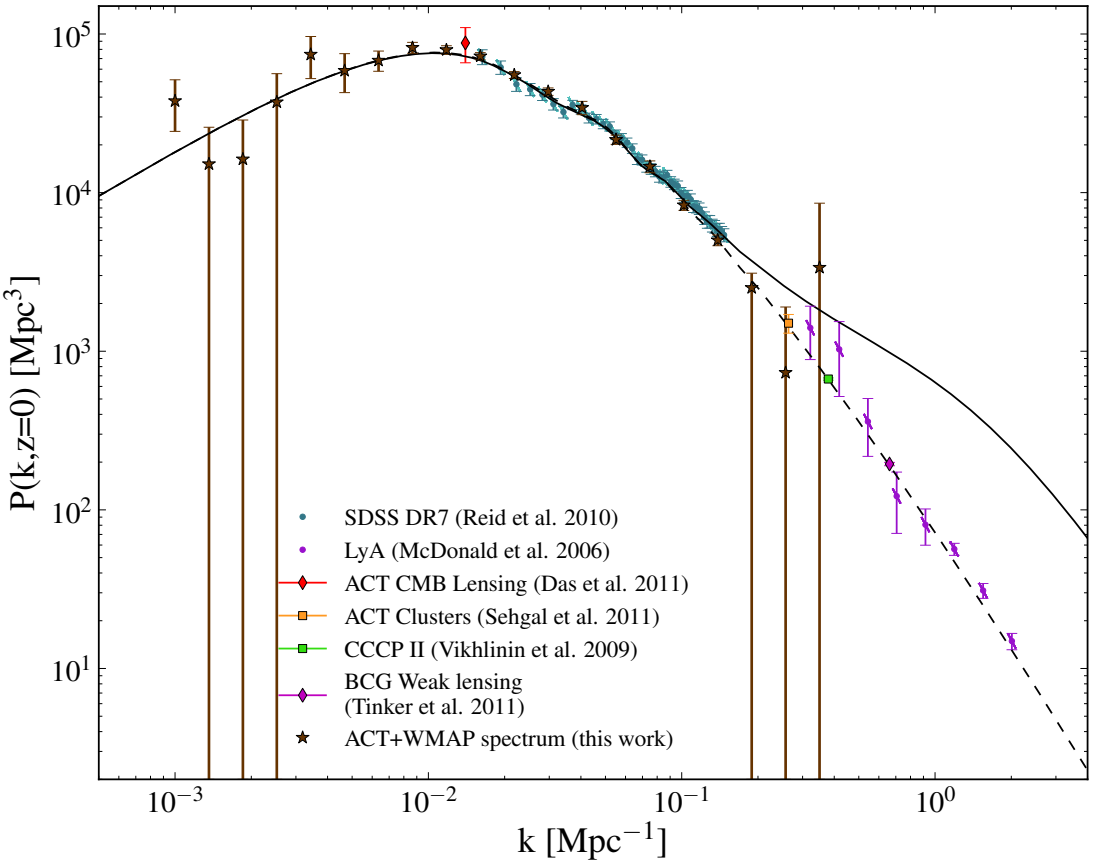

Figure 1.6: The matter power spectrum today ($z = 0$). The dotted line is the linear power spectrum and the solid line is the non-linear power spectrum; both are the best-fit ΛCDM model as inferred by CMB observations from ACT. The non-linearities are computed using HALOFIT (Smith et al., 2003). The data-points are as indicated in the legend; details of the datasets are given in Hlozek et al. (2012). The Lyman-alpha forest (LyA) is discussed in § 1.6.3. This figure is adapted from Hlozek et al. (2012).

(on large scales $k \lesssim 0.15\,h\,\mathrm{Mpc}^{-1}$; Sefusatti and Scoccimarro, 2005) is encapsulated by their two-point function[15]; in Fourier space, this is the matter power spectrum. Figure 1.6 shows the matter power spectrum as it is inferred to be today in the ΛCDM model. The dotted line is the linear power spectrum and is the expectation from linear perturbation theory (as set out in § 1.3.3); in the linear regime, the Fourier modes evolve independently. The main features of the linear power spectrum are determined by when the modes enter the horizon. The very largest modes are super-horizon and unaffected by causal processes. On intermediate scales, the modes enter the horizon after matter-radiation equality and grow proportional to the scale factor during matter-domination. All modes stop growing in Λ-domination and stay constant as gravitational potentials decay in the accelerated expansion. However, the smallest scales enter the horizon before matter-radiation equality and during radiation-domination, their growth is suppressed by the radiation pressure. This means that the power in these modes is suppressed by a factor of $\sim k^{-4}$ relative to the larger scales. When this is superposed on the (nearly) scale-invariant initial power spectrum $P_\mathrm{m}(k) \propto k$ (Harrison, 1970; Peebles and Yu, 1970; Zeldovich, 1972)[16], the small-scale tail scales $\propto k^{-3}$, forming the characteristic turn-over at the horizon scale at matter-radiation equality. The other significant feature of the matter distribution is the imprint from the acoustic oscillations in the primordial baryon-photon plasma (BAO; see § 1.3.3 and 1.6.1) at the scale of the sound horizon at recombination. Although this feature is more distinctive in the real space two-point function (the correlation function), it is observable in the power spectrum as wiggles at $k \sim 0.04\,\mathrm{Mpc}^{-1}$. In Fig. 1.6, the solid line is the non-linear power spectrum. It is enhanced on small scales relative to the linear power spectrum due to the non-linear gravitational collapse of structures and is sensitive to the baryonic physics of galaxy formation. It is understood through cosmological simulations, which I will discuss in § 1.7. The measurements of the linear power spectrum from different observables and surveys are in excellent agreement with the ΛCDM theory; details of the different datasets are given in Hlozek et al. (2012). I highlight here the measurements of the small-scale linear power spectrum from the Lyman-alpha forest (LyA), which I will discuss in § 1.6.3.

The total matter distribution (which includes baryons and dark matter) is inferred by the observation of tracers. These are in general luminous objects like galaxies, galaxy clusters or quasars (high-redshift active galactic nuclei); or diffuse fields like neutral hydrogen through *e. g.*, 21 cm emission or Lyman-alpha forest absorption (see § 1.6.3). The dark matter distribution can be more

[15]There is more information in the higher-order point functions and this is an active area of research.

[16]The (near) scale-invariance refers to the dimensionless primordial curvature power spectrum (Eq. (1.11)). This implies in the weak-field limit that the dimensionless power in fluctuations in the gravitational potential ($k^3 P_\Phi$) is constant, where Φ is the fluctuation in the potential. By Poisson's equation, the matter fluctuation $\delta_\mathrm{m} \propto k^2\Phi$ and so the matter power spectrum $P_\mathrm{m}(k) = \langle|\delta_\mathrm{m}|^2\rangle \propto k^4\langle|\Phi|^2\rangle \propto k^4 P_\Phi \propto k$.

directly probed by measuring the weak gravitational lensing of the CMB or galaxies, whose observed shapes are distorted by intervening matter. Galaxy surveys have historically informed us about the large-scale distribution of matter. One of the first such datasets was the Lick galaxy catalogue (1967; re-mapped by Seldner et al., 1977). As early as the 1970s, it was realised that galaxies are not randomly distributed but form distinct structures (what we today term the "cosmic web"; see § 1.3.3). The Lick catalogue only gave the angular positions of galaxies on the sky (being formed only from photographic plates). The first major three-dimensional survey (measuring redshifts as well as angular positions) was the CfA Redshift Survey (1977 - 1982; Davis et al., 1982), which further revealed the clusters, filaments and voids that galaxies assemble. This was followed by CfA2 (1985 - 1995; de Lapparent et al., 1986), which, among other things, discovered the "Great Wall" (Geller and Huchra, 1989), a sheet of galaxies spanning hundreds of Mpc. The formation of such large structures was not understood at the time, but these are now understood to have been sourced in primordial inflationary perturbations (see § 1.4). These surveys were superseded by *e. g.*, APM (1990s; Maddox et al., 1990) and the 2dF Galaxy Redshift Survey (1997 - 2002; Peacock et al., 2001), whose large survey volumes allowed the precise determination of cosmological parameters like the matter fraction. The major modern galaxy survey is SDSS (1998 - today; Abazajian et al., 2003), within which latterly BOSS (2008 - 2014; Dawson et al., 2013) has the task of mapping the three-dimensional distribution of luminous red galaxies and quasars.

Indeed, there are numerous current and upcoming photometric, spectroscopic and radio surveys. Spectroscopic surveys (*e. g.*, BOSS[17], DESI[18]) measure the full spectrum of the EM radiation from galaxies or quasars and are able to accurately measure their redshifts. Photometric surveys (*e. g.*, SDSS[19], DES[20], LSST[21], *Euclid*[22]) only measure the spectrum of the light from these objects in a few integrated photometric bands and have to estimate the redshifts from this limited information. Radio surveys (*e. g.*, SKA[23]) *e. g.*, aim to measure the high-redshift 21 cm emission.

These observables do not exactly map the dark matter distribution and in general will be biased tracers. On large scales, this bias can be linearly modelled (Kaiser, 1984), such that the overdensity in the tracer δ_s can be related to the matter overdensity δ_m by a linear bias parameter b:

$$\delta_\mathrm{s} = b(1 + \beta\mu^2)\delta_\mathrm{m}, \tag{1.13}$$

[17] http://www.sdss3.org/surveys/boss.php
[18] http://desi.lbl.gov
[19] http://www.sdss.org
[20] http://www.darkenergysurvey.org
[21] https://www.lsst.org
[22] https://www.euclid-ec.org
[23] https://skatelescope.org

where μ is the cosine of the angle away from the line-of-sight, such that along the line-of-sight $\mu = 1$. Here, I have additionally included the multiplicative term $(1 + \beta\mu^2)$ to model large-scale linear redshift-space distortions (RSD; Kaiser, 1987), where β is the linear RSD parameter. This is an observational effect where on large scales, overdensities appear compressed towards the line-of-sight. This is due to the coherent peculiar velocities of objects on the edge of an overdensity pointing towards the centre. This changes their measured redshift such that objects on the far side of an overdensity along the line-of-sight appear closer than they are and objects on the near side appear further away. This "flattens" the overdensity along the axis of the line-of-sight and enhances the apparent overdensity by a factor β, the redshift-space distortion parameter[24]. There is no effect in the direction transverse to the line-of-sight since only line-of-sight velocities are measured. Objects like galaxies or galaxy clusters in general have a bias $b > 1$, since they form at the peaks of the underlying matter density distribution (although this is not true for all populations). However, some tracers are less clustered than the dark matter and have a bias $b < 1$. Indeed, *e. g.*, the transmitted flux of the Lyman-alpha forest has $|b| < 1$ and, in fact, the bias is negative because more flux is transmitted in matter underdensities (since there is less absorption by neutral hydrogen; see § 1.6.3 for more details).

1.6.3 The Lyman-alpha forest

The Lyman-alpha forest is the absorption in quasar spectra from neutral hydrogen (HI) along the line-of-sight due to the Lyman-alpha transition from the ground to the first excited state. Quasars are very luminous sources at the centres of galaxies (essentially the emission from central accretion discs around supermassive black holes) and as a consequence can be observed at high redshifts. As their radiation travels towards us, the Lyman-alpha line (at a rest wavelength of 1216 Å) is absorbed by intervening HI (and re-emitted isotropically so as to leave the absorption line). However, the lines from previous absorption are continuously redshifted such that a series (or "forest") of lines is formed. The first proposed cosmological use was to search quasar spectra for the Gunn-Peterson trough (Gunn and Peterson, 1965). If the intergalactic medium (IGM) is completely neutral then the absorption lines of the forest would form a single absorption "trough" in the spectrum. The lack of a Gunn-Peterson trough at low redshifts showed that the Universe reionised and the identification of troughs at high redshifts helps to determine the redshift of (hydrogen) reionisation (see § 1.3.3). The first Gunn-Peterson trough was identified at $z = 6.3$ (Becker et al., 2001); the exact connection

[24]This is the opposite effect as the small-scale non-linear RSD which form the "fingers of God." Here, random peculiar velocities in collapsed structures make overdensities appear elongated along the line-of-sight.

to the redshift of reionisation is complicated by its patchy nature. Croft et al. (1998) proposed that the statistics of the absorption could be used to infer the clustering of matter at the high redshifts at which the Lyman-alpha forest is observed ($2 \lesssim z \lesssim 6$). This was supported by the advances in simulating the Lyman-alpha absorption in a cosmological context (Hernquist et al., 1996, see § 1.7), allowing the absorption statistics to be related to cosmological parameters.

The gas forming the Lyman-alpha forest is low-density (specifically defined as having column density $< 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$; Wolfe et al., 1986). This means that there exists a tight relation between the density (ρ) and temperature (T) of the gas ($T \propto (\rho/\overline{\rho})^{\alpha}$, where $\overline{\rho}$ is the mean density and α is a constant), determined by the rates of photoionisation heating and adiabatic cooling (Hui and Gnedin, 1997); although, as discussed below, uncertainty in this relation limits cosmological parameter inference. Simulations of the Lyman-alpha forest show that there is no clear distinction between individual absorption lines and the Lyman-alpha forest is better characterised as absorption formed by a smoothly fluctuating IGM (Hernquist et al., 1996). Although the redshifts at which the Lyman-alpha forest is observed are later than the epoch of reionisation, because Lyman-alpha absorption has a particularly large cross-section, even a very small fraction of neutral hydrogen remaining will form a saturated absorption line. At $z \sim 6$, when the Gunn-Peterson trough is giving way to the Lyman-alpha forest, the neutral fraction is already $\sim 10^{-6}$ (Haehnelt, 1995) and it continues to decrease as the Universe evolves.

When quasar sightlines are closely separated, modern surveys can use the absorption lines of the Lyman-alpha forest to tomographically map the three-dimensional distribution of HI and hence infer the total matter distribution (Caucci et al., 2008; Krolewski et al., 2017; Lee et al., 2014; Pichon et al., 2001). Wide-angle surveys (*e. g.*, BOSS and, in the future, DESI) measure the statistics (*i. e.*, the two-point function) of the fluctuations in the transmitted flux. On large scales, these fluctuations are modelled as biased tracers of the matter density fluctuation distribution with redshift space distortions (Eq. (1.13)) and the correlation function (and cross-correlation to the quasar distribution) is used to measure the BAO feature (see § 1.6.2) at high redshift ($z \sim 2.3$; Bautista et al., 2017; Busca et al., 2013; Delubac et al., 2015; du Mas des Bourboux et al., 2017; Font-Ribera et al., 2013, 2014a; Kirkby et al., 2013; Slosar et al., 2011, 2013). On small scales, the one-dimensional flux power spectrum (integrated over directions transverse to the line-of-sight) is used to measure the small-scale linear matter power spectrum (see Fig. 1.6; Armengaud et al., 2017; Iršič et al., 2017a,b; Palanque-Delabrouille et al., 2015; Seljak et al., 2005; Yeche et al., 2017). The one-dimensional power spectrum can actually probe a wide range of (line-of-sight) scales ($0.1\,h\,\mathrm{Mpc}^{-1} \lesssim k_{||} \lesssim 10\,h\,\mathrm{Mpc}^{-1}$). This is achieved by combining a large number of lower

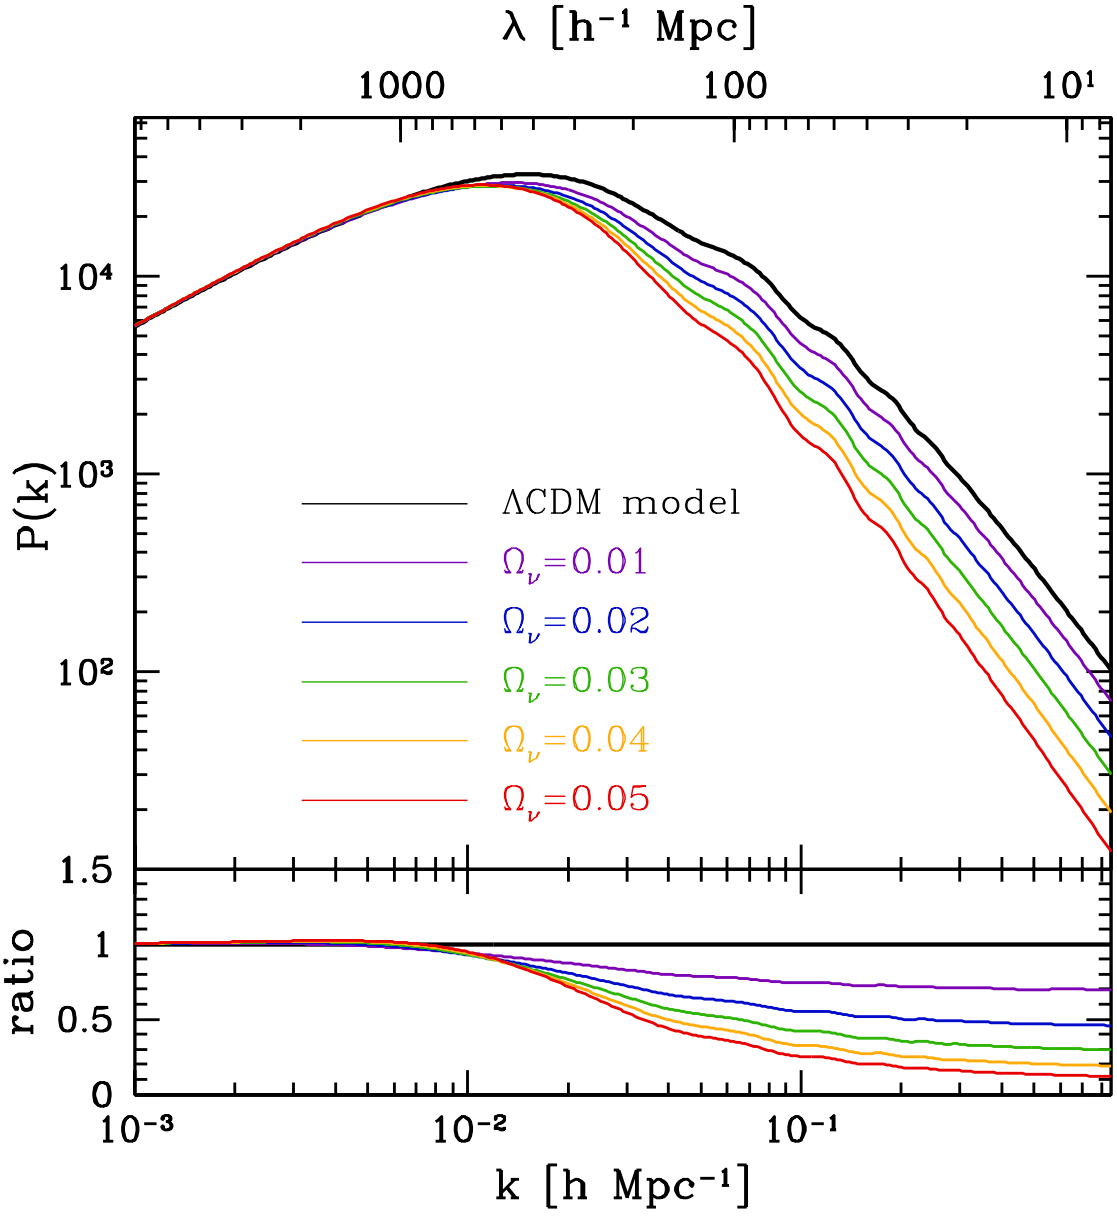

Figure 1.7: The ΛCDM model matter power spectrum for different values of the neutrino energy density today as a fraction of the total Ω_ν (for fixed total (cold and hot) dark matter fraction $\Omega_c + \Omega_\nu$). This figure is adapted from Park et al. (2012).

resolution and lower signal-to-noise spectra (*e. g.*, BOSS) on larger scales with a smaller number of high resolution and high signal-to-noise spectra on smaller scales (*e. g.*, Kim et al., 2004; Viel et al., 2004b, 2013). In current (eBOSS) and future (DESI) surveys, the power spectrum measurements will be expanded to include the full three-dimensional flux power spectrum, providing many more modes (Font-Ribera et al., 2017).

Figure 1.7 shows the effect of massive neutrinos on the linear matter power spectrum (see also § 1.5). The greater the neutrino energy density fraction today Ω_ν (and hence the greater the sum of the masses $\sum m_\nu$), the greater the small-scale suppression in power. The small-scale Lyman-alpha forest flux power spectrum is very sensitive to this suppression. Indeed, the best upper limit on $\sum m_\nu$ comes from the one-dimensional Lyman-alpha forest flux power spectrum in combination with the CMB angular power spectrum to measure the large scales (Palanque-Delabrouille et al., 2015). As mentioned in § 1.5, WDM similarly suppresses small-scale power and the Lyman-alpha forest is sensitive to the mass of the WDM particle (*e. g.*, Iršič et al., 2017b). The small-scale Lyman-alpha forest power spectrum is also sensitive to the shape of the primordial power spectrum, *e. g.*, the running of the spectral index (see § 1.4). There is more discussion about the current and future status of Lyman-alpha forest observations and their constraining power in § 4.2 and 5.2.

The Lyman-alpha forest is not only a powerful cosmological probe, but it is also sensitive to gas and ionisation physics. This requires cosmology and the physics of the IGM to be simultaneously constrained, relying on accurate simulations of the IGM (see § 1.7). Indeed, degeneracies arise between the two, *e. g.*, between neutrino or WDM masses and the thermal history of the IGM. A large systematic uncertainty arises from estimation of the mean transmitted flux, which depends on the neutral fraction of hydrogen and the quasar emitted continuum. This affects the amplitude of the absorption correlations and so is partially degenerate with the amplitude of the density perturbations (McDonald et al., 2005b), weakening cosmological constraints. There is also uncertainty in estimating the temperature and temperature-density relation of the gas, reflecting a lack of knowledge about the ionising photon background from stars and quasars (see § 1.3.3).

In the data analysis of the Lyman-alpha forest, the leading "foreground" contamination is the presence of broadened absorption features from high column density (HCD) systems of HI. These HCD absorbers have damping wings which extend along the spectrum, causing absorption away from the location of the absorber (Font-Ribera and Miralda-Escudé, 2012; McDonald et al., 2005a), which will bias inference from the Lyman-alpha forest. HCD absorbers are regions of HI gas that are dense enough to self-shield their cores from the ionising ultra-violet (UV) background (Cen, 2012) and diffuse enough to have a low star-formation rate (Fumagalli et al., 2015). Self-shielding occurs

for column densities $N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$ and sets the lower limit in column density for HCD absorbers. By contrast, lower column density absorbers form the Lyman-alpha forest and are heavily ionised by the UV background. The relation between optical depth and density changes with the onset of self-shielding; broadening from damping wings causes absorption in the spectrum away from the location of the absorbing gas. The damping wings have a characteristic Voigt profile, which is a convolution of a Gaussian profile (*e. g.*, caused by Doppler broadening) and a Lorentzian profile (*e. g.*, caused by natural or collision broadening). The width of these wings in velocity space increases with the column density of the absorbing system. High column density absorbers are then usually classified as either damped Lyman-alpha absorbers (DLAs), whose damping wings are considered significantly broadened and which correspond to $N(\mathrm{HI}) > 2 \times 10^{20}\,\mathrm{atoms\,cm^{-2}}$ (Wolfe et al., 1986); or Lyman-limit systems (LLS), which correspond to column densities in the range $2 \times 10^{20}\,\mathrm{atoms\,cm^{-2}} > N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$. The systems are formed at peaks of the underlying density distribution; consequently, they cluster more strongly than the forest itself (Font-Ribera et al., 2012b; Pérez-Ràfols et al., 2017). They are usually identified with the gas in or around galaxies, but this remains an active area of research because observation of the galaxies hosting DLAs is complicated by the presence of the illuminating background quasar on the line-of-sight. Current observations that have been made suggest that the hosts are massive star-forming galaxies, which are situated within metal-enriched HI gas reservoirs that extend far beyond the star-forming interstellar medium (*e. g.*, Neeleman et al., 2017). Attempts are made to remove these systems in quasar spectra before cosmologcal inference (*e. g.*, Garnett et al., 2017; Lee et al., 2013), but there is always a residual contamination. The construction of models from cosmological hydrodynamical simulations (§ 1.7) for their effect on one-dimensional and three-dimensional correlations in the Lyman-alpha forest is the focus of Chapters 4 and 5.

1.7 Cosmological hydrodynamical simulations

The work presented in Chapters 4 and 5 makes use of cosmological hydrodynamical simulations from the Illustris project (Nelson et al., 2015; Vogelsberger et al., 2014a). These simulations need to replicate the co-evolution of dark matter and baryons on cosmological distances and timescales. I first discuss some of the methods used to simulate the formation of structure under Newtonian gravity. Newtonian gravity is expected to be a good approximation because the Universe is dominated by dark matter and baryons, whose velocities and matter densities are non-relativistic. The full problem is to solve N-body gravitational dynamics for N particles. N-body simulations

approximate a large complex object as a point particle at its centre-of-mass (*e. g.*, Illustris-1, the simulation used in Chapters 4 and 5, has 1820^3 dark matter particles, each of mass $6.3 \times 10^6\ M_\odot$; Vogelsberger et al., 2014a). This neglects internal structure on scales smaller than the mass scale of the N-body particle (Press and Schechter, 1974; von Hoerner, 1960). In order to check that this unresolved smaller-scale structure does not have a strong impact on results, an identical simulation is run with a slightly lower mass resolution. In order to minimise boundary effects, periodic boundary conditions are used, which mimics the large-scale homogeneity of the Universe.

The full N-body dynamics as implemented numerically essentially requires evaluating the total force exerted on each particle i of mass M_i at position $\mathbf{x}_i$ in a series of small time-steps: $F_i = GM_i \sum_j M_j |\mathbf{x}_i - \mathbf{x}_j|^{-2}$ (Hernquist and Katz, 1989). This direct summation method is computationally unfeasible for a cosmological simulation. The gravitational dynamics can be described by Poisson's equation $\nabla^2\phi(\mathbf{x}) = 4\pi G\rho(\mathbf{x})$, where ϕ is the gravitational potential and ρ is the density field. Once the density field is evaluated on a grid, this equation can be solved using Fourier methods, which are much faster than direct summation (Miller and Prendergast, 1968). This however requires the approximation of the particle-mesh approach, where particles are smoothed to the width of a grid cell and interpolated onto a grid, which is then Fourier transformed and the forces computed (Doroshkevich et al., 1980). This can accurately resolve long-range forces, but has finite resolution corresponding to the grid spacing. In order to maintain the necessary dynamic range, early codes (*e. g.*, Efstathiou and Eastwood, 1981) used a particle mesh grid on large scales and resolved small scales with direct summation: the so-called particle-particle-particle-mesh, or $\mathrm{P^3M}$, algorithm. The modern development is the TreePM algorithm (Springel, 2005), which uses a particle mesh on large scales, but uses a tree algorithm on small scales. This is essentially a multipole method, which groups particles into a hierarchy of cells. TreePM is the gravity solver used in Illustris.

Hydrodynamical simulations require a description of gas physics and pressure forces. This is achieved by applying the fluid approximation and tracking fluid elements across a grid of cells. The traditional use of a fixed Eulerian grid (Richtmyer and Morton, 1967) suffers from limited resolution (like the bare particle-mesh approach above). The high dynamic range necessary for cosmological simulations is maintained by adaptive mesh refinement (AMR; Anninos et al., 1994). In this approach, grid cells are dynamically refined to a finer sub-grid as is necessary. There are many simulation codes that couple hydrodynamics using AMR to a particle-mesh gravity solver, *e. g.*, `RAMSES` (Teyssier, 2002) and `ENZO` (O'Shea et al., 2004). However, these codes suffer computational overhead from having to combine a particle description of gravity and a grid-based fluid description.

This is overcome by the use of smoothed particle hydrodynamics (SPH; Evrard, 1988; Gingold and Monaghan, 1977; Lucy, 1977), where the fluid is sampled by a number of tracer particles, each of which corresponds to a Lagrangian fluid element of constant density. SPH discretises the fluid in mass (in Illustris-1, the gas particle masses are each $1.3 \times 10^6\, M_\odot$) rather than spatial coordinate and so it combines well with the particle description of gravity. It automatically has a high dynamic range because particles concentrate in areas of high matter density. A fluid description is obtained by interpolating smoothly between particles. The formulation of SPH in the code `GADGET` explicitly conserves entropy and energy (Springel and Hernquist, 2002) and thus avoids spurious mixing. The code `AREPO` (Springel, 2010) combines the advantages of SPH and grid-based codes by discretising the fluid using a Lagrangian moving mesh; this is the code used in Illustris. There is some discussion of the galaxy formation physics implemented in Illustris of relevance to the work in Chapters 4 and 5 in § 4.4.1.

In Chapters 4 and 5, simulated spectra containing only the Lyman-alpha absorption line (*i. e.*, the Lyman-alpha forest; see § 1.6.3) are generated and used. These are formed using snapshots (*i. e.*, a view of the simulation box at a fixed redshift) from the Illustris simulation. For each spectrum, the optical depth τ is measured along a chosen line-of-sight in a series of bins, effectively as given by Eq. (A.1) except transformed to velocity space. More details are given in § 4.4.1 and Appendix A.1. The full Voigt profile of each absorption line is not needed for the majority of the HI gas, which has low column density ($N(\mathrm{HI}) < 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$) and which forms the Lyman-alpha forest, and is replaced by only the Gaussian Doppler broadening term. The necessary velocities are measured from the simulation particles. The transmitted flux in the spectrum $\mathcal{F} = e^{-\tau}$. Figure 1.8a shows some examples of the simulated Lyman-alpha forest spectra used in Chapters 4 and 5. The same sightline is shown for two different redshifts ($z = 2.44$ and 3.49), highlighting the increased amount of absorption at higher redshift due to the higher abundance of neutral hydrogen. Figure 1.8b shows an example real observed quasar spectrum (Q1422+2309), where the (currently usable) Lyman-alpha forest is the series of absorption lines from 1216 Å to 1026 Å. There is a remarkable qualitative similarity between the real and simulated Lyman-alpha forests. The remaining differences are that in the real spectrum, the quasar emitted continuum has not been normalised. Estimation of this continuum is a source of uncertainty in data analysis, since most observed spectra do not have the remarkable signal-to-noise and resolution of this example. Furthermore, the length of the simulated spectra is much shorter than the full forest observed in panel (b); this only has the visual effect of "stretching-out" the absorption lines. The conversion from comoving position to wavelength is explained in Chapter 4 and Appendix A.1.

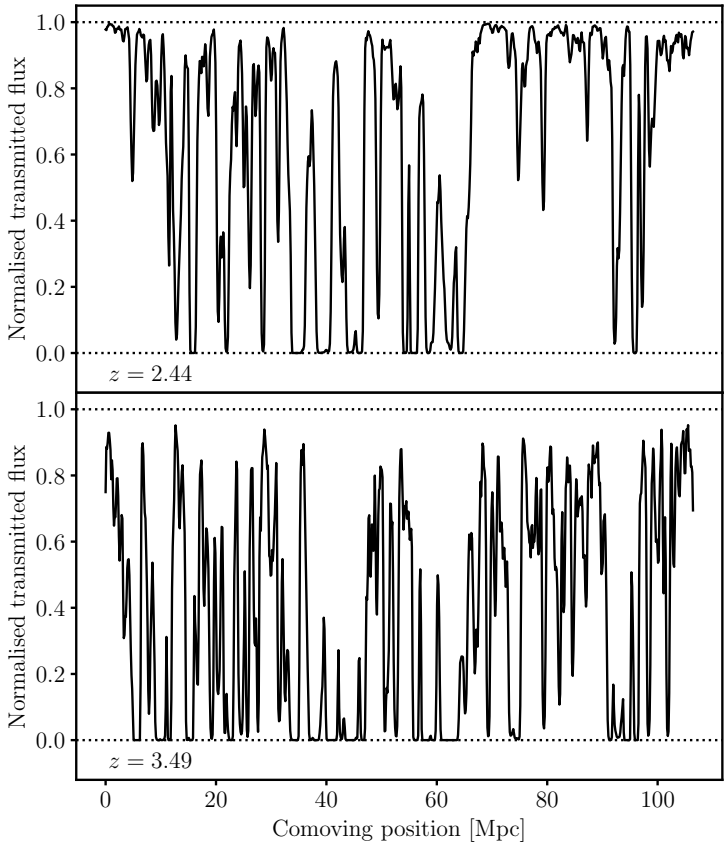

(a) Simulated Lyman-alpha forest.

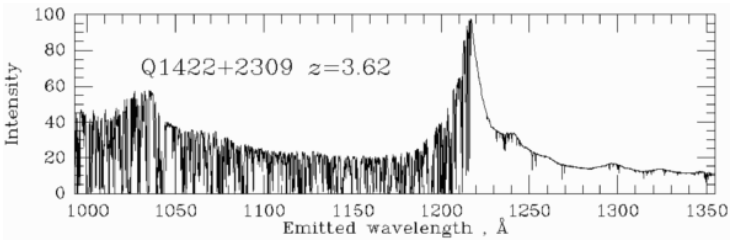

(b) Real Lyman-alpha forest.

Figure 1.8: (a) Examples of simulated Lyman-alpha forest spectra as generated from a cosmological hydrodynamical simulation (Illustris-1; Nelson et al., 2015; Vogelsberger et al., 2014a, using the `fake_spectra` code; Bird 2017). The normalised transmitted flux is shown as a function of comoving position along the sightline within the simulation box. The same sightline is shown in the two panels for (*top*) $z = 2.44$ and (*bottom*) $z = 3.49$. (b) An example of a real Lyman-alpha forest spectrum (as observed by the Keck I HIRES instrument). The normalised intensity is shown as a function of wavelength in the quasar's (Q1422+2309) rest-frame at $z = 3.62$. The full quasar spectrum is shown and the Lyman-alpha forest comprises the series of absorption lines from 1216 Å to 1026 Å; at shorter wavelengths, the Lyman-alpha forest superposes with the absorption from the Lyman-beta transition and this region is usually excised in data analysis. This sub-figure is modified from an original available at `http://pages.astronomy.ua.edu/keel/agn/forest.html`.

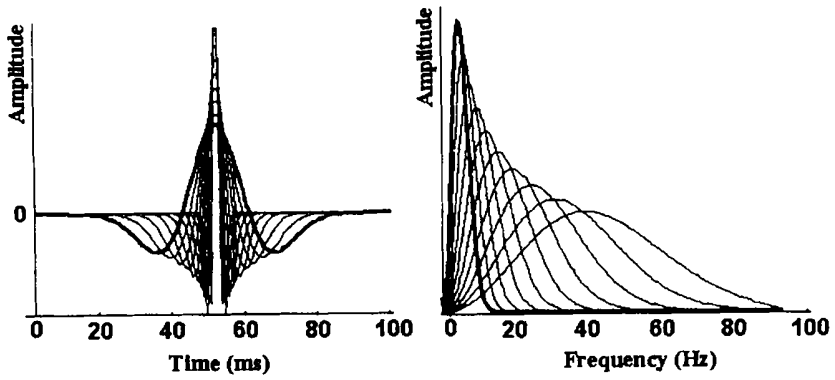

Figure 1.9: A set of Mexican hat wavelets, showing their localisation in time (*left*) and frequency (*right*) domains. This figure is taken from Couderc and Zareba (1998).

1.8 Wavelets

By the uncertainty principle, functions cannot be perfectly localised simultaneously in both real and frequency space. *I. e.*, a Dirac delta function (which is perfectly localised) in one space is unbounded (has infinite support) in the other. Wavelets are the compromise solution: they are bounded (or band-limited) in both spaces. Figure 1.9 shows a set of Mexican hat wavelets, a particular type of one-dimensional wavelet (higher-dimensional forms can be constructed). Their localisation in both real (time) and frequency space is shown. When a wavelet has more localised (compact) support in one domain, it is less localised in the other. Many different types of wavelet exist with different localisation properties and also in higher dimensions. They are very useful for analysing data with localised features. In Chapters 2 and 3, I use particular types of wavelets called (spin) scale-discretised directional wavelets (McEwen et al., 2015b) to analyse CMB temperature and polarisation anisotropies. These wavelets are defined on the two-dimensional surface of the sphere and are additionally localised in the orientation of their (non-axisymmetric) spatial kernels (*i. e.*, they are "directional"). Details of the wavelets and an introduction to related wavelet analyses on the CMB are given in Chapters 2 and 3. Indeed, wavelets are commonly used in cosmological analyses, especially for the CMB. They have been used *e. g.*, to search in CMB maps for the ISW effect (see § 1.6.1; McEwen et al., 2007, 2008), the signatures of primordial non-Gaussianity (Cayón et al., 2001; Vielva et al., 2004) and other anomalous features (*e. g.*, Feeney et al., 2011).

1.9 Thesis outline

Each of the following chapters (2 to 5; their associated appendices are collated in Appendices A and B) is a journal article of which I am the lead author, reproduced here with minor formatting changes. There is some repetition in the introductory material presented in each chapter because of this reproduction.

- **Chapter 2**: *SILC: a new* Planck *Internal Linear Combination CMB temperature map using directional wavelets*.

  This work was published as Keir K. Rogers, Hiranya V. Peiris, Boris Leistedt, Jason D. McEwen, and Andrew Pontzen, Monthly Notices of the Royal Astronomical Society, 460: 3014-3028, 2016, and was carried out in collaboration with the named co-authors.

- **Chapter 3**: *Spin-SILC: CMB polarisation component separation with spin wavelets*.

  This work was published as Keir K. Rogers, Hiranya V. Peiris, Boris Leistedt, Jason D. McEwen, and Andrew Pontzen, Monthly Notices of the Royal Astronomical Society, 463: 2310-2322, 2016, and was carried out in collaboration with the named co-authors.

- **Chapter 4**: *Simulating the effect of high column density absorbers on the one-dimensional Lyman-alpha forest flux power spectrum*.

  This work was published as Keir K. Rogers, Simeon Bird, Hiranya V. Peiris, Andrew Pontzen, Andreu Font-Ribera, and Boris Leistedt, Monthly Notices of the Royal Astronomical Society, 474: 3032-3042, 2018, and was carried out in collaboration with the named co-authors.

- **Chapter 5**: *Correlations in the three-dimensional Lyman-alpha forest contaminated by high column density absorbers*.

  This work was submitted as Keir K. Rogers, Simeon Bird, Hiranya V. Peiris, Andrew Pontzen, Andreu Font-Ribera, and Boris Leistedt, Monthly Notices of the Royal Astronomical Society, submitted, 2018, and was carried out in collaboration with the named co-authors.

2

SILC: a new Planck Internal Linear Combination CMB temperature map using directional wavelets

2.1 Abstract

We present new clean maps of the CMB temperature anisotropies (as measured by *Planck*) constructed with a novel internal linear combination (ILC) algorithm using directional, scale-discretised wavelets — Scale-discretised, directional wavelet ILC or SILC. Directional wavelets, when convolved with signals on the sphere, can separate the anisotropic filamentary structures which are characteristic of both the CMB and foregrounds. Extending previous component separation methods, which use the frequency, spatial and harmonic signatures of foregrounds to separate them from the cosmological background signal, SILC can additionally use morphological information in the foregrounds and CMB to better localise the cleaning algorithm. We test the method on *Planck* data and simulations, demonstrating consistency with existing component separation algorithms, and discuss how to optimise the use of morphological information by varying the number of directional wavelets as a function of spatial scale. We find that combining the use of directional and axisymmetric wavelets depending on scale could yield higher quality CMB temperature maps. Our results set the stage for the application of SILC to polarisation anisotropies through an extension to spin wavelets.

2.2 Introduction

Accurate measurements of the cosmic microwave background (CMB) arguably form the bedrock of modern precision cosmology. In particular, the full-sky multifrequency CMB maps provided by three generations of satellite experiments — *COBE* (Boggess et al., 1992; Mather et al., 1990), *WMAP*

(Bennett et al., 2003a) and *Planck* (Planck Collaboration et al., 2011) — represent milestones in our understanding of the cosmological model. However, to obtain a full-sky map of the CMB requires removing instrumental noise and signals due to astrophysical foregrounds (primarily in the Milky Way). Full-sky foreground-cleaned CMB maps are used for a wide variety of scientific purposes (see *e. g.*, Planck Collaboration et al., 2015d,g).

There are numerous methods to perform foreground component separation. They broadly fall into two categories: blind methods which make minimal physical assumptions about the contributing signals and the so-called mixing matrix (which quantifies the strength of different components at different frequencies) and non-blind methods which are based on a physical modelling of the sky components. Examples of non-blind methods include the Maximum Entropy Method (MEM) (Hobson et al., 1998) and the parametric Bayesian CMB Gibbs sampler Commander (Eriksen et al., 2006, 2008). Correlated Component Analysis (CCA) (Bedini et al., 2005) is a semi-blind method that estimates the mixing matrix based on second-order statistics. Spectral Estimation via Expectation Maximisation (SEVEM) (Fernández-Cobos et al., 2012; Leach et al., 2008; Martínez-González et al., 2003) is a template fitting technique. Examples of so-called blind source separation include the sparsity-based method Local-Generalized Morphological Component Analysis (L-GMCA) (Bobin et al., 2008, 2013) and the Spectral Matching Independent Component Analysis (SMICA) (Cardoso et al., 2008), although the latter work does discuss how the choice of component model affects the blindness of this method. Of particular interest to this work is another blind method, the Internal Linear Combination (ILC), most recently implemented by the Needlet ILC (NILC) (Delabrouille et al., 2009). In its component separation analysis, the *Planck* Collaboration used Commander, NILC, SEVEM and SMICA (Planck Collaboration et al., 2015c). See, *e. g.*, Bobin et al. (2013); Delabrouille et al. (2009) for reviews of CMB component separation methods.

The ILC computes a weighted sum of CMB maps as measured at multiple frequencies. These weights are constrained to sum to unity at each point in the map, ensuring that the CMB signal is conserved, assuming that it is equal at each frequency. Under this constraint, the weights are calculated by minimising the empirical variance of the ILC map, which in turn minimises the variance of the error in CMB reconstruction (assuming the CMB and foregrounds and the CMB and noise are respectively uncorrelated). The variances we minimise are empirical in that they are calculated using the data themselves. In this way, the weights are calculated to remove foreground and noise, revealing the underlying primordial CMB anisotropies. The ILC method was originally used by the *WMAP* Collaboration (Bennett et al., 2003b) and then extended by Eriksen et al. (2004) through an analytical calculation of the weights. One limitation of the original ILC approach is the

extent of localisation of the weights. The initial versions calculated different weights in separate parts of the sky (*e. g.*, Bennett et al. 2003b split the Galactic region into 11 parts). In order to further remove local contamination, the weights can be allowed to vary across the sky and also at different spatial scales. Tegmark et al. (2003) made an ILC map allowing the weights to vary at each multipole, as well as within different regions of the sky. A direct extension of this work is to make use of both spatial and frequency information simultaneously using wavelets. The weights are then defined across wavelet scales and within wavelet coefficient maps on the sky.

Wavelets are functions that are localised in both real and frequency space. To analyse full-sky CMB maps, wavelets defined on the surface of a sphere are required. A number of wavelet frameworks on the sphere have been developed recently (Antoine and Vandergheynst, 1998, 1999; Baldi et al., 2009; Barreiro et al., 2000; Geller and Marinucci, 2010, 2011; Geller et al., 2008; Leistedt et al., 2013, 2015; Marinucci et al., 2008; McEwen and Scaife, 2008; McEwen et al., 2006, 2011; McEwen et al., 2013, 2014; McEwen et al., 2015a,b; Narcowich et al., 2006; Sanz et al., 2006; Starck et al., 2009, 2006; Wiaux et al., 2005; Wiaux et al., 2008). In particular, needlets (Baldi et al., 2009; Marinucci et al., 2008; Narcowich et al., 2006) have been used in the latest generation of ILC methods. Needlets are a set of axisymmetric kernels defined on the surface of a sphere. Each member of the set has compact support in harmonic space over different multipole ranges. When each needlet is convolved with a signal defined on the sphere, the resulting signal (*i. e.*, needlet coefficients) also has compact harmonic support. NILC (Delabrouille et al., 2009) computes its weights by considering needlet scales separately (harmonic localisation) and then different parts of each needlet coefficient map separately (spatial localisation). The needlets are constructed in such a way that the original signal can be recovered from its needlet coefficients with no loss of information (in practice, small losses can be introduced by approximate spherical harmonic transforms). NILC has been very successful at forming clean full-sky CMB maps, which contain very little residual foreground and noise contamination.

In this work, we introduce the Scale-discretised, directional wavelet ILC or SILC, which extends the wavelet ILC framework by localising the calculation of ILC weights in an additional domain. We use wavelets which are not only harmonically-localised but also directional (McEwen et al., 2013; McEwen et al., 2015b; Wiaux et al., 2008). Unlike needlets, which are axisymmetric on the sphere, directional wavelets are non-axisymmetric, *i. e.*, the kernels are "squeezed." This means that for one wavelet scale, one axisymmetric kernel is replaced by a number of complementary directional kernels, each with a different orientation. When these directional wavelets are convolved with a signal on the sphere, (within each scale) different orientations of signal structure are separated.

Table 2.1: The harmonic band-limits $[\ell^j_{\rm min}, \ell^j_{\rm max}]$ of the directional wavelets used in this work. $\ell^j_{\rm peak}$ is the multipole at which each wavelet has its maximum response. The final column shows the number of equiangular samples per wavelet coefficient map $N^j_{\rm samp}$.

Wavelet scale j	$\ell^j_{\rm min}$	$\ell^j_{\rm peak}$	$\ell^j_{\rm max}$	$N^j_{\rm samp}$
Scal.	0	64	64	8,385
0	32	64	128	33,153
1	64	128	256	131,841
2	128	256	512	525,825
3	256	512	706	998,991
4	542	705	918	1,688,203
5	705	917	1193	2,850,078
6	917	1192	1551	4,815,856
7	1192	1550	2015	8,126,496
8	1550	2015	2540	12,910,821
9	2116	2539	3048	18,589,753
10	2539	3047	3600	25,930,801
11	3047	3600	3600	25,930,801

This directional localisation allows the ILC weights to be additionally fine-tuned to better remove foreground and noise, in particular for signals with filamentary structure. Furthermore, directional wavelets exhibit exact reconstruction, allowing them to be embedded in an ILC such that no signal is lost.

SILC is being developed with the goal of analysing CMB polarisation components through an extension to spin, directional wavelets (Leistedt et al., 2015; McEwen et al., 2014; McEwen et al., 2015b), which are expected to be well-suited to localising the complex filamentary morphologies of polarised foregrounds. As a precursor step, in this work we test SILC on the scalar temperature field in order to demonstrate the quantitative consistency of its foreground cleaning performance compared with existing component separation methods, and to identify possible optimisations for the extension to spin fields.

Directional wavelets are explained briefly in § 2.3. In § 2.4, the SILC algorithm is explained in detail. Various sources of error in the method are considered in § 2.5. In § 2.6, we compare our method to previous component-separation methods. The application to *Planck* simulations (§ 2.7) is followed by application to *Planck* data (§ 2.8). We discuss the results and error estimation based on the data in § 2.9 and conclude in § 2.10.

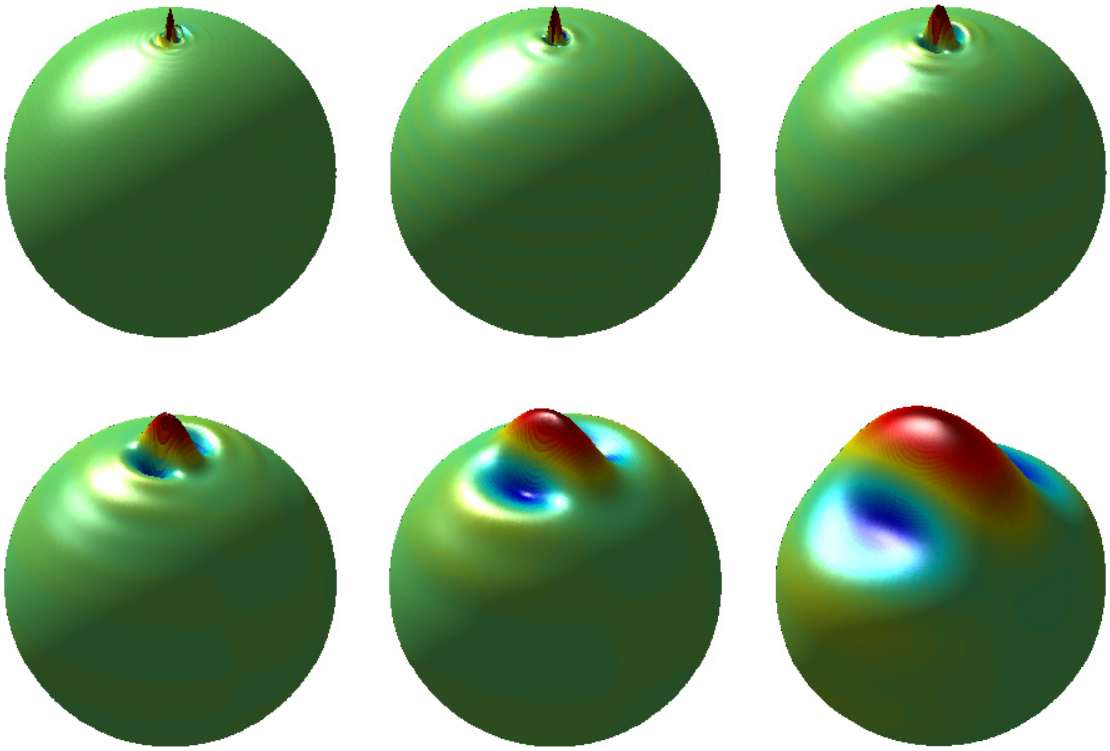

Figure 2.1: The spatial localisation on the sphere of directional, scale-discretised wavelets. Each sub-plot shows a representation of a directional wavelet kernel at different scales, where red, raised parts show positive wavelet response and blue, depressed parts show negative wavelet response. *From left to right, top to bottom*: wavelet scale index j decreases. The number of directions per wavelet scale $N = 3$. Therefore, for complete reconstruction at each scale, the above wavelets would be complemented by two more wavelets of the same size but of a different orientation on the sphere. This figure is adapted from McEwen et al. (2013).

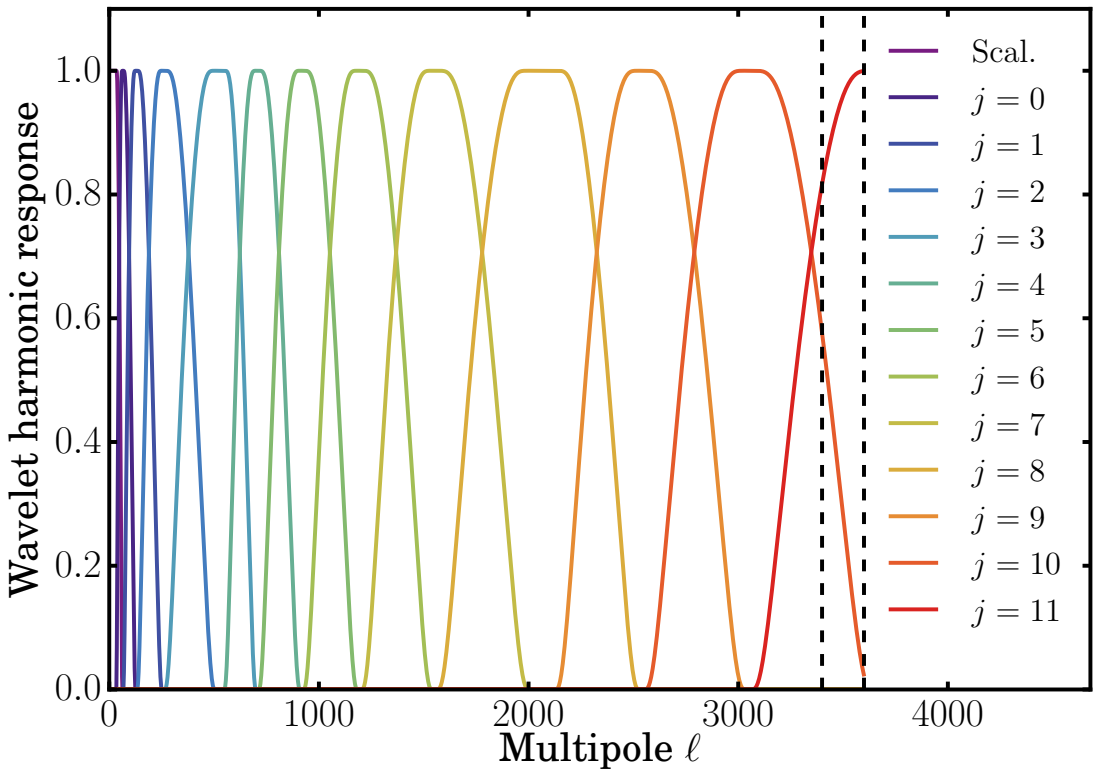

Figure 2.2: The harmonic response of the directional wavelets used in this work, where j specifies the wavelet scale. Increasing j corresponds to a smaller wavelet kernel and so a multipole range on smaller scales (*i. e.*, larger multipoles ℓ). The largest wavelet scale (Scal.) is the scaling function (§ 2.4.4). The two smallest wavelets are harmonically truncated at $\ell = 3600$ but are smoothly tapered to zero from $\ell = 3400$ to $\ell = 3600$ (the two dotted lines) by the beam tapering discussed in § 2.4.3. The band-limits of the above wavelets are given in Table 2.1.

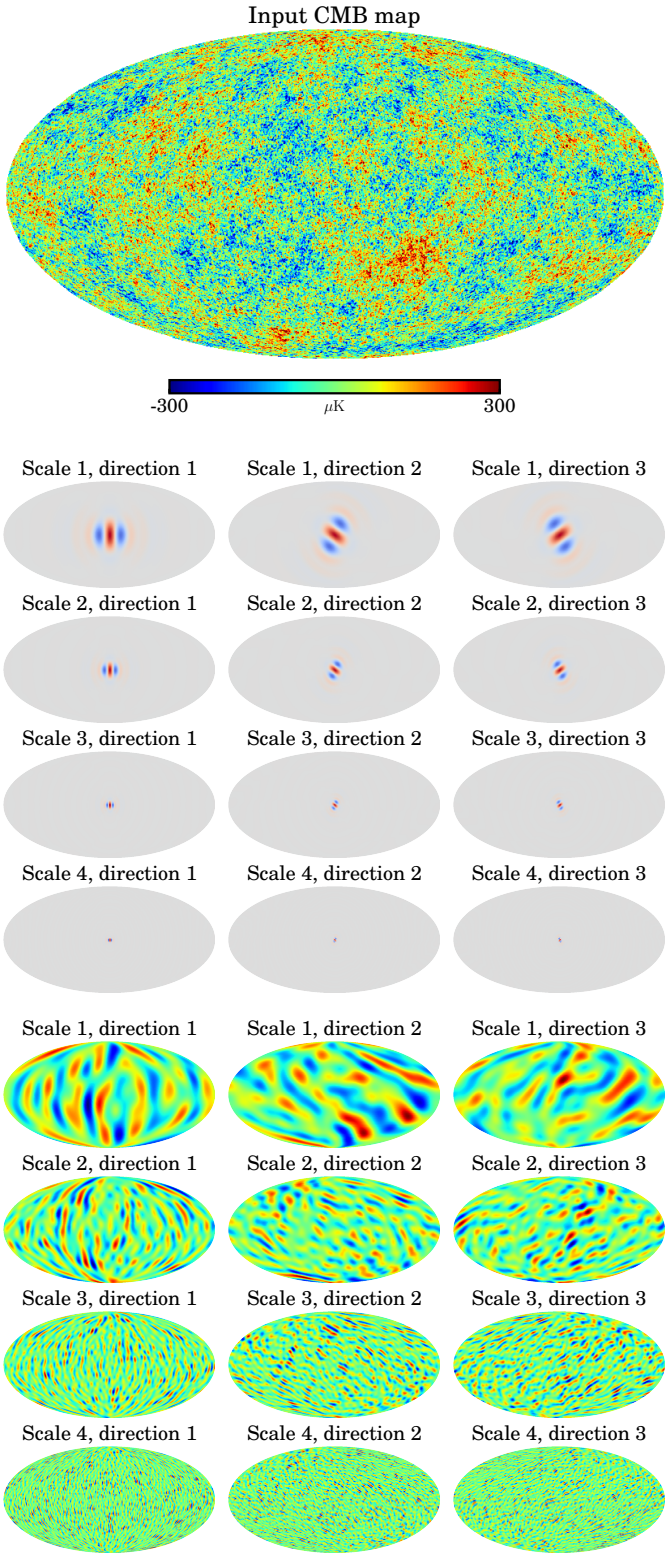

Figure 2.3: The CMB (*top map*) decomposed into directional wavelet coefficient maps (*bottom section*). The wavelet kernels are shown (*middle section*), where red indicates positive response and blue indicates negative. In the full analysis, we also include smaller wavelets than we show above.

2.3 Directional wavelets

Directional, scale-discretised wavelets on the sphere that support exact reconstruction have been developed in Wiaux et al. (2008), McEwen et al. (2013) and McEwen et al. (2015b), while their localisation properties have been studied in McEwen et al. (2015a). Figure 2.1 shows an example of the spatial localisation of directional wavelets. Larger wavelet scales have larger kernels, and when these are convolved with signals defined on the sphere (such as the CMB and astrophysical foregrounds), signal structure with the same scale and orientation as the wavelet is isolated. The kernels in Fig. 2.1 are shown for a single direction and (for complete reconstruction) would be complemented by two more sets of kernels of the same sizes but rotated to different orientations. Figure 2.2 shows an example of the harmonic localisation of directional wavelets (for the wavelets used in this work). The harmonic supports of the wavelets overlap, with each wavelet covering a finite set of multipoles. Figure 2.3 shows an example of directional wavelet decomposition as applied to the CMB. Although the CMB anisotropies are statistically Gaussian, the CMB spots on the sky demonstrate anisotropy as a function of scale (Bond and Efstathiou, 1987). When the CMB is convolved with directional wavelets, structure of different orientations is separated. This further supports the use of directional wavelets in CMB analysis: both filamentary foreground structure and the CMB itself are better localised. This particularly applies in the case of polarisation, as will be discussed in § 2.10 when we consider extensions to our method. For a mathematical description of directional wavelets, see § 2.4.4.

In the spherical harmonic transforms used in the computation of directional wavelet coefficient maps, we adopt the sampling scheme on the sphere of McEwen and Wiaux (2011) (hereafter MW sampling), rather than, *e. g.*, HEALPix sampling (Górski et al., 2005), although in principle HEALPix could be used if desired. The corresponding sampling theorem of McEwen and Wiaux (2011) shows that the MW sampling scheme requires fewer samples for a band-limited signal than any other sampling theorem. Additionally, the use of a separation of variables and fast Fourier transforms (FFTs) yields a numerically efficient algorithm. In particular, our spherical harmonic transforms are theoretically exact, unlike HEALPix. This allows one to manipulate signals with the minimal number of samples and to perform the numerous spherical harmonic transforms involved in the ILC algorithm without any loss of information (other than that due to the finite representation of floating point numbers). Our final map products, however, are provided in HEALPix format. Finally, MW sampling of spin signals on the sphere requires no additional computational complexity and this will be vital in the extension of our method to polarisation E and B modes (§ 2.10). Further

details on MW sampling are given in § 2.4.4.

2.4 Method

We start by outlining the SILC algorithm. The steps are explained in more detail in the subsequent subsections (§ 2.4.1 to 2.4.6). We discuss our numerical implementation in § 2.4.7.

(1). The raw input data are multifrequency full-sky maps of CMB temperature fluctuations. These maps use the HEALPix format. (See § 2.4.1.) The model we employ for the raw data is explained in § 2.4.2.

(2). The maps are "pre-processed" by inpainting in a small point source mask (see § 2.4.6).

(3). The input maps are converted to thermodynamic (CMB) temperature (if necessary). For *Planck* temperature data, the 545 GHz and 857 GHz maps are converted from spectral flux density per unit solid angle (MJy Sr^{-1}) to CMB temperature ($\mathrm{K_{CMB}}$) by the unit conversions given in the *Planck* 2015 Release Explanatory Supplement[1].

(4). The maps are each convolved to have the same effective beam (see § 2.4.3).

(5). Each input map is converted into a set of wavelet coefficient maps. This separates both the scale and orientation of structure within each map. These wavelet coefficient maps use MW sampling (McEwen and Wiaux, 2011). (See § 2.4.4.)

(6). The ILC method is then applied separately to each wavelet scale and orientation. For each scale and orientation, the multifrequency wavelet coefficient maps are weighted and added to form a single wavelet coefficient map that contains mainly CMB signal, as well as some residual foreground and noise. These weights are allowed to vary from wavelet coefficient to wavelet coefficient. The calculation of these weights is explained in § 2.4.5.

(7). The final ILC wavelet coefficient maps are synthesised to form the final product: a full-sky map of CMB temperature fluctuations (with some residual foreground and noise). The final map uses the HEALPix format. (See § 2.4.4.)

(8). The final map is inpainted in a small point source mask (see § 2.4.6).

[1] *Planck* 2015 Release Explanatory Supplement: UC CC tables (`http://wiki.cosmos.esa.int/planckpla2015/index.php/UC_CC_Tables`). For the 545 GHz map, the unit conversion is (58.0356 ± 0.0278) MJy Sr^{-1} $\mathrm{K_{CMB}^{-1}}$ and for the 857 GHz map, the unit conversion is (2.2681 ± 0.0270) MJy Sr^{-1} $\mathrm{K_{CMB}^{-1}}$.

2.4.1 Input data

Our main CMB temperature map products use full-mission 2015 release *Planck* temperature maps as their input[2]. All nine frequency channels are used. At 70 GHz, we use the higher-resolution version at $N_{\rm side} = 2048$. We also use the full-mission Full Focal Plane 8 (FFP8) simulations (Planck Collaboration et al., 2015h) without bandpass mismatch[3].

2.4.2 Data model

Each full-sky temperature map can be modelled (*e. g.*, Basak and Delabrouille, 2012) as

$$T^{{\rm OBS},c}(\hat{\mathbf{n}}) = \int_{\hat{\mathbf{n}}'} {\rm d}\hat{\mathbf{n}}' B^c(\hat{\mathbf{n}}, \hat{\mathbf{n}}') T^{{\rm SIG},c}(\hat{\mathbf{n}}') + T^{{\rm N},c}(\hat{\mathbf{n}}), \tag{2.1}$$

where the signal component can further be decomposed as

$$T^{{\rm SIG},c}(\hat{\mathbf{n}}) = a^c T^{\rm CMB}(\hat{\mathbf{n}}) + T^{{\rm FG},c}(\hat{\mathbf{n}}). \tag{2.2}$$

$T^{\rm CMB}(\hat{\mathbf{n}})$ is the CMB component at a point on the sky $\hat{\mathbf{n}}$. $T^{{\rm FG},c}(\hat{\mathbf{n}})$ and $T^{{\rm N},c}(\hat{\mathbf{n}})$ are respectively the foreground and detector noise components for frequency channel c. a^c is the calibration coefficient for the CMB for each channel. The overall signal component is smoothed by a beam function $B^c(\hat{\mathbf{n}}, \hat{\mathbf{n}}')$ due to the finite resolution of the observations. However, the noise component is not smoothed by the beam. Here we assume the beam to be circularly symmetric. Therefore, the beam can be represented as a sum over Legendre polynomials,

$$B^c(\hat{\mathbf{n}}, \hat{\mathbf{n}}') = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} B_\ell^c P_\ell(\hat{\mathbf{n}}.\hat{\mathbf{n}}'). \tag{2.3}$$

We can recast Eq. (2.1) in the spherical harmonic representation as

$$a_{\ell m}^{{\rm OBS},c} = a^c B_\ell^c a_{\ell m}^{\rm CMB} + B_\ell^c a_{\ell m}^{{\rm FG},c} + a_{\ell m}^{{\rm N},c} \tag{2.4}$$

where $a_{\ell m}$ are the coefficients of spherical harmonics $Y_{\ell m}(\hat{\mathbf{n}})$.

[2] `http://pla.esac.esa.int/pla`

[3] FFP8 simulations are also available with bandpass mismatch, accounting for differences in the bandpasses of detectors nominally at the same frequency, leading to spurious signals in the frequency maps.

2.4.3 Beam convolution

Equation (2.4) shows that each frequency channel c has a different beam transfer function B_ℓ^c. To replace each beam with a channel-independent resolution, we perform a deconvolution/convolution procedure to give spherical harmonic coefficients

$$a_{\ell m}^c = \frac{B_\ell^{\rm EFF}}{B_\ell^c} a_{\ell m}^{{\rm OBS},c}, \tag{2.5}$$

where $B_\ell^{\rm EFF}$ is the final (effective) beam transfer function of our map products. For *Planck* data, we use a Gaussian beam with a FWHM of 5' as our input beam. We taper this beam to zero from $\ell = 3400$ to $\ell = 3600$ using a Fermi function. This suppresses any small-scale power aliasing due to having harmonically-truncated wavelets in this multipole range. Convolving with beam transfer functions ignores the non-axisymmetric component of the beams; these will remain in the input maps but are assumed to be small.

This deconvolution/convolution procedure does not correctly handle the noise component of our input maps. Equation (2.5) can be expanded (using Eq. (2.4)) as

$$a_{\ell m}^c = B_\ell^{\rm EFF}(a^c a_{\ell m}^{\rm CMB} + a_{\ell m}^{{\rm FG},c}) + \frac{B_\ell^{\rm EFF}}{B_\ell^c} a_{\ell m}^{{\rm N},c}, \tag{2.6}$$

where a^c are the CMB calibration coefficients (not to be confused with the inverse spherical harmonic transform of harmonic coefficients a_{lm}^c). The final resolution of an ILC map is usually chosen to match the best resolution of the input maps. Therefore, for all but the highest resolution channel and for all ℓ, $B_\ell^{\rm EFF} > B_\ell^c$. This has the effect of increasing the noise contribution of the input maps, particularly at high ℓ and for low-resolution maps, where $B_\ell^{\rm EFF} \gg B_\ell^c$. We use the *Planck* beam transfer functions as provided in the Reduced Instrument Model (RIMO)[4]. For the LFI beams, we use Gaussian approximations with FWHM 32.33', 27.01' and 13.25' for 30, 44 and 70 GHz respectively. Following Planck Collaboration et al. (2014c), the deconvolved beams are thresholded such that the B_l^c is set to the value given in the RIMO or 0.001, whichever is larger. This prevents the last term in Eq. (2.6) from becoming so large that numerical errors occur. Although we lose accuracy in the deconvolution process, the contribution of the channels in the multipole ranges affected is highly attenuated in the ILC weights in any case.

[4] *Planck* 2015 Release Explanatory Supplement: The 2015 instrument model (`http://wiki.cosmos.esa.int/planckpla2015/index.php/The_RIMO`).

2.4.4 Wavelet analysis and synthesis

The wavelet ILC method requires the decomposition of each band-limited temperature map $T^c(\hat{\mathbf{n}})$ into a set of wavelet coefficient maps W^{Ψ^j}: in our case, directional, scale-discretised wavelets (McEwen et al., 2013; McEwen et al., 2015b; Wiaux et al., 2008). A general introduction was provided in § 2.3 — here we provide some technical details of the implementation. We drop the c superscript on T for the rest of this subsection since each map is analysed using the same wavelets. The wavelet coefficients are defined as the directional convolution of T with wavelets defined on the sphere $\Psi^j \in \mathrm{L}^2(\mathbb{S}^2)$ (specifically those shown in Fig. 2.2) where index j denotes the wavelet scale. Importantly, directional wavelets yield coefficients $W^{\Psi^j}(\hat{\rho})$ that live on the space of three-dimensional rotations, *i. e.*, the rotation group SO(3):

$$W^{\Psi^j}(\hat{\rho}) \equiv \langle T, \mathcal{R}_{\hat{\rho}} \Psi^j \rangle = \int_{\mathbb{S}^2} \mathrm{d}\hat{\mathbf{n}}\, T(\hat{\mathbf{n}}) (\mathcal{R}_{\hat{\rho}} \Psi^j)^*(\hat{\mathbf{n}}), \tag{2.7}$$

where $\mathrm{d}\hat{\mathbf{n}}$ is the usual invariant measure on the sphere, $\cdot^*$ denotes complex conjugation and the rotation operator is defined by

$$(\mathcal{R}_{\hat{\rho}} \Psi^j)(\hat{\mathbf{n}}) \equiv \Psi^j(\mathbf{R}_{\hat{\rho}}^{-1} \hat{\mathbf{n}}), \tag{2.8}$$

where $\mathbf{R}_{\hat{\rho}}$ is the three-dimensional rotation matrix corresponding to $\mathcal{R}_{\hat{\rho}}$. In these equations, $\hat{\rho} = (\theta, \phi, \chi) \in \mathrm{SO}(3)$ denotes the Euler angles (in the zyz convention) with colatitude $\theta \in [0, \pi]$, longitude $\phi \in [0, 2\pi)$ and direction $\chi \in [0, 2\pi)$[5]. In other words, the wavelet coefficients probe directional structure in T with χ corresponding to the orientation about each point (θ, ϕ) on the sphere.

Following the directional construction of scale-discretised wavelets (McEwen et al., 2013; McEwen et al., 2015b; Wiaux et al., 2008), wavelets are defined by their spherical harmonic coefficients in factorised form:

$$\Psi^j_{\ell n} \equiv \kappa^j(\ell) s_{\ell n}, \tag{2.9}$$

where $\kappa^j(\ell)$ sets the harmonic localisation (Fig. 2.2) and $s_{\ell n}$ sets the directional localisation.

In the original definition of scale-discretised wavelets, the size of all harmonic kernels (setting the harmonic localisation of the wavelets) is parameterised by a unique wavelet dilation parameter $\lambda \in \mathbb{R}^+_*$, $\lambda > 1$. Similarly, the number of directions is set by a unique azimuthal band-limit N. These two parameters respectively characterise $\kappa^j(\ell)$ and $s_{\ell n}$ for all j. In this work, we vary

[5]We adopt the zyz Euler convention corresponding to the rotation of a physical body in a *fixed* coordinate system about the z, y and z axes by χ, θ and ϕ, respectively.

λ as a function of multipole in order to allow more flexible harmonic localisation. We achieve this by defining different values of λ in different multipole regions and then stitching together harmonically-truncated wavelets at the region boundaries. We use the values $\lambda = 2, 1.3, 1.2$ with transitions at the multipoles $\ell = 512, 2015$. If at a transition multipole the harmonic peak of the larger wavelet doesn't equal the peak of the smaller wavelet, then a small amount of unit response is used so that the two wavelets can be continuously combined. Wavelets constructed in this manner satisfy the standard admissibility criterion required for exact reconstruction. The harmonic tiling of the resulting wavelets is shown in Fig. 2.2. The technical details of the construction of each kernel is described in McEwen et al. (2015a). Finally, we use a single parameter N for all scales, *i. e.*, each wavelet is divided into the same number of directions. However, a possible extension of this work is to vary N as a function of scale j, *e. g.*, by using curvelet kernels (Chan et al., 2015) or other directional optimisations.

In the case of a single parameter λ, the limits of the wavelet harmonic window for scale j are simply $(\ell^j_{\rm min}, \ell^j_{\rm max}) = (\lambda^{j-1}, \lambda^{j+1})$, with their peak response at λ^j. In our hybrid scheme, this property remains but j and λ must be adjusted in each harmonic region. The full details of our tiling are given in Table 2.1. When the limits of the harmonic windows of the maximum wavelet scales extend beyond the overall band-limit $\ell_{\rm max}$, the windows are truncated at $\ell_{\rm max}$. Finally, note that a scaling function W^Φ is needed to capture the very low frequency content of the signal. It is axisymmetric and the corresponding scaling coefficients therefore live on the sphere. Here we do not give the full details of the construction of the scaling function or the factors $\kappa^j(\ell)$ and $s_{\ell n}$ since these can be straightforwardly reproduced by following previous approaches (McEwen et al., 2013; McEwen et al., 2015b; Wiaux et al., 2008) and using Table 2.1.

To apply the ILC algorithm, the above continuous wavelet coefficients must be discretised. Since they live on the rotation group SO(3), we represent them using the sampling scheme of McEwen et al. (2015), which is a generalisation of the MW sampling scheme (McEwen and Wiaux, 2011). Because our wavelets have well-defined band-limits, this approach allows a multiresolution scheme where each scale is pixellated with a minimal number of samples. In practice, the j-th wavelet scale has a band-limit $\ell^j_{\rm max}$ and is only evaluated at locations $(\theta^j_t, \phi^j_p, \chi_n)$ with $t \in \{0, 1, \ldots, \ell^j_{\rm max}\}$, $p \in \{0, 1, \ldots, 2\ell^j_{\rm max}\}$ and $n \in \{0, 1, \ldots, N-1\}$. Although wavelet coefficients are evaluated at discrete samples only, for a band-limited signal they capture the total information content of the underlying continuous wavelet coefficient representation, probed up to harmonic band-limit $\ell^j_{\rm max}$ and azimuthal band-limit N. This is thanks to the sampling theory on the rotation group $\mathrm{SO}(3)$ of McEwen et al. (2015). In the full ensemble of realisations, the ILC (see § 2.4.5 for details) has no

sensitivity to the choice of coordinate convention for directions χ. In a single realisation, there will be some marginal sensitivity to this choice manifesting in the localisation of the empirical covariances we use. However, this effect is sub-dominant to the choice of N, on which we concentrate our analysis.

After the ILC method (see § 2.4.5) has been applied to the sets of wavelet coefficient maps, there is one final map $W^{\Psi^j,\mathrm{ILC}}(\hat{\rho})$, for each wavelet scale j, living on SO(3) and including the multiple orientations $\chi_0, \ldots, \chi_{N-1}$. The additional axisymmetric scaling coefficients $W^{\Phi,\mathrm{ILC}}(\hat{\mathbf{n}})$ live on the sphere. The final temperature map $T^{\mathrm{ILC}}(\hat{\mathbf{n}})$ is synthesised by

$$\begin{aligned} T^{\mathrm{ILC}}(\hat{\mathbf{n}}) = &\int_{\mathbb{S}^2} \mathrm{d}\hat{\mathbf{n}}' W^{\Phi,\mathrm{ILC}}(\hat{\mathbf{n}}')(\mathcal{R}_{\hat{\mathbf{n}}'}\Phi)(\hat{\mathbf{n}}) \\ &+ \sum_{j=j_{\min}}^{j_{\max}} \int_{\mathrm{SO}(3)} \mathrm{d}\hat{\rho}\; W^{\Psi^j,\mathrm{ILC}}(\hat{\rho})(\mathcal{R}_{\hat{\rho}}\Psi^j)(\hat{\mathbf{n}}), \end{aligned} \tag{2.10}$$

where $\mathrm{d}\hat{\rho}$ is the usual invariant measure on the rotation group. This final ILC temperature map is pixellated using the HEALPix format from its spherical harmonic coefficients $T^{\mathrm{ILC}}_{\ell m}$.

The wavelet analysis and synthesis are performed using the latest version of the S2LET[6] code (Leistedt et al., 2013; McEwen et al., 2015), which in turn relies on the SSHT[7] (McEwen and Wiaux, 2011) and S03[8] (McEwen et al., 2015) codes to compute spin spherical harmonics and Wigner transforms exactly and efficiently using the MW sampling scheme. Thanks to the sampling theorem, the wavelet coefficients can be transformed using Wigner transforms without any loss of information (McEwen et al., 2015).

2.4.5 ILC method

Following the wavelet analysis of the input maps (see § 2.4.4), there is a wavelet coefficient map W^c_{jnk} for each channel c, scale j and orientation n with a pixel index k. Using this more compact notation, we conflate the scaling coefficient map with the wavelet coefficient maps as the ILC method applies in exactly the same fashion. The ILC estimate of the CMB signal at each wavelet scale and orientation is defined as a weighted sum of the wavelet coefficient maps at that scale and orientation

$$W^{\mathrm{ILC}}_{jnk} \equiv \sum_{c=1}^{N_c} \omega^c_{jnk} W^c_{jnk}\,, \tag{2.11}$$

[6] http://www.s2let.org
[7] http://www.spinsht.org
[8] http://www.sothree.org

where ω^c_{jnk} are the weights (which are allowed to vary across the scale and orientation of the signal as well as pixel space) and N_c is the number of input channels.

We impose a constraint on the weights (to ensure that the CMB signal is preserved) such that

$$\sum_{c=1}^{N_c} a^c \omega^c_{jnk} = 1. \tag{2.12}$$

Assuming that the CMB and foregrounds and the CMB and noise are respectively uncorrelated, the variance of the error in the result is minimised when the variance of the ILC map itself is minimised. The resulting weights are given by

$$\omega^c_{jnk} = \frac{\sum_{c'=1}^{N_c} (R^{-1}_{jnk})^{cc'} a^{c'}}{\sum_{c=1}^{N_c} \sum_{c'=1}^{N_c} a^c (R^{-1}_{jnk})^{cc'} a^{c'}}, \tag{2.13}$$

where the true covariance matrices at scale j, orientation n and pixel k, $(R_{jnk})^{cc'} = \langle W^c_{jnk} W^{c'}_{jnk} \rangle$ (where the angled brackets indicate an ensemble average). For a derivation of Eq. (2.13), see Eriksen et al. (2004); Tegmark et al. (2003).

In this work, we estimate covariance matrices empirically by the following procedure (as used in Basak and Delabrouille 2012; Planck Collaboration et al. 2015c). We start by calculating at each pixel k:

$$(R^{\rm approx}_{jnk})^{cc'} = W^c_{jnk} W^{c'}_{jnk}. \tag{2.14}$$

We then smooth each element of the above matrix by a Gaussian beam $w_j(k, k')$ in pixel space to form the empirical estimates of covariance matrices

$$(\hat{R}_{jnk})^{cc'} = \sum_{k'=1}^{N^j_{\rm samp}} w_j(k, k') (R^{\rm approx}_{jnk'})^{cc'} \tag{2.15}$$

where $N^j_{\rm samp}$ is the total number of pixels in a given map at scale j. For computational efficiency, we perform this smoothing in harmonic space:

$$(\hat{R}_{jnk})^{cc'} = \sum_{\ell=0}^{2\ell^j_{\rm max}} \sum_{m=-\ell}^{\ell} w^\ell_j (r^{\ell m}_{jn})^{cc'} Y^{\ell m}_k, \tag{2.16}$$

where $(r^{\ell m}_{jn})^{cc'}$ are the harmonic coefficients of the maps formed by the elements of matrices $(R^{\rm approx}_{jnk})^{cc'}$, w^ℓ_j is a Gaussian beam transfer function and $Y^{\ell m}_k$ are the spherical harmonics evaluated at pixel k.

The size of the Gaussian kernel used to smooth the covariance matrices is chosen to be proportional to the size of the wavelet used to form a particular set of wavelet coefficient maps[9]. In general the estimation of covariance matrices in ILC methods could be further optimised. It may be preferable to dynamically adapt the smoothing kernel used based on local data. Delabrouille et al. (2009) suggested using a larger kernel at high Galactic latitudes where Galactic emission does not vary so much and a smaller kernel towards the Galactic equator where emission is more complex. It could involve masking equatorial regions when estimating the covariance at higher latitudes (somewhat akin to Planck Collaboration et al. (2014c)). It could involve convolving the maps of elements of covariance matrices with the same directional wavelet in order to pick out how the local covariance follows the directional structure of the underlying signal. As mentioned above, in this work, we use a similar method as in previous work for ease of comparison.

It is also worth discussing the upper limit on the summation over ℓ in Eq. (2.16). We first note the general rule that for the product of two spherical harmonics (Driscoll and Healy, 1994)

$$Y_{\ell_1,m_1}(\hat{\mathbf{n}})Y_{\ell_2,m_2}(\hat{\mathbf{n}}) = \sum_{L=|\ell_1-\ell_2|}^{\ell_1+\ell_2} a_{L,m_1+m_2} Y_{L,m_1+m_2}(\hat{\mathbf{n}}) \tag{2.17}$$

where $Y_{L,m_1+m_2}(\hat{\mathbf{n}})$ is defined to be zero if $|m_1 + m_2| > L$. It follows that for the product of two band-limited maps $M(\hat{\mathbf{n}}) = \sum_{\ell=\ell_1}^{\ell_2} \sum_{m=-\ell}^{\ell} m_{\ell m} Y_{\ell m}(\hat{\mathbf{n}})$ and $N(\hat{\mathbf{n}}) = \sum_{\ell=\ell_3}^{\ell_4} \sum_{m=-\ell}^{\ell} n_{\ell m} Y_{\ell m}(\hat{\mathbf{n}})$ (where, without loss of generality, $\ell_1 \leq \ell_4$):

$$M(\hat{\mathbf{n}})N(\hat{\mathbf{n}}) = \sum_{L=0}^{\ell_2+\ell_4} \sum_{M=-L}^{L} p_{LM} Y_{LM}(\hat{\mathbf{n}}) \tag{2.18}$$

for $\ell_3 < \ell_2$, *i. e.*, the limits on ℓ in the two maps overlap (the p_{LM} are the new harmonic coefficients). (If the limits do not overlap, the lower limit on L in Eq. (2.18) becomes $\ell_3 - \ell_2$.) The limits on ℓ in the wavelet coefficient maps W^c_{jnk} are $(\ell^j_{\rm min}, \ell^j_{\rm max})$ (see § 2.4.4). Therefore, by Eq. (2.18) and Eq. (2.15), the limits on ℓ in the covariance matrix element maps $(\hat{R}_{jnk})^{cc'}$ are $(0, 2\ell^j_{\rm max})$; hence the limits on ℓ in Eq. (2.16).

Having established the main equations governing the ILC method, we now present the main steps in the ILC algorithm that we use:

1. Form the $(R^{\rm approx}_{jnk})^{cc'}$ by Eq. (2.14).

2. Smooth the $(R^{\rm approx}_{jnk})^{cc'}$ in harmonic space by Eq. (2.16).

[9] $\mathrm{FWHM}^j = 50\sqrt{\frac{1200}{N^j_{\rm samp}}}$. This value is the same as used in the NILC implementation on *Planck* data.

3. Take the inverse of each covariance matrix at each pixel to form $(\hat{R}^{-1}_{jnk})^{cc'}$.

4. Calculate the ILC weights ω^c_{jnk} by Eq. (2.13), where we assume that $a^c = 1$ for all c and we substitute the empirical estimates for the inverse covariance matrices.

5. Finally, calculate the ILC estimate wavelet coefficient maps $W^{\rm ILC}_{jnk}$ by applying Eq. (2.11).

2.4.6 Point source masking

The input frequency maps are diffusively inpainted in a small point source mask following the method employed by Planck Collaboration et al. (2015g). This recognises that the ILC fails when the CMB is obscured by bright extragalactic sources or complex emission near the Galactic equator. The inpainting removes these sources and attempts to replace them with an extrapolation of the surrounding signal. The mask supplied is taken from the NILC section of Planck Collaboration et al. (2014a) and is constructed from the Planck Catalogue of Compact Sources (PCCS) (Planck Collaboration et al., 2014b, 2015f)[10]. It masks about 2.2% of the whole sky, predominantly along the Galactic equator towards the Galactic centre.

Because of this inpainting, the final ILC map is inpainted within the point source mask. For the purposes of this inpainting, we have split the mask into two, based on the size of its constituent individual contiguous holes[11]. For holes consisting of less than or equal to 800 pixels, we inpaint with a constrained Gaussian realisation following the method of Benoit-Lévy et al. (2013), itself an approximate implementation of the Hoffman-Ribak algorithm (Hoffman and Ribak, 1991). For holes consisting of more than 800 pixels (the largest 131 out of 10031), we inpaint with a standard diffusive algorithm (in particular, following the method employed by Planck Collaboration et al. (2015g)). The result is that the ILC map is 1.3% Gaussian inpainted and 0.9% diffusively inpainted. This follows Benoit-Lévy et al. (2013), who do not recommend using their Gaussian inpainting for large holes near the Galactic equator.

2.4.7 Numerical implementation

SILC is implemented in `Python` and is parallelised. At full *Planck* resolution ($N_{\rm side} = 2048$, $\ell_{\rm max} = 3600$), when run on a 60-core symmetric multiprocessor (SMP) with 1.5 TB RAM and a 24-core cluster node with 256 GB RAM[12], the pipeline takes approximately 12 hours per direction. As shown

[10] The details of its construction are given in Planck Collaboration et al. (2014a). It can be downloaded from `http://pla.esac.esa.int/pla` and is labelled `I_MASK` in the NILC data products.

[11] Note that many holes can be large and irregularly-shaped due to the overlapping of smaller circular holes.

[12] The exact specification for our infrastructure is an Intel Xeon E7-4890 2.8 GHz SMP with 4 × 15-core CPUs with 25.6 GB RAM per core, and an Intel Xeon E5-2697 2.7 GHz node with 2 × 12-core CPUs with 10.7 GB RAM per core.

by Eq. (2.16), we perform spherical harmonic transforms to $2\ell_{\rm max} = 7200$. For a given number of directions N, the full pipeline takes approximately N times as long as the axisymmetric limit of our method (when $N = 1$). In our infrastructure, the code was usually memory-limited (due to the very high resolution of the covariance matrix maps $(R_{jnk})^{cc'}$ at harmonic band-limit $2\ell^j_{\rm max}$); the amount of parallelisation sometimes had to be reduced to prevent memory overloads on a single node. As mentioned in § 2.4.4, the wavelet transforms employ the latest version of `S2LET` (Leistedt et al., 2013; McEwen et al., 2015), written in `C` with `Python` wrappers, itself employing `SSHT` (McEwen and Wiaux, 2011) and `SO3` (McEwen et al., 2015). Despite the use of MW sampling and FFTs, spherical harmonic transforms are the most time-consuming part of the pipeline, again due to the very high resolution of the $(R_{jnk})^{cc'}$ (for the smallest wavelets, these covariance maps are band-limited at $\ell = 7200$). There is scope to further optimise the implementation. Our wavelet analysis and synthesis functions do not respectively output and take as input wavelet coefficient maps at double-resolution (*i. e.*, a map band-limited at $\ell^{\rm max}_j$ sampled at $2\ell^{\rm max}_j$), requiring additional spherical harmonic transforms to double the resolution. Also, our spherical harmonic transform function does not calculate harmonic coefficients to a multipole less than the band-limit of the input map (*i. e.*, to only calculate $a_{\ell m}$ for $\ell < L$ where $L < \ell^j_{\rm max}$), resulting in excess computation at certain steps in the algorithm. These optimisations are left as further work.

2.5 Sources of error in the ILC

By the linearity of the wavelet transform in Eq. (2.7), the data model in Eqs. (2.1) and (2.2) can be recast in wavelet space as

$$W^c_{jnk} = a^c W^{\rm CMB}_{jnk} + W^{{\rm FG},c}_{jnk} + W^{{\rm N},c}_{jnk}\,, \tag{2.19}$$

where $W^{\rm CMB}_{jnk}$, $W^{{\rm FG},c}_{jnk}$ and $W^{{\rm N},c}_{jnk}$ are respectively the CMB, foreground and instrumental noise contributions to each wavelet coefficient map. The beams within each component have been absorbed into the component wavelet coefficient maps. Substituting Eq. (2.19) into Eq. (2.11) gives

$$\begin{aligned} W^{\rm ILC}_{jnk} &= \sum_{c=1}^{N_c} a^c \omega^c_{jnk} W^{\rm CMB}_{jnk} + \sum_{c=1}^{N_c} \omega^c_{jnk} (W^{{\rm FG},c}_{jnk} + W^{{\rm N},c}_{jnk}) \\ &= W^{\rm CMB}_{jnk} + \frac{\sum_{c,c'=1}^{N_c} (W^{{\rm FG},c}_{jnk} + W^{{\rm N},c}_{jnk}) (R^{-1}_{jnk})^{cc'} a^{c'}}{\sum_{c,c'=1}^{N_c} a^c (R^{-1}_{jnk})^{cc'} a^{c'}}\,, \end{aligned} \tag{2.20}$$

where the second equality follows by applying the constraint given in Eq. (2.12) and expanding the weights as given in Eq. (2.13). Even when the calibration a^c and the covariance matrices $(R_{jnk})^{cc'}$

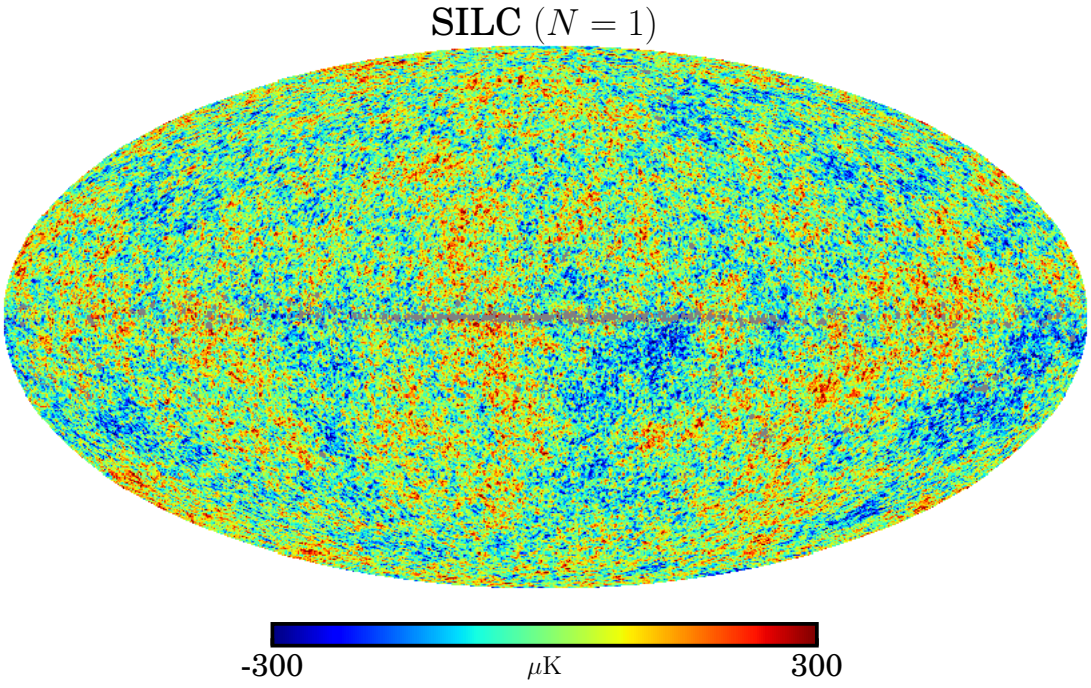

Figure 2.4: Planck *data*. The CMB temperature anisotropies reconstructed using SILC in the axisymmetric limit $(N = 1, \mathrm{FWHM} = 5', N_{\mathrm{side}} = 2048)$. The grey pixels are the point source mask.

are correct, there is always residual signal in the final ILC wavelet coefficient maps, given by the second term on the RHS of Eq. (2.20). Due to the linearity of the inverse wavelet transforms, this residual signal will propagate linearly into the final ILC temperature map as calculated by Eq. (2.10). As explained in § 2.4.5, this error term is reduced by minimising the empirical variance of the ILC map assuming that the CMB and foregrounds and the CMB and noise are respectively uncorrelated.

There are additional sources of error in the ILC method. The first is due to inaccuracy in the calculation of covariance matrices $(R_{jnk})^{cc'}$, *i. e.*, deviations in the empirical estimate $(\hat{R}_{jnk})^{cc'}$ from the true covariance $(R_{jnk})^{cc'}$. Delabrouille et al. (2009) estimated the first order expansion of the reconstruction error in the ILC map estimate due to this covariance error. They showed that the covariance of the ILC error with the CMB is inversely proportional to the number of "effective modes" used in the ILC calculation. This covariance bias is negative. In our directional wavelet decomposition, our "effective modes" are spherical harmonic coefficients weighted to take account of the fact that the harmonic responses of wavelets overlap in both scale and direction. As N, the number of orientations probed, increases and so does the number of wavelets, each wavelet coefficient map contains fewer "effective modes." We therefore conclude that our directional wavelet ILC may be susceptible to this negative ILC bias by increasing N. Delabrouille et al.

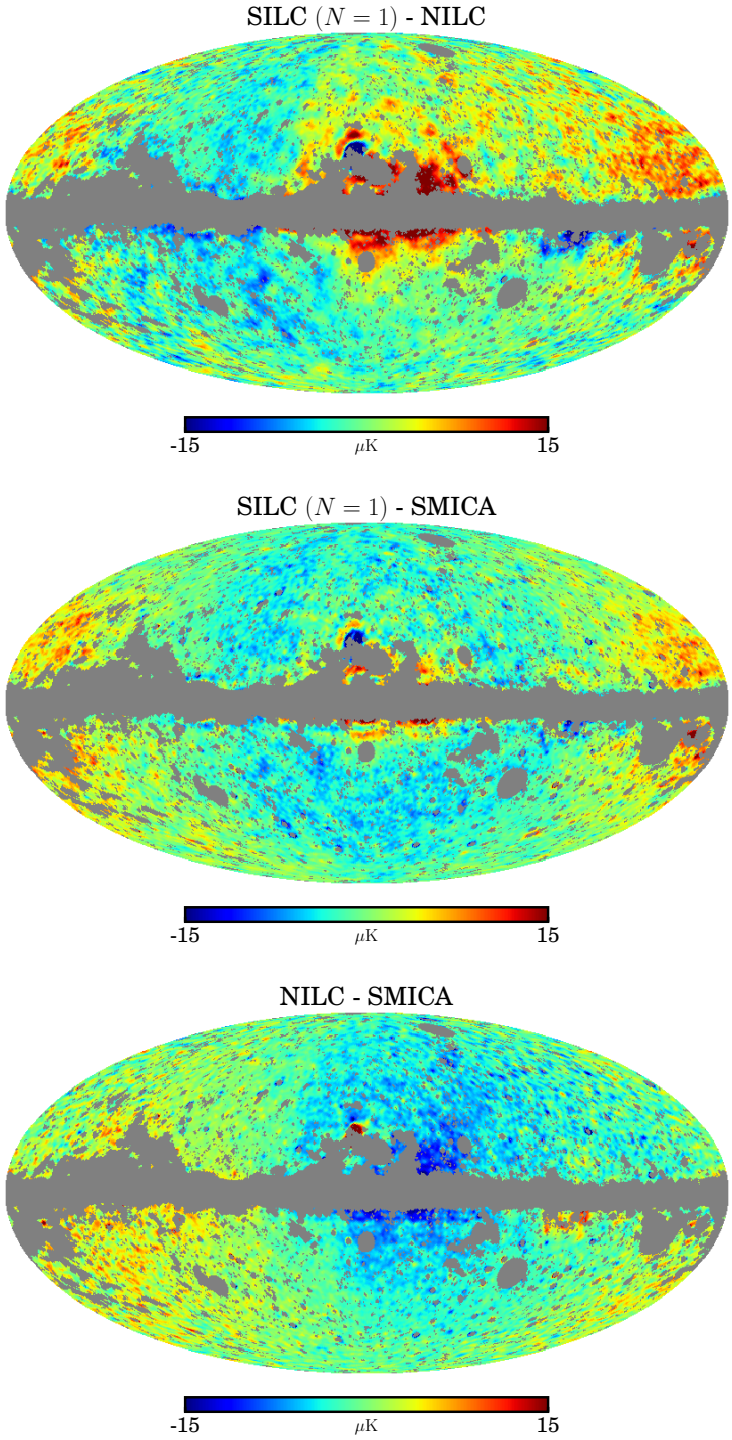

Figure 2.5: Planck *data.* Differences between the axisymmetric limit ($N = 1$) of SILC, NILC and SMICA. The maps have been smoothed to $\mathrm{FWHM} = 80'$ and downgraded to $N_{\mathrm{side}} = 128$. The grey pixels are the UT78 confidence mask from Planck Collaboration et al. (2015c), which masks the regions of the NILC and SMICA maps not recommended for cosmological analysis. The differences are (*from top to bottom*) (a) SILC ($N = 1$) - NILC, (b) SILC ($N = 1$) - SMICA and (c) NILC - SMICA.

(2009) also showed that due to chance correlations between the CMB and foregrounds, the variance minimisation leads to the unintentional cancellation of $N^c - 1$ CMB modes. For *Planck*, $N^c = 9$, whereas for *WMAP*, $N^c = 5$. We therefore expect the magnitude of this negative bias to double simply by using more input frequency channels. Also, since this covariance bias is due to the cancellation of CMB modes, Delabrouille et al. (2009) showed that the absolute value is proportional to the CMB power. Therefore, the absolute value of the bias is greatest on large scales where CMB power is concentrated. In general, these biases are best estimated through suites of Monte Carlo simulations.

Another source of error is due to inaccuracy in the calibration a^c of the CMB. Dick et al. (2010) calculated the consequence of a first order error in a^c on a multiplicative correction to the CMB term in Eq. (2.20). They showed that even a small error in calibration can lead to a significant negative multiplicative bias in the CMB term, when the signal-to-noise ratio is large. (Here, the noise in this ratio also includes foreground signal.) They consider the implications for using an ILC on *Planck* data, where the signal-to-noise ratio is larger than for *WMAP* data. They estimate that a 1% error in a^c can cancel about a third of the CMB signal, while even a 0.1% error in a^c can remove about 1% of the CMB. Since our main map products use *Planck* data as input, they will be susceptible to this additional negative calibration bias. In this work, we assume that the CMB is calibrated to have unit response for all frequency channels, *i. e.*, $a^c = 1$ for all c.

As mentioned in § 2.4.2 and 2.4.3, we assume all beams to be circularly symmetric. Therefore, non-axisymmetric beam components will propagate into the ILC calculation but are assumed sufficiently small to be ignored.

2.6 Comparison to previous work

We now consider how SILC compares with existing component separation methods, particularly those adopted for the *Planck* analysis. We applied the axisymmetric limit (when $N = 1$) of SILC to full-mission *Planck* data and compared the results to existing *Planck* analyses using the NILC and SMICA methods: the former because it is the closest in spirit to SILC, and the latter because it is the baseline method adopted by the Planck Collaboration for high-resolution analyses. Figure 2.4 shows the CMB reconstructed by the axisymmetric limit of SILC, while Fig. 2.5 shows the differences between this map and the NILC and SMICA (full-mission 2015 release) CMB maps and the difference between NILC and SMICA. The differences between the three maps are small in magnitude and mostly concentrated at the edges of the Galactic mask towards the Galactic centre,

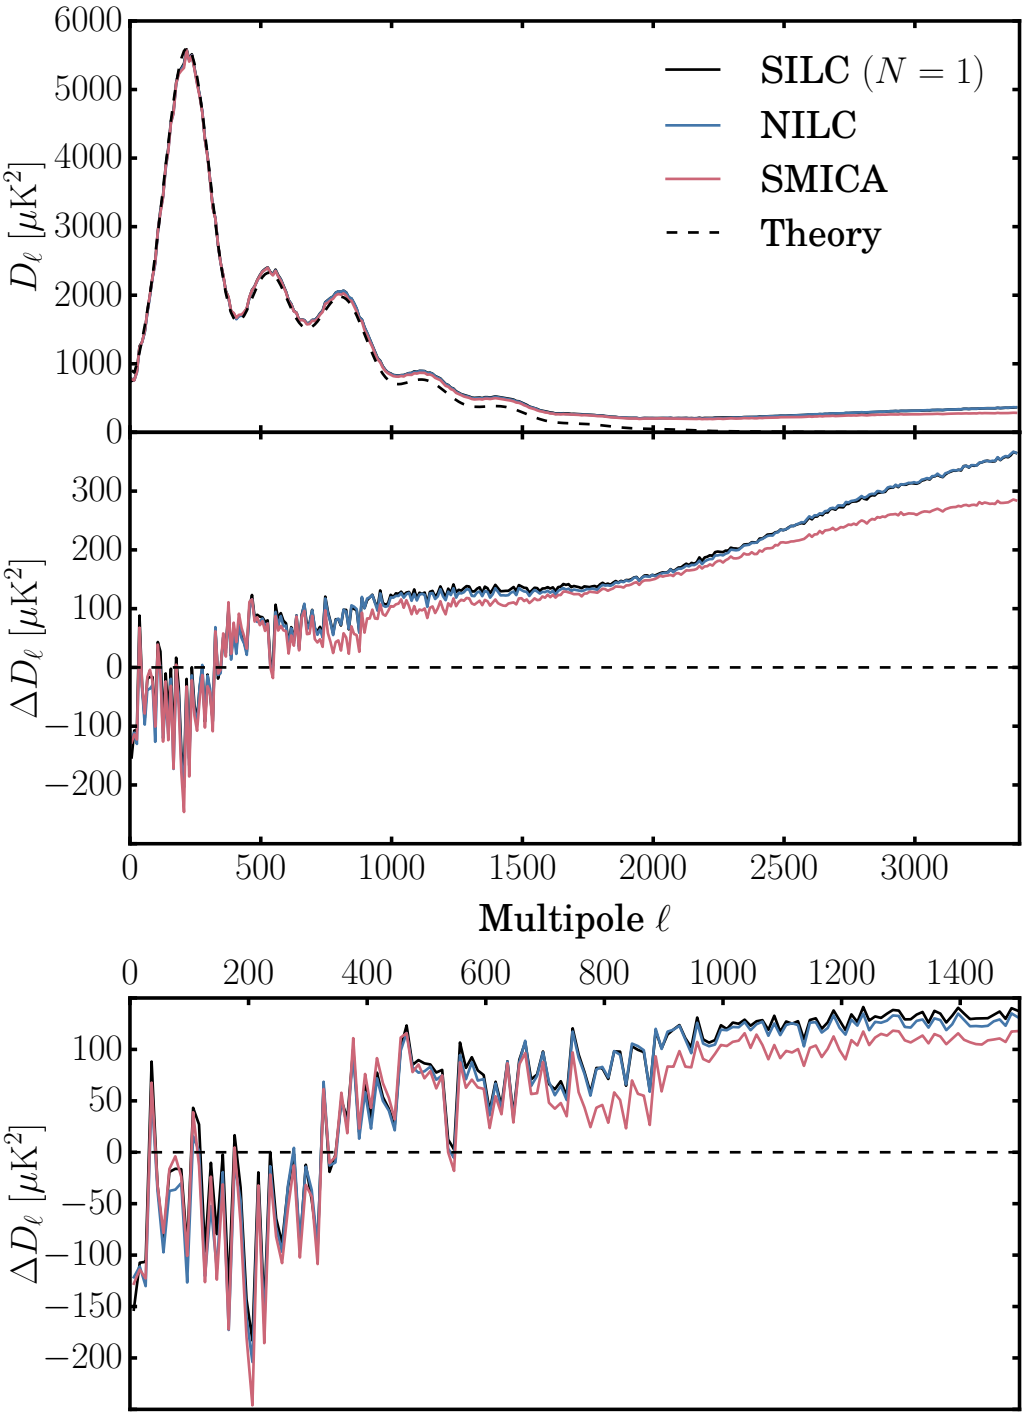

Figure 2.6: Planck *data*. TT angular power spectra comparing the axisymmetric limit ($N = 1$) of SILC to NILC and SMICA. The top panel (a) shows point source masked spectra. The middle panel (b) shows residuals after subtracting the best-fit ΛCDM model from the *Planck* 2015 likelihood. The bottom panel (c) shows the same residuals at low multipoles only ($\ell < 1500$).

where foreground emission is most intense and complex. Quantitatively, we can compare the mean values and standard deviations of the masked difference maps. The mean values of Figs. 2.5 (a), (b) and (c) are respectively 0.44, -0.63, and -1.07 μK, while the standard deviations are respectively 4.24, 3.38 and 3.43 $\mu\mathrm{K}^2$. These values are small and similar, suggesting a strong consistency between the three methods. These difference maps have been formed from maps which have been smoothed and downgraded in resolution and so visually highlight differences at the lowest multipoles.

Figure 2.6 compares point source masked TT angular power spectra ($D_\ell = \ell(\ell+1)C_\ell/2\pi$) at the full multipole range of the three maps (up to $\ell = 3400$)[13] with a CMB spectrum derived from the *Planck* 2015 TT and low TEB likelihood[14]. The SILC spectrum is remarkably similar to that of NILC. This is unsurprising since the axisymmetric limit of SILC (when $N = 1$) is very similar to the NILC method. Nonetheless, there are a number of pipeline differences. In particular, we use a different set of wavelets than the needlets employed in NILC (as discussed in § 2.4.4), even in the axisymmetric limit, with different harmonic responses. Figure 2.2 shows the harmonic response of the wavelets used in this work and Table 2.1 lists their harmonic band-limits ℓ^j_{min} and ℓ^j_{max}. The SMICA spectrum has lower residuals at higher multipoles than both the axisymmetric limit of SILC and NILC.

Figure 2.7 compares full-sky angular power spectra of the three maps, including the inpainted point source regions. The spectra are similar to those in Fig. 2.6. The main difference is the lower noise tail in the SILC map at high multipoles above $\ell = 1500$ (where all component separation CMB maps are dominated by residual instrumental noise). This is because, unlike NILC and SMICA, we do not Gaussian inpaint the very largest point source holes, but rather use diffusive inpainting (as discussed in § 2.4.6). The Gaussian inpainting of large irregular holes is poorly constrained and adds residual noise relative to diffusive inpainting.

We have shown that the axisymmetric limit of SILC gives comparable performance to NILC and SMICA. In § 2.7 and § 2.8, we "turn on" the directionality of the wavelets and consider the impact on CMB reconstruction from simulated and real data respectively.

[13] In order to estimate full-sky spectra from a masked map, we correct the C_ℓ by dividing by $f_{\mathrm{sky}} = 0.978$, a good approximation for a small mask. We elect to use point source masked spectra in order to concentrate our analysis on foreground and noise removal, rather than how maps are inpainted; all three maps are inpainted (at least) within the mask used.

[14] The parameters come from the `base_plikHM_TT_lowTEB` likelihood. The values are available in the *Planck* 2015 Release Explanatory Supplement: 2015 Cosmological parameters and MC chains (`http://wiki.cosmos.esa.int/planckpla2015/images/f/f7/Baseline_params_table_2015_limit68.pdf`).

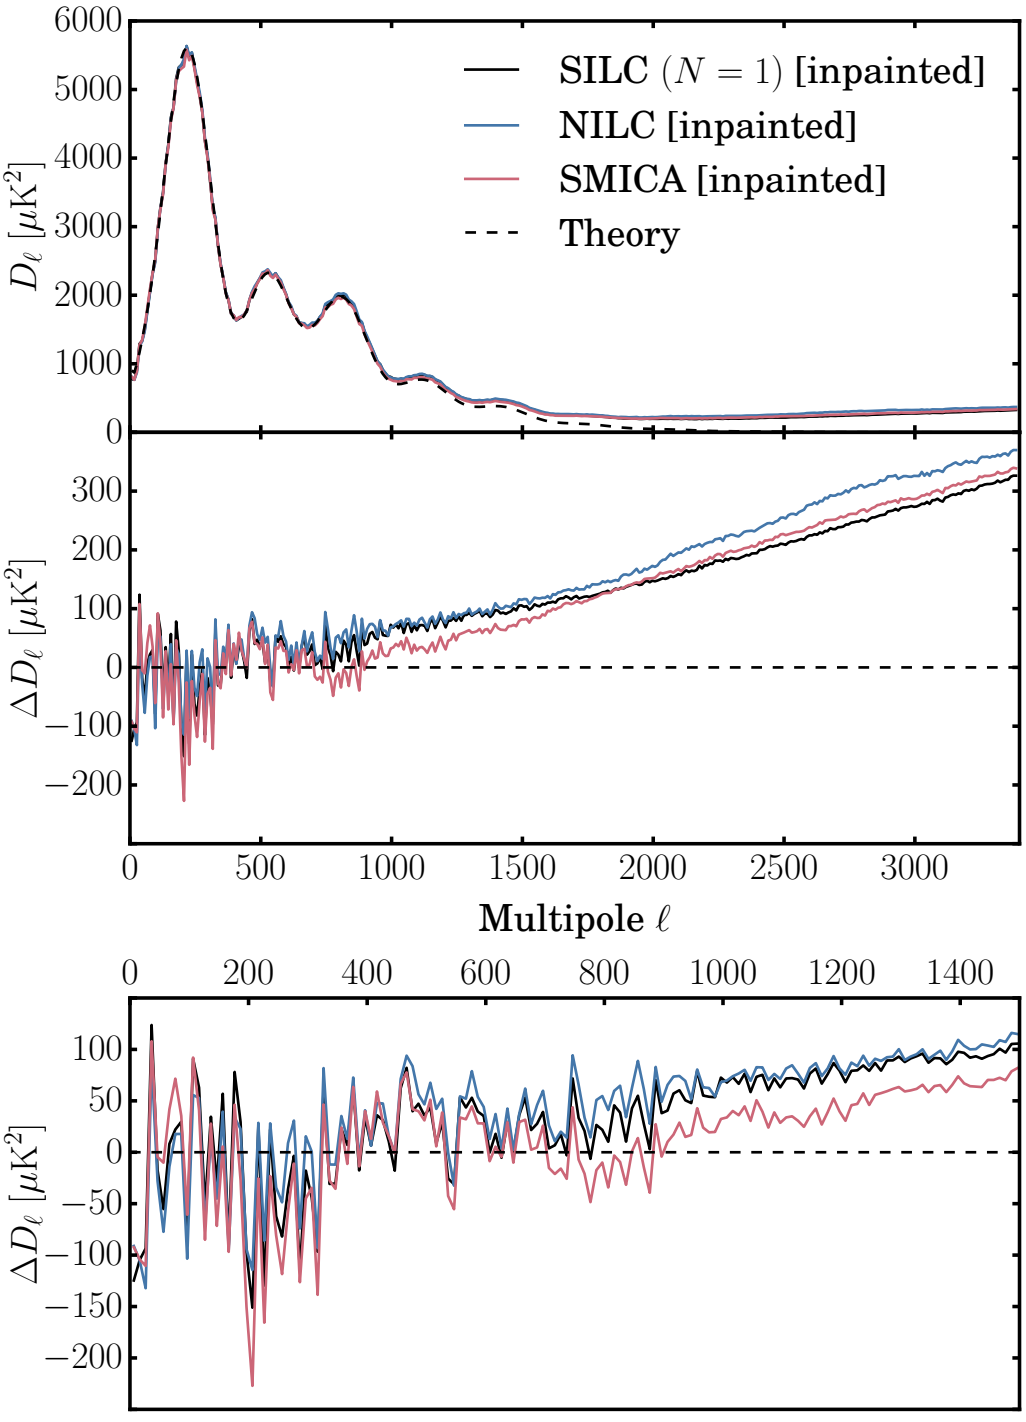

Figure 2.7: Planck *data*. TT angular power spectra comparing the axisymmetric limit ($N = 1$) of SILC to NILC and SMICA. The top panel (a) shows full-sky spectra of inpainted maps. The middle panel (b) shows residuals after subtracting the best-fit ΛCDM model from the *Planck* 2015 likelihood. The bottom panel (c) shows the same residuals at low multipoles only ($\ell < 1500$).

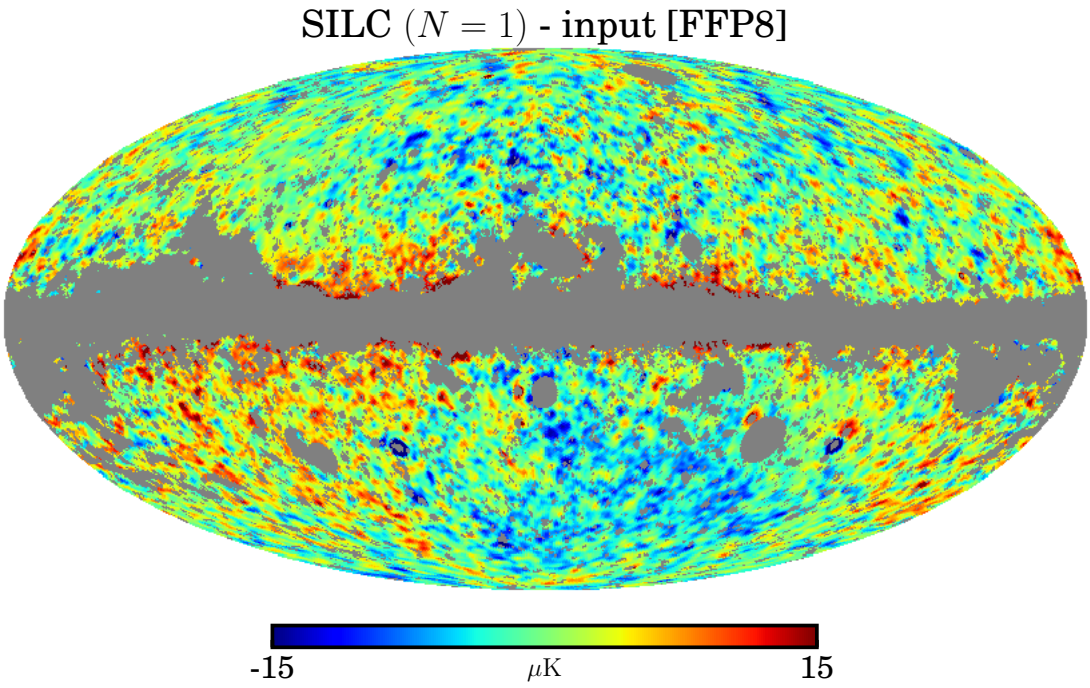

Figure 2.8: Planck *simulations.* Difference between output ILC and input CMB temperature maps from FFP8 simulations. The maps have been smoothed to $\mathrm{FWHM} = 80'$ and downgraded to $N_{\mathrm{side}} = 128$. The grey pixels are the UTA76 confidence mask from Planck Collaboration et al. (2015c), which masks the Galactic region in FFP8 simulations where foreground emission is strongest.

2.7 Application to *Planck* simulations

We now apply SILC to the fiducial full-mission *Planck* FFP8 simulated sky maps, focusing on the impact on CMB reconstruction by increasing directionality as a function of scale. Figure 2.8 shows the difference between our reconstructed CMB (using $N = 1$) and the input simulated CMB. There are small-magnitude differences particularly at the edge of the Galactic mask where the strength and complexity of foreground emission is greatest. As in Fig. 2.5, this difference map is at low resolution and so highlights residuals at the lowest multipoles. Figure 2.9 compares point source masked TT angular power spectra (up to $\ell = 3400$) of CMB maps reconstructed using values of N from 1 to 5. It can be seen that the introduction of directionality has the greatest effect at multipoles around $\ell = 800$; the residuals are beginning to converge for $\ell \gtrsim 2000$. Figure 2.10 shows the differences between simulated CMB maps reconstructed using $N = [2, 3, 4, 5]$ minus the input CMB. The four maps and the axisymmetric difference map in Fig. 2.8 are almost identical with small magnitude residuals. This is because these low-resolution difference maps again highlight residuals on the very largest scales. However, as discussed in § 2.4.4, the wavelets we use are constructed to have an axisymmetric scaling function at the very lowest multipoles. The scaling function we use (as detailed in Table 2.1) means that no directionality is applied for $\ell < 32$.

Figure 2.11 shows equivalent difference maps as in Fig. 2.10 but for the simulated CMB reconstructed using directional wavelets at all scales[15], including for $\ell < 32$. It can be seen that the reconstruction errors are significantly larger in magnitude and cover almost the entire sky. The errors are also dominated by the largest scales, in particular a large error in the quadrupole increasing with magnitude as the amount of directionality N increases. We attribute this effect most probably to the ILC "biases" discussed in § 2.5, in particular the cancellation of CMB modes due to chance correlations with foregrounds in the ILC variance minimisation. Delabrouille et al. (2009) showed that the absolute value of this effect is largest on large scales where CMB power is concentrated, since the cancelled CMB modes on large scales have the greatest magnitude. Further, as discussed in § 2.5 and shown in Fig. 2.11, these errors are expected to increase in magnitude as a function of N. This is because as N increases, each directional wavelet coefficient map (the space in which our ILC operates) contains fewer "effective modes" of the input data and so the error in our empirical covariance estimation is expected to increase. This error propagates to the final maps.

These map reconstruction errors due to the implementation of directionality on the very largest scales are accompanied by increasingly negative power spectrum residuals as N increases, in

[15] In particular, the scaling function and $j = 0$ wavelet are replaced by two directional wavelets with harmonic band-limits $[1, 60]$ and $[1, 128]$.

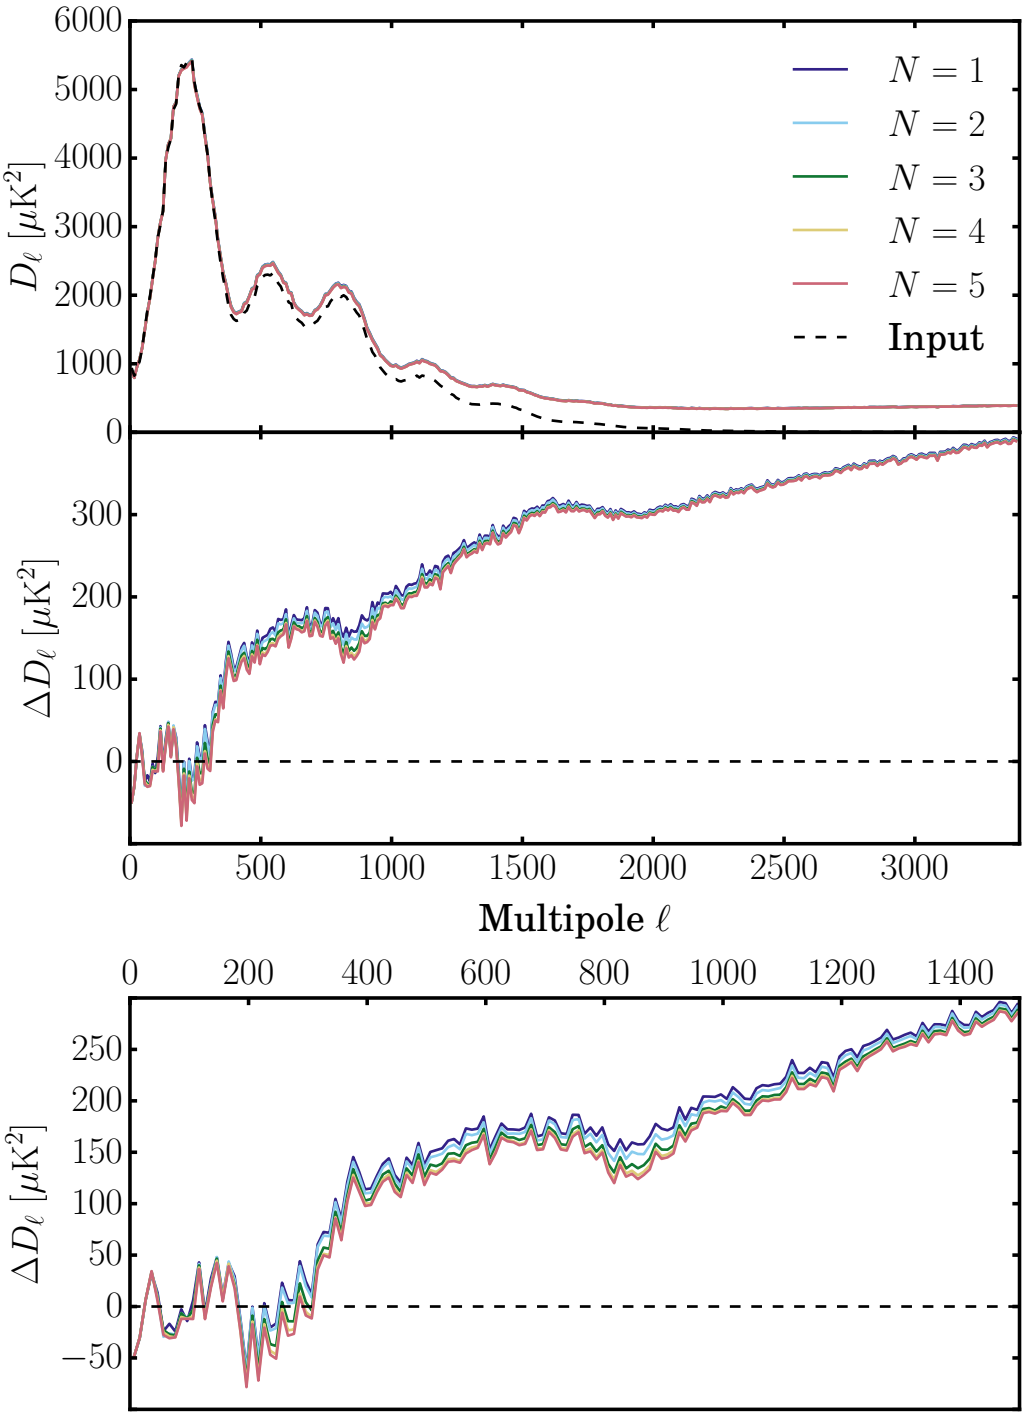

Figure 2.9: Planck *simulations*. TT angular power spectra comparing output ILC using different values of N and input CMB from FFP8 simulations. The top panel (a) shows point source masked spectra. The middle panel (b) shows residuals after subtracting the input CMB spectrum. The bottom panel (c) shows the same residuals at low multipoles only ($\ell < 1500$).

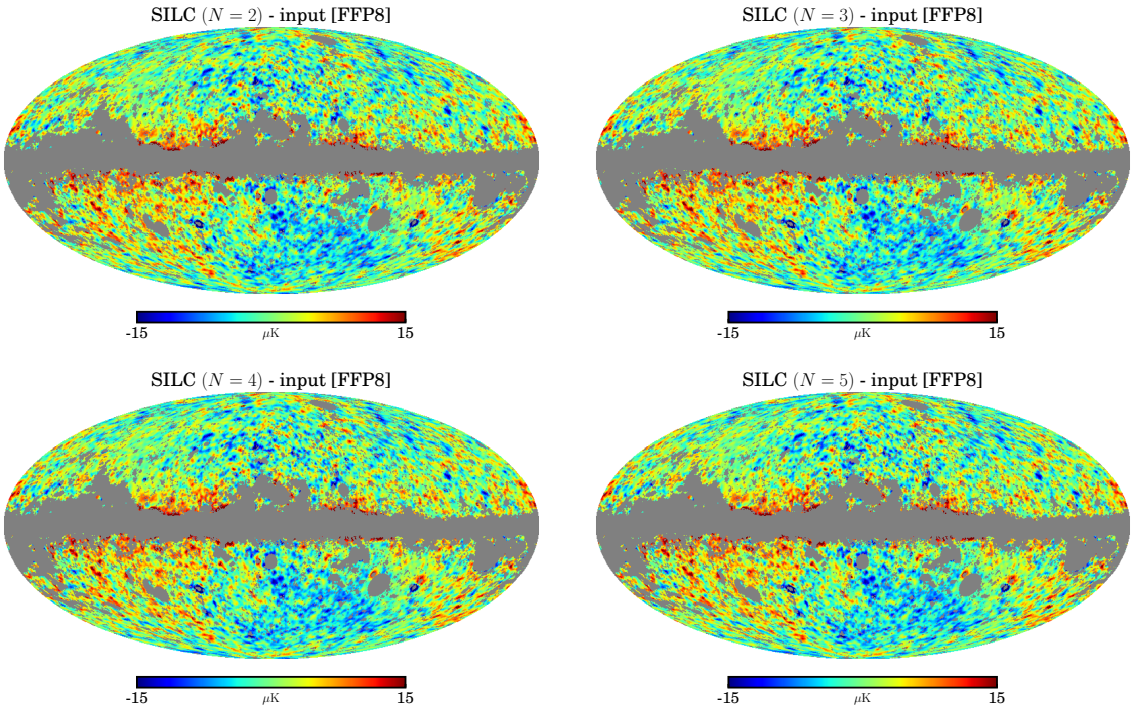

Figure 2.10: Planck *simulations.* Differences between output ILC reconstructed using different values of N and input CMB temperature maps from FFP8 simulations. The maps have been smoothed to $\mathrm{FWHM} = 80'$ and downgraded to $N_{\mathrm{side}} = 128$. The grey pixels are the UTA76 confidence mask. The differences are (*from left to right, top to bottom*) (a) $N = 2$, (b) $N = 3$, (c) $N = 4$, (d) $N = 5$ minus the input CMB.

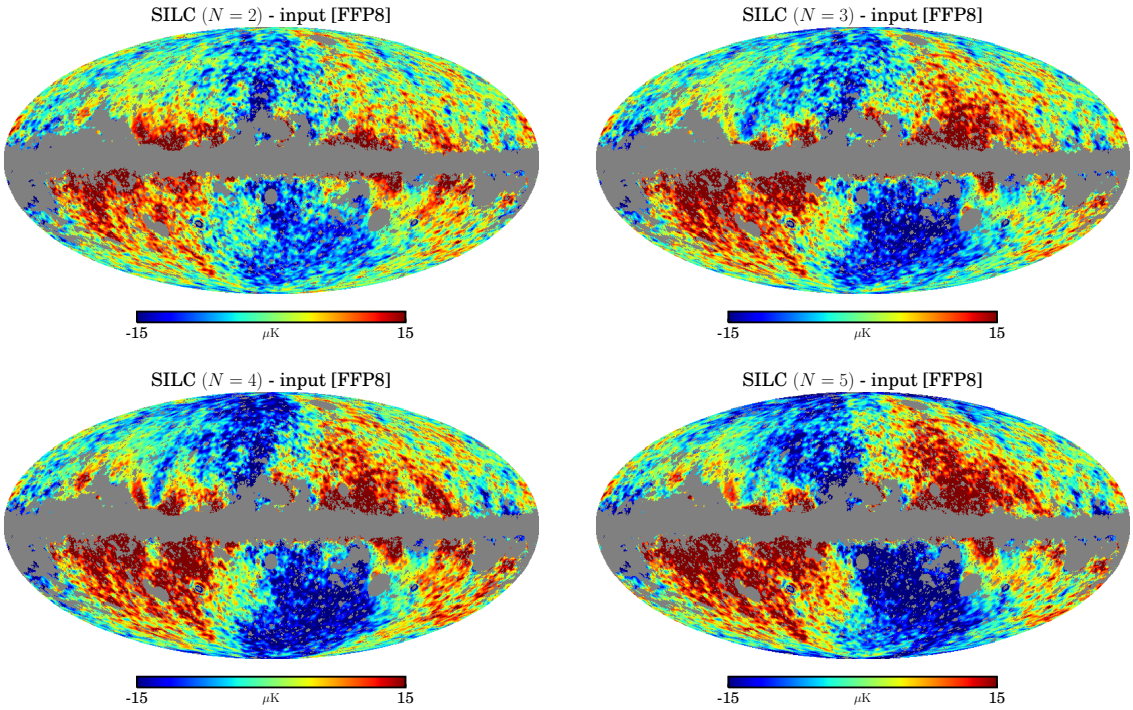

Figure 2.11: Planck *simulations*. Same as Fig. 2.10 (which uses the recommended wavelets) but here using directional wavelets on large scales ($\ell < 32$), which is not recommended as it leads to increased CMB reconstruction errors as seen above.

particular in the first multipole bin from $\ell = 2$ to $\ell = 11$. This is also indicative of the negative ILC bias due to empirical CMB cancellation, as discussed in § 2.5 and Delabrouille et al. (2009). The results in Fig. 2.11 thus motivate the use of an axisymmetric scaling function, which ensures that no directionality is used for $\ell < 32$ and so reduces the errors in CMB reconstruction. In principle these biases can be estimated and corrected through large suites of simulations, which is beyond the scope of this work.

2.8 Application to *Planck* data

We now study the application of SILC with increasing directionality to the full-mission *Planck* sky maps. The left column of Fig. 2.12 shows the full-resolution reconstructed CMB maps as calculated with different values of N from 2 to 5, which visually appear very similar. The right column of Fig. 2.12 shows the differences between the CMB reconstructed using $N = [2, 3, 4, 5]$ minus the axisymmetric limit (when $N = 1$), highlighting the differences at the lowest multipoles. The differences are of largest magnitude towards the Galactic plane where foreground emission is concentrated. This shows how the different wavelet kernels are localising the ILC weights differently in response to the directional structure of the foregrounds and CMB. The differences are small, reflecting the implementation of an axisymmetric scaling function, meaning that no directionality is applied at $\ell < 32$. Figure 2.13 compares point source masked TT angular power spectra of the CMB reconstructed using values of N from 1 to 5. The power spectrum residuals from a *Planck* best-fit ΛCDM model remain small for most scales until the reconstructed spectra reach a characteristic noise spectrum for $\ell \gtrsim 1500$ where the different values of N converge. At these high multipoles, the ILC solution is dominated by residual instrumental noise. We see the biggest impact from directionality at intermediate multipoles (from $\ell = 400$ to $\ell = 1500$). For comparison, we plot the SMICA power spectrum. In further support to the discussion in § 2.6, SILC matches the performance of SMICA. We note that, as with the simulations in § 2.7, directionality changes the reconstructed CMB power spectrum most significantly at intermediate multipoles around $\ell = 800$.

2.9 Discussion

The comparisons in § 2.6 demonstrate that SILC matches the performance of two previous methods, NILC and SMICA, in both maps and power spectra, with particular similarity between the axisymmetric limit of SILC and NILC, as expected. Both map residuals and power spectra in § 2.7

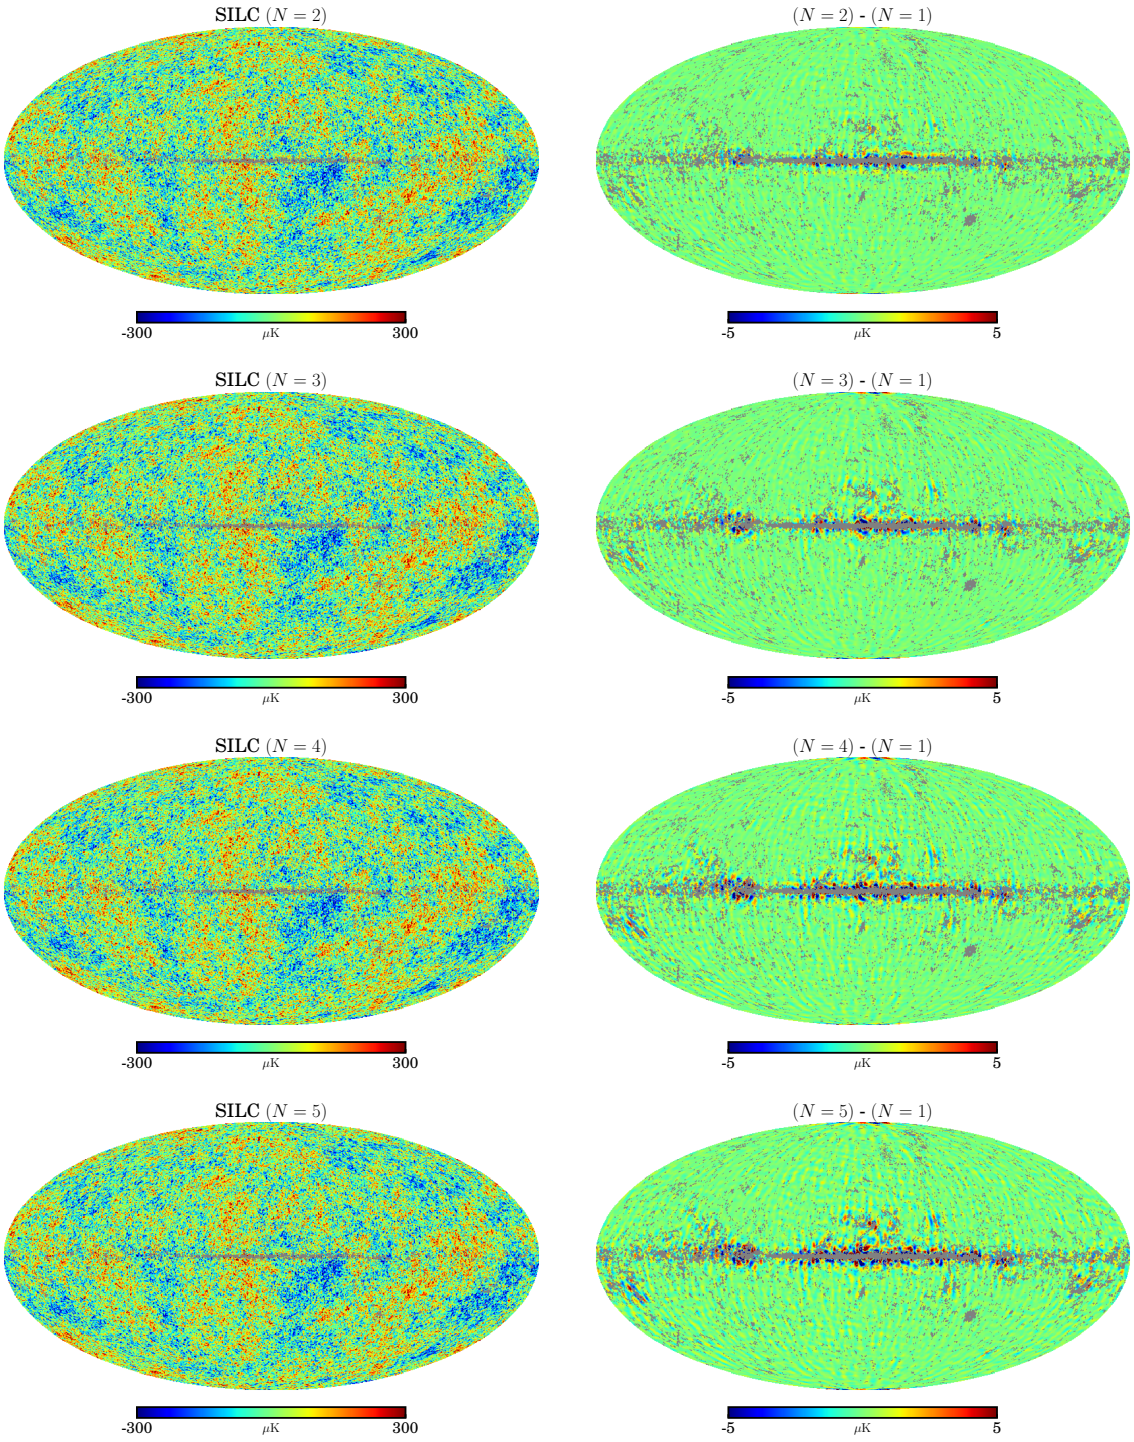

Figure 2.12: Planck *data*. *Left*: CMB temperature anisotropies reconstructed using SILC with different values of N (FWHM = 5', $N_{\rm side} = 2048$). *Right*: differences between CMB temperature maps reconstructed using different values of N minus the axisymmetric limit $N = 1$. The maps have been smoothed to $\mathrm{FWHM} = 80'$ and downgraded to $N_{\rm side} = 128$. *In both columns*: the grey pixels are the point source mask (downgraded in resolution as appropriate). *From top to bottom*: (a) $N = 2$, (b) $N = 3$, (c) $N = 4$, (d) $N = 5$.

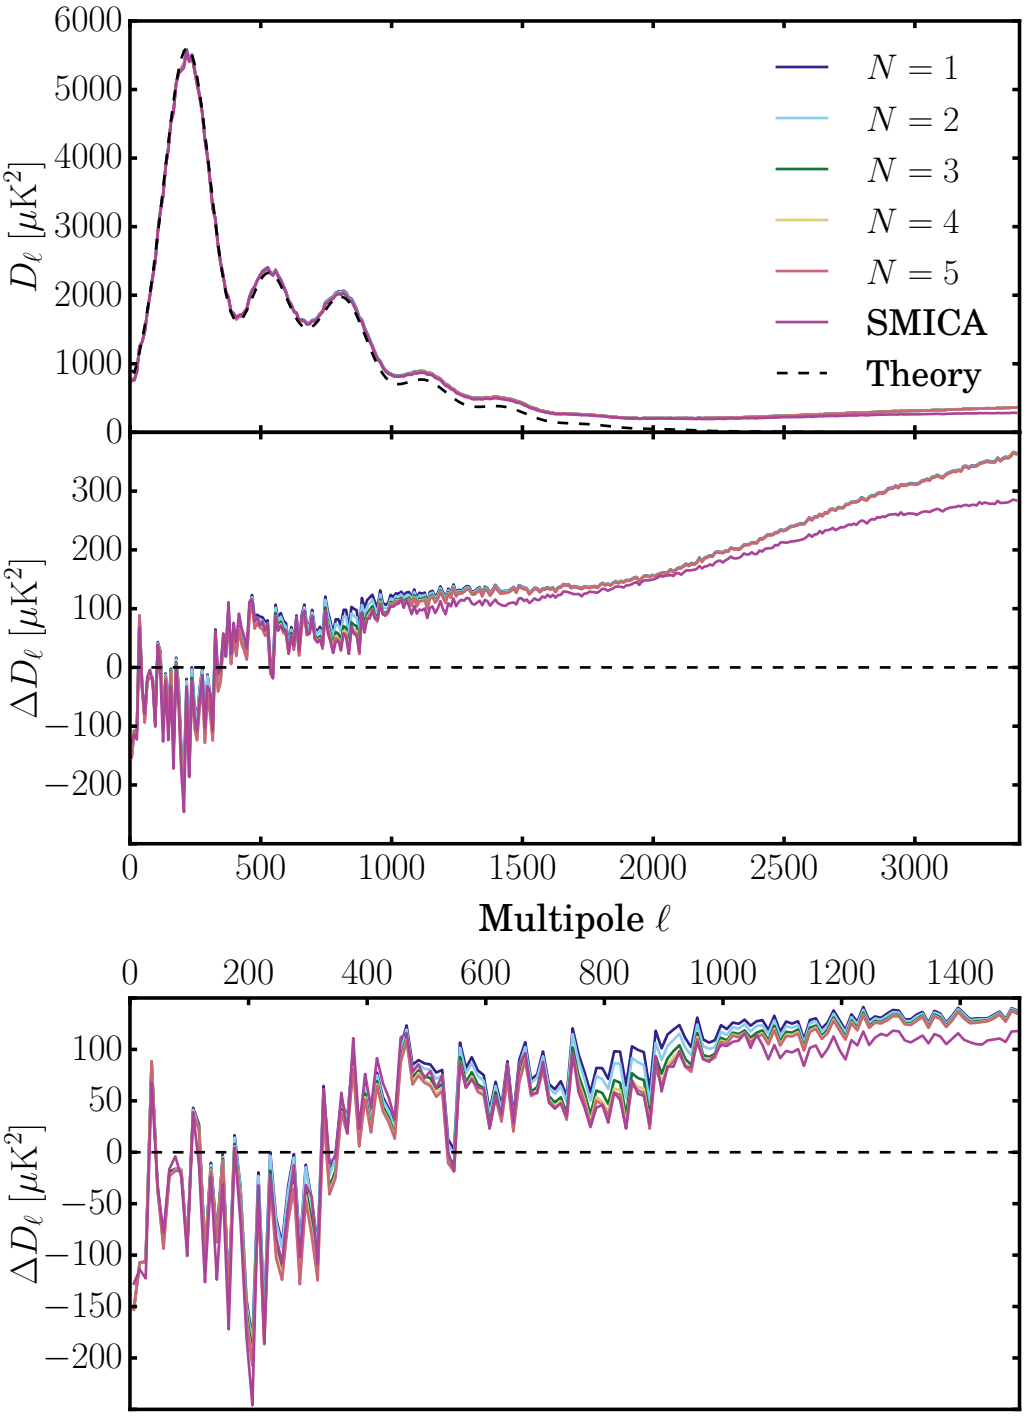

Figure 2.13: Planck *data*. TT angular power spectra comparing different values of N from 1 to 5 and SMICA. The top panel (a) shows point source masked spectra. The middle panel (b) shows residuals after subtracting the best-fit ΛCDM model from the *Planck* 2015 likelihood. The bottom panel (c) shows the same residuals at low multipoles only ($\ell < 1500$).

and § 2.8 show that switching on directionality changes CMB reconstruction most significantly at intermediate multipoles $\ell = 400$ – 1500. There appears to be little benefit in localising the ILC with directional wavelets at the very smallest scales, where the ILC result is noise-limited. We also adopt an axisymmetric scaling function on the very largest scales, meaning that there is no directionality at $\ell < 32$. In Fig. 2.11, we show the large CMB reconstruction errors arising from using directionality on the largest scales. This motivates the use of an axisymmetric scaling function, which significantly reduces the errors as seen in Fig. 2.10. In § 2.7, we sketched out an argument that attributes these errors to empirical CMB cancellation (§ 2.5). However, the precise source and exact magnitude of any ILC errors are best estimated through suites of simulations.

We have presented this analysis by producing CMB maps (in § 2.7 and § 2.8) each with a different single value of N at all wavelet scales. Our method can be simply extended to allow different values of N at each wavelet scale. In the same way that each wavelet scale has different harmonic band-limits, they can also have different azimuthal band-limits, optimised as identified above to reduce foreground and noise residuals.

The negative ILC power spectrum biases discussed in § 2.5 must be quantified in parallel to this directionality optimisation if using the resulting map for power spectrum analyses. It is possible to estimate variance biases in the data through suites of realistic simulations. However, we can also calculate this using the data itself and a fiducial CMB spectrum. In wavelet space, the variance estimator at each wavelet coefficient is $\langle W^{\rm ILC} W^{\rm ILC\dagger}\rangle = \omega^\dagger \langle \mathbf{W}\mathbf{W}^\dagger\rangle \omega = \omega^\dagger \mathbf{R} \omega = (\mathbf{a}^\dagger \mathbf{R}^{-1} \mathbf{a})^{-1} = (\sum_{c,c'=1}^{N_c} (R^{-1})^{cc'})^{-1}$, where each equality follows by applying in turn Eqs. (2.11), (2.14) and (2.13) (from § 2.4.5) and then expanding the vector notation (the vectors span the space defined by the number of input channels; explicitly, we assume unit CMB calibration $\mathbf{a} = (1, 1, 1, 1, 1, 1, 1, 1, 1)$). In order to calculate the variance bias in the ILC, we can subtract the expected CMB variance $\sum_{\ell m} C_\ell |\Psi^j_{\ell m}|^2$. If analysis was done in wavelet space, the above would define the variance bias. If only considering the diagonal terms in wavelet space, it is possible to straightforwardly transform this estimate to real space through an inverse wavelet transform as in Eq. (2.10), substituting $\Phi(\hat{\mathbf{n}})$ and $\Psi^j(\hat{\rho})$ respectively for $|\Phi(\hat{\mathbf{n}})|^2$ and $|\Psi^j(\hat{\rho})|^2$. However, for a full treatment of the variance bias including off-diagonal terms, full wavelet space covariances need to be calculated. Although many off-diagonal terms would decay, this would still be computationally demanding and will be the focus of future research. However, we reiterate that if analysis is carried out in wavelet space, then variance biases can be straightforwardly calculated from the information already contained in the results.

2.10 Conclusions

We have presented SILC, a new form of internal linear combination that uses directional, scale-discretised wavelets to localise the ILC weighting according to the frequency, spatial, harmonic and, for the first time, morphological information in the CMB and its foregrounds. This is motivated by the anisotropic or filmentary morphology of both the CMB and astrophysical foregrounds in the microwave sky. We have tested SILC on 2015-release *Planck* data and simulations, demonstrating comparable performance to two existing component separation algorithms, NILC and SMICA, and investigated how to optimise the use of morphological information through directionality. We have explored increasing the amount of directionality in the algorithm, showing that on the largest and the smallest scales, the axisymmetric limit of the ILC works well, while at intermediate multipoles (from $\ell = 400 - 1500$), increasing N (the number of directions per scale) leads to lower residuals. At high multipoles ($\ell \gtrsim 1500$), the input data are already noise-limited, as is the ILC reconstruction, and directionality does not reduce the reconstruction error, as instrumental noise has no directional structure. We adopt an axisymmetric scaling function to analyse the largest scales, meaning that we use no directionality for $\ell < 32$. This is motivated by the observation that increasing directionality on large scales gives increased reconstruction errors over the axisymmetric limit. We argue that these errors are due to empirical CMB cancellation in the ILC calculation, though the exact source must be estimated through large suites of realistic simulations. Allowing N to vary with wavelet scale is analogous to the choice of different harmonic band-limits at different scales.

We conclude that the introduction of directional wavelets allows greater flexibility in the ILC to make use of morphological information at targeted scales. Our multiprocessing implementation takes advantage of the wavelet scales to allow large-scale results to be analysed while small scales are still being processed. Moreover, our wavelet transforms are quick and exact, using MW sampling and fast Fourier transforms (see § 2.3). As discussed in § 2.5, the ILC is prone to several sources of error and variance bias, including empirical CMB cancellation. This bias can be estimated through suites of Monte Carlo simulations, but we have also outlined (in § 2.9) the ability to estimate biases directly from the data, most straightforwardly in wavelet space. We make our map products available at `http://www.silc-cmb.org`[16].

This work on scalar signals (*i. e.*, the temperature I component of the CMB) can be extended to spin signals (*i. e.*, the polarisation Q and U components of the CMB, or, equivalently, the E and B modes), by using spin wavelets (Leistedt et al., 2015; McEwen et al., 2014; McEwen et al., 2015b).

[16]The DOI for our data release is 10.5281/zenodo.44373.

These are an extension of directional, scale-discretised wavelets to represent spin signals, such as CMB polarisation, a spin ± 2 signal. We expect that the directionality will be particularly suited to the anisotropic, filamentary nature of polarised foregrounds when observed on the sky, and in future work will present the application of SILC to CMB polarisation data.

Acknowledgements

KKR thanks Franz Elsner for valuable discussions. 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.

3

Spin-SILC: CMB polarisation component separation with spin wavelets

3.1 Abstract

We present Spin-SILC, a new foreground component separation method that accurately extracts the cosmic microwave background (CMB) polarisation E and B modes from raw multifrequency Stokes Q and U measurements of the microwave sky. Spin-SILC is an internal linear combination method that uses spin wavelets to analyse the spin-2 polarisation signal $P = Q + iU$. The wavelets are additionally directional (non-axisymmetric). This allows different morphologies of signals to be separated and therefore the cleaning algorithm is localised using an additional domain of information. The advantage of spin wavelets over standard scalar wavelets is to simultaneously and self-consistently probe scales and directions in the polarisation signal $P = Q + iU$ and in the underlying E and B modes, therefore providing the ability to perform component separation and E-B decomposition concurrently for the first time. We test Spin-SILC on full-mission *Planck* simulations and data and show the capacity to correctly recover the underlying cosmological E and B modes. We also demonstrate a strong consistency of our CMB maps with those derived from existing component separation methods. Spin-SILC can be combined with the pseudo- and pure E-B spin wavelet estimators presented in a companion paper to reliably extract the cosmological signal in the presence of complicated sky cuts and noise. Therefore, it will provide a computationally-efficient method to accurately extract the CMB E and B modes for future polarisation experiments.

3.2 Introduction

The polarisation of the cosmic microwave background (CMB) is a powerful cosmological observable, providing deep insights into the physics of the early universe. The decomposition of the linear polarisation into curl-free (E mode) and divergence-free (B mode) components allows the detection of tensor perturbations to the metric. Specifically, a non-zero BB power spectrum on degree scales would support the existence of a stochastic background of gravitational waves predicted by inflationary theory (Kamionkowski et al., 1997; Seljak and Zaldarriaga, 1997). Accurate measurement of B mode polarisation on arcminute scales also gives strong constraints on the neutrino sector via the weak gravitational lensing of CMB E modes (Zaldarriaga and Seljak, 1998). There are numerous existing and planned ground-based, balloon-borne and satellite experiments designed to precisely measure CMB polarisation (see *e. g.*, Errard et al. 2015 for a recent forecast on the cosmological constraining power of current and upcoming missions). However, as in measuring the temperature T anisotropies of the CMB, the polarised background needs to be separated from instrumental noise and signals due to astronomical foregrounds (in particular, synchrotron radiation and thermal radiation from Galactic dust). This foreground component separation is more difficult compared with the case of CMB temperature, due to the relative strength and morphological complexity of polarised foregrounds, which are poorly understood.

Foreground component separation has been performed in numerous ways but, on real observational data, always by removing foreground contamination from scalar signals. For example, in the polarised setting, foreground contamination is removed from the Stokes Q and U or from E and B mode maps by treating Q and U or E and B as independent scalar fields. We presented a thorough discussion of blind and non-blind component separation methods in Rogers et al. (2016) (see also *e. g.*, Bobin et al. 2013; Delabrouille et al. 2009 for reviews). In this work, we highlight only the four component separation methods employed in Planck Collaboration et al. (2015c). Commander (Eriksen et al., 2006, 2008) and SEVEM (Fernández-Cobos et al., 2012; Leach et al., 2008; Martínez-González et al., 2003) operate on the Q and U maps, while NILC (Delabrouille et al., 2009) and SMICA (Cardoso et al., 2008) operate on the E and B mode maps.

Recently, Fernández-Cobos et al. (2016) explored an extension of the Internal Linear Combination (ILC) method to act fully on the spin-2 signal formed by the Q and U maps and applied their method to simulations. In general, the ILC method estimates the CMB as a weighted sum of maps of the sky at different microwave frequencies. The weights are constrained to conserve the CMB signal but minimise foreground and noise residuals by minimising the variance of the output map. The

weights can be localised in various domains, but most usefully in wavelet space (*e. g.*, Rogers et al., 2016), which allows the weights to vary simultaneously with position on the sky and harmonic scale. Fernández-Cobos et al. (2016) minimised the covariant quantity $\langle |P|^2 \rangle$, where $P = Q + iU$, in map space. Consequently, they do not consider any harmonic localisation.

In this work, we introduce the Spin, Scale-discretised, directional wavelet ILC or Spin-SILC. This is an extension of the SILC method we introduced in Rogers et al. (2016), where to analyse spin signals such as CMB polarisation we now use spin scale-discretised wavelets, the complete construction of which is presented in McEwen et al. (2015b) (see also Leistedt et al., 2015; McEwen et al., 2014). Wavelets are functions that are localised in both real and harmonic space and, in particular, scale-discretised wavelets satisfy excellent localisation properties (McEwen et al., 2015a). In the scalar SILC method, we use directional scale-discretised wavelets (McEwen et al., 2013; McEwen et al., 2015b; Wiaux et al., 2008). Directional wavelets are spatially and harmonically localised and additionally "directionally-localised," *i. e.*, the spatial kernels are non-axisymmetric and can be rotated to pick out a preferred direction on the surface of the sphere. Rogers et al. (2016) gives an introductory summary of directional wavelets. The spin wavelets we use are still spatially, harmonically and directionally localised but are now constructed in the space of spin spherical harmonics. When spin wavelets are convolved with spin signals defined on the sphere, the output wavelet coefficients isolate signal structure of different scale and orientation, while maintaining the spatial information. The spin can in general be arbitrary, but since we are interested in analysing the spin-2 signal $P = Q + iU$ we adopt spin-2 wavelets. By the particular construction of the spin wavelets, the complex spin-2 wavelet coefficients can be separated (by their real and imaginary parts) into scalar wavelet coefficients of the E and B fields, where the scalar wavelet coefficients correspond to a scalar wavelet that is a spin-lowered version of the original spin-2 wavelet (McEwen et al., 2015b). Hence, the spin-2 wavelet transform at the heart of Spin-SILC performs E-B decomposition from input Q and U maps. The ILC method is then applied to the complex wavelet coefficients, with complex weights, and jointly minimises the variance of the reconstructed E and B fields. Moreover, the weights vary spatially, harmonically and according to different orientations, fine-tuning the cleaning algorithm to remove foreground and noise contamination.

It follows that Spin-SILC introduces two main novelties to CMB polarisation component separation. Firstly, the use of spin scale-discretised wavelets allows the full analysis of the polarisation spin signal P. By their construction, we can then perform component separation and E-B decomposition simultaneously and self-consistently. Secondly, the use of directional wavelets allows the additional flexibility to localise the foreground removal according to the morphological structure of the CMB

and the foregrounds.

There is a third novel attribute to Spin-SILC of interest to future polarisation observations. Although in this work, we have tested Spin-SILC on the full-sky multifrequency maps provided by the Planck Collaboration, these frequency maps are dominated by instrumental noise and hence so are also our estimates of the CMB polarisation. Future polarisation measurements will have high signal-to-noise, but will usually cover only a fraction of the sky. However, E-B decomposition (from the measured Q and U modes) on the cut-sky is not uniquely defined unlike the full-sky case. This leads to leaking or mixing between the E and B modes. This is of particular concern in extracting the B field since the E field is orders of magnitude larger. As presented in a companion paper (Leistedt et al., 2016), the spin scale-discretised wavelets we use can be employed to construct pure estimators of the masked E and B modes (pure E (B) modes are orthogonal to all B (E) modes on the partial sky, respectively). This builds on the work of Bunn et al. (2003); Ferté et al. (2013); Grain et al. (2012); Lewis et al. (2002); Smith and Zaldarriaga (2007) (see in particular Bunn et al. 2003 for a discussion of pure modes at the map level). This only requires calculating additional wavelet transforms of the input data subject to a suitably apodised mask. One of the main advantages of this approach is the possibility of optimising the mask as a function of scale and direction, therefore yielding a more efficient cancellation of the systematic E-B mixing due to masking (Leistedt et al., 2016). Hence, Spin-SILC can produce accurate estimates of the cosmological E and B fields, even on the cut-sky, in conjunction with the E-B estimators presented in Leistedt et al. (2016). More details of how Spin-SILC can operate on partial sky observations are given in § 3.4.8.

We provide an introduction to spin scale-discretised wavelets in § 3.3. In § 3.4, the Spin-SILC algorithm is explained in detail. We test the method on *Planck* simulations in § 3.5 and *Planck* data in § 3.6. In § 3.7, we compare our method to previous component-separation methods. We discuss the results in § 3.8 and conclude in § 3.9.

3.3 Spin wavelets

Spin, directional, scale-discretised wavelets on the sphere that support exact reconstruction have been constructed in McEwen et al. (2015b) (and discussed briefly in Leistedt et al. 2015; McEwen et al. 2014). These are an extension of the scalar, directional wavelets developed in McEwen et al. (2013); Wiaux et al. (2008), which are used in the scalar version of the SILC method for the analysis of CMB temperature anistropies (Rogers et al., 2016). They maintain the properties of spatial, harmonic and directional localisation, but can now additionally analyse spin fields by being

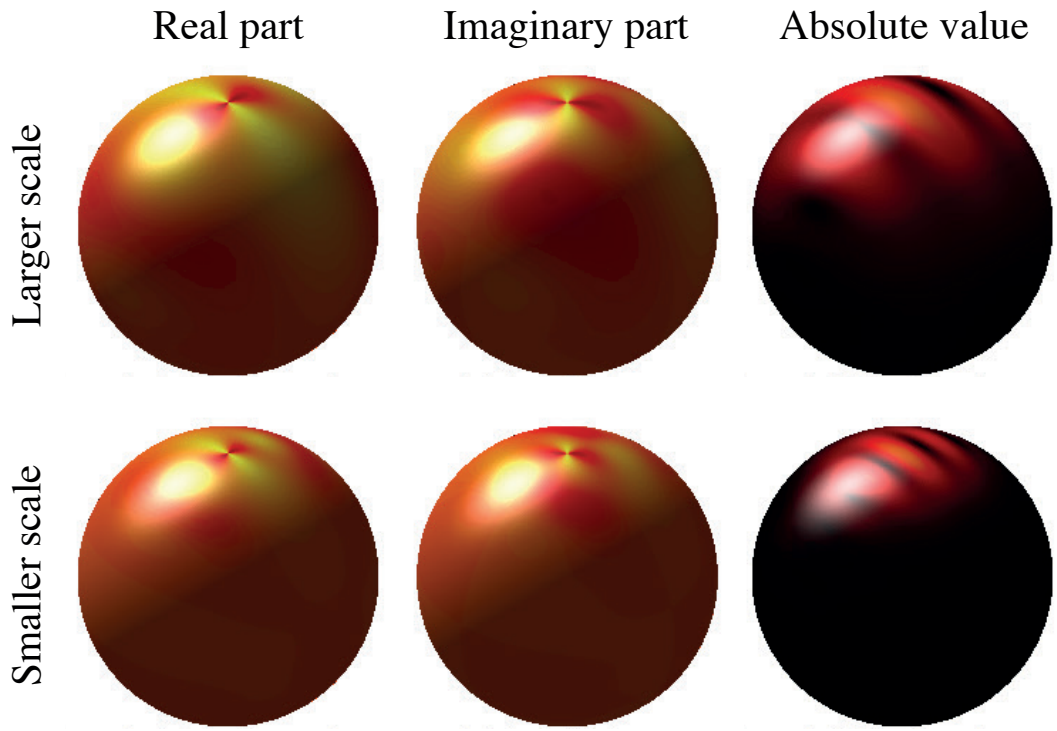

Figure 3.1: The spatial localisation on the sphere of spin, directional, scale-discretised wavelets. The top row shows larger scale wavelets than the bottom row. The left column shows the real part of the wavelet, the middle column shows the imaginary part of the wavelet and the right column shows the absolute value of the wavelet. The number of directions per wavelet scale $N = 5$. Therefore, for complete reconstruction at each scale, the above wavelets would be complemented by four more wavelets of the same size but of a different orientation on the sphere. The spin number $s = 2$, which is what is required for the analysis of Stokes Q and U modes. This figure is adapted from McEwen et al. (2015b).

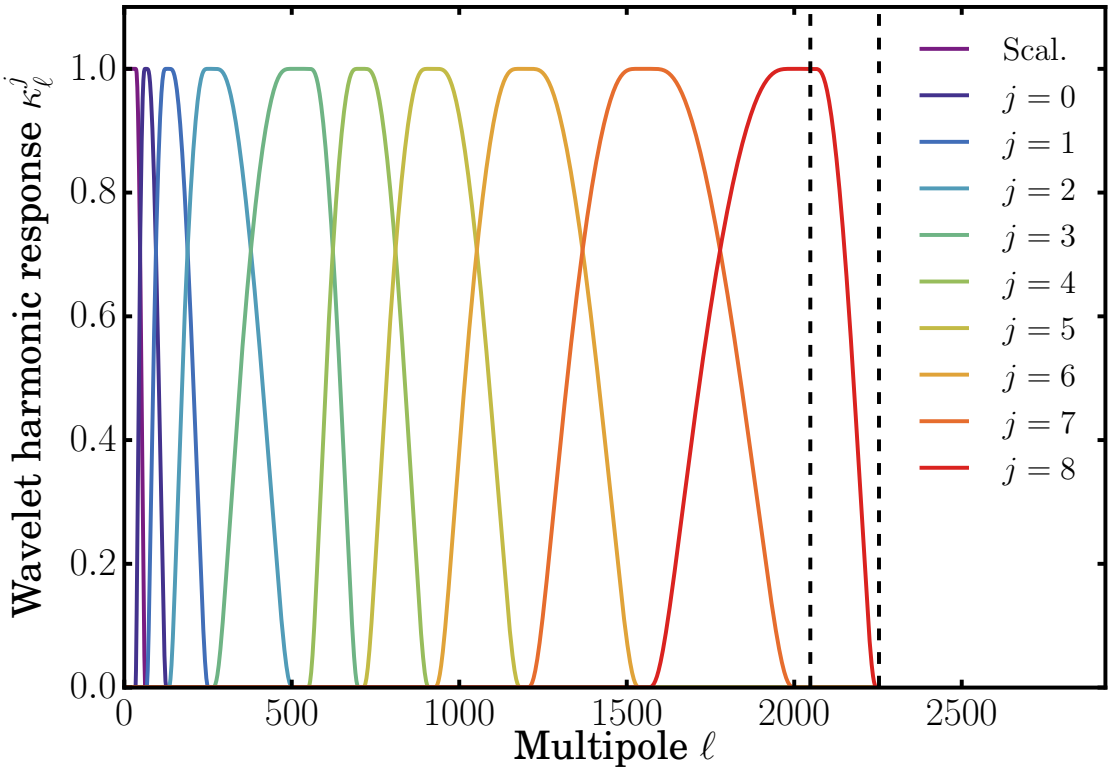

Figure 3.2: The harmonic localisation of the spin wavelets used in this work (κ_ℓ^j as defined in Eq. (3.11)), where j specifies the wavelet scale. Increasing j corresponds to a smaller wavelet kernel and so a multipole range on smaller scales (*i. e.*, larger multipoles ℓ). The largest wavelet scale (Scal.) is the scaling function (§ 3.4.4). This choice of wavelets deliberately ensures exact reconstruction only for $\ell \leq 2048$. The tapering of the smallest wavelet for $2048 < \ell \leq 2253$ suppresses the smallest-scale power within the algorithm. The band-limits of the above wavelets are given in Table 3.1.

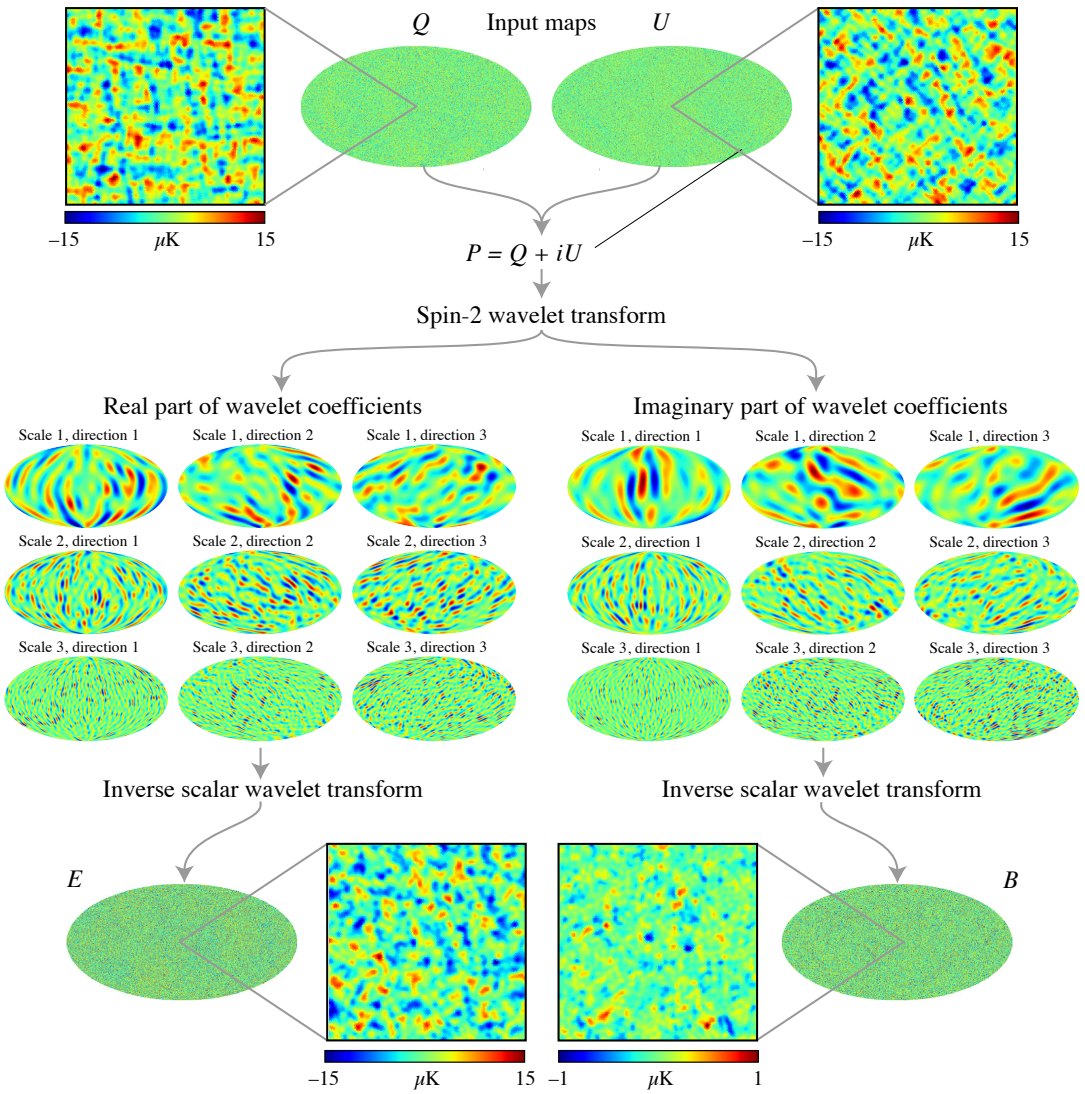

Figure 3.3: An illustration of the spin, directional wavelet decomposition of the CMB Stokes Q and U maps and the E - B separation that automatically occurs as a consequence. The top row shows example input Q and U maps, simulated with lensed scalar perturbations, with zoomed regions to show structure in the fields. The middle row shows the real and imaginary parts of the spin, directional wavelet coefficient maps, formed by the spin-2 wavelet transform of $P = Q + iU$. The ILC algorithm acts on such wavelet coefficients (calculated for multiple polarisation channels) and produces clean wavelet coefficients of the CMB polarisation. By the construction of the wavelets, the real and imaginary parts are respectively equal to scalar wavelet transforms of the E and B fields (with a different scalar wavelet). The bottom row shows the output E and B maps, also with zoomed regions, formed respectively by inverse scalar wavelet transforms of the real and imaginary parts of the wavelet coefficient maps. In our Spin-SILC analysis we include wavelets on smaller scales than those used in the simple demonstration shown above.

Table 3.1: The harmonic band-limits $[\ell^j_{\rm min}, \ell^j_{\rm max}]$ of the spin wavelets used in this work. $\ell^j_{\rm peak}$ is the multipole at which each wavelet has its maximum response. The final column shows the number of equiangular samples per wavelet coefficient map $N^j_{\rm samp}$.

Wavelet scale j	$\ell^j_{\rm min}$	$\ell^j_{\rm peak}$	$\ell^j_{\rm max}$	$N^j_{\rm samp}$
Scal.	0	64	64	8,385
0	32	64	128	33,153
1	64	128	256	131,841
2	128	256	512	525,825
3	256	512	706	998,991
4	542	705	918	1,688,203
5	705	917	1193	2,850,078
6	917	1192	1551	4,815,856
7	1192	1550	2015	8,126,496
8	1550	2015	2253	10,158,778

constructed on the basis of spin spherical harmonics. In particular, spin-2 wavelets can be convolved with the spin-2 field $P = Q + iU$, where Q and U are the Stokes parameters of the CMB's linear polarisation. Figure 3.1 shows an example of the spatial localisation of spin wavelets. Unlike scalar wavelets which are real-valued, spin wavelets are complex-valued. Figure 3.2 shows an example of the harmonic localisation of spin wavelets (for the wavelets used in this work). Figure 3.3 shows an example of spin-2 wavelet decomposition as applied to a simulated CMB polarisation field P. By the construction of the spin-2 wavelets, the real and imaginary parts of the complex wavelet coefficient maps of P are respectively scalar wavelet transforms of E and B fields (with a different scalar wavelet). It can be seen that the spin-2 wavelet transform in the Spin-SILC method carries out the decomposition of the CMB polarisation into E and B modes. Details about the use of Spin-SILC on partial sky observations are given in § 3.4.8.

3.4 Method

We start by outlining the Spin-SILC algorithm. The steps are explained in more detail in the subsequent subsections (§ 3.4.1 to 3.4.7). The use of Spin-SILC on partial sky observations is discussed in § 3.4.8. We discuss our numerical implementation in § 3.4.9.

(1). The raw input data are full-sky frequency maps of the anisotropies in the linear polarisation of the CMB, *i. e.*, Stokes Q and U fields. These maps use the HEALPix format (Górski et al., 2005). (See § 3.4.1.) The model we employ for the raw data is explained in § 3.4.2.

(2). The maps are "pre-processed" by inpainting in a small point source mask and each convolved

to have the same effective beam (see § 3.4.3).

(3). At each frequency band, the complex spin-2 polarisation field $P = Q + iU$ is formed. Each P map is converted into a set of complex-valued spin-2 wavelet coefficient maps. This separates both the scale and orientation of structure within each map. These wavelet coefficient maps are sampled according to the sampling theorem of McEwen et al. (2015). (See § 3.4.4.)

(4). A spin-2 ILC method is then applied separately to each wavelet scale and orientation. For each scale and orientation, the multifrequency wavelet coefficient maps are weighted and added to form a single (complex-valued) wavelet coefficient map that contains mainly CMB signal, as well as some residual foreground and noise. These weights are allowed to vary at each wavelet coefficient. The calculation of these weights is explained in § 3.4.5.

(5). By the construction of the spin-2 wavelets we use, the real and imaginary parts of the final ILC wavelet coefficient maps are respectively ILC estimates of the scalar wavelet transforms of the CMB E and B maps (with a different scalar wavelet). Therefore, the real and imaginary parts are separately synthesised with scalar wavelets to form the final products: full-sky maps of the CMB E and B anisotropies (with some residual foreground and noise). (See § 3.4.7.) Q and U maps are also formed by a standard spin-2 inverse wavelet transform of the ILC results (see § 3.4.6). All final maps use the HEALPix format.

3.4.1 Input data

Our CMB polarisation map products use full-mission 2015 release *Planck* Q and U polarisation maps as their input[1] (Planck Collaboration et al., 2015b,i). All seven polarisation frequency channels are used. At 70 GHz, we use the higher-resolution version at $N_{\rm side} = 2048$. As noted in Planck Collaboration et al. (2015a), the 100, 143 and 217 GHz polarisation maps have been high-pass filtered due to insufficient characterisation of residual systematic effects on large scales, in particular leakage between temperature and polarisation measurements. We therefore follow Planck Collaboration et al. (2015c) in also high-pass filtering the spherical harmonic coefficients of

[1] `http://pla.esac.esa.int/pla`

our output data products with a harmonic cosine filter:

$$w_\ell = \begin{cases} 0, & \text{if } \ell < 20, \\ \frac{1}{2}\left[1 - \cos\left(\frac{\pi}{20}(\ell - 20)\right)\right], & \text{if } 20 \leq \ell \leq 40, \\ 1, & \text{otherwise.} \end{cases} \tag{3.1}$$

We use the full-mission Full Focal Plane 8 (FFP8) simulations (Planck Collaboration et al., 2015h) with lensed scalar perturbations and without bandpass mismatch. These consist of a superposition of a CMB realisation, a noise realisation and full simulations of diffuse and point source astrophysical foregrounds.

While we do not expect an algorithm developed for next generation precision CMB polarisation observations to demonstrate its full capabilities with the *Planck* dataset, this setting comprises the best publicly-available simulations and data, and benefits from the availability of comparison data products from well-studied and highly tested component separation algorithms used by the Planck Collaboration. Thus we use the *Planck* setting to benchmark our algorithm.

3.4.2 Data model

Each of the full-sky Stokes Q and U polarisation maps ($X = Q, U$) can be independently[2] physically modelled (*e. g.*, Basak and Delabrouille, 2013) as

$$X^{\mathrm{OBS},c}(\hat{\mathbf{n}}) = \int_{\hat{\mathbf{n}}'} \mathrm{d}\hat{\mathbf{n}}' B^c(\hat{\mathbf{n}}, \hat{\mathbf{n}}') X^{\mathrm{SIG},c}(\hat{\mathbf{n}}') + X^{\mathrm{N},c}(\hat{\mathbf{n}}), \tag{3.2}$$

where the signal component can further be decomposed as

$$X^{\mathrm{SIG},c}(\hat{\mathbf{n}}) = a^c X^{\mathrm{CMB}}(\hat{\mathbf{n}}) + X^{\mathrm{FG},c}(\hat{\mathbf{n}}). \tag{3.3}$$

$X^{\mathrm{CMB}}(\hat{\mathbf{n}})$ is the CMB component at a point on the sky $\hat{\mathbf{n}}$. $X^{\mathrm{FG},c}(\hat{\mathbf{n}})$ and $X^{\mathrm{N},c}(\hat{\mathbf{n}})$ are respectively the foreground and detector noise components for frequency channel c. a^c is the calibration coefficient for the CMB for each channel. The overall signal component is smoothed by a beam function $B^c(\hat{\mathbf{n}}, \hat{\mathbf{n}}')$ due to the finite resolution of the observations. However, the noise component is not smoothed by the beam. Here we assume the beam to be circularly symmetric. Therefore, the beam

[2]The independent modelling is based on the accurate assumption that any mixing of Q and U modes in their measurement has been previously corrected in any given experiment.

can be represented as a sum over Legendre polynomials,

$$B^c(\hat{\mathbf{n}}, \hat{\mathbf{n}}') = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} B_\ell^c P_\ell(\hat{\mathbf{n}}.\hat{\mathbf{n}}'). \tag{3.4}$$

We can recast Eq. (3.2) in the scalar spherical harmonic representation as

$${}_0a_{\ell m}^{\mathrm{OBS},c} = a^c B_\ell^c \, {}_0a_{\ell m}^{\mathrm{CMB}} + B_\ell^c \, {}_0a_{\ell m}^{\mathrm{FG},c} + {}_0a_{\ell m}^{\mathrm{N},c} \tag{3.5}$$

where ${}_0a_{\ell m}$ are the coefficients of scalar spherical harmonics ${}_0Y_{\ell m}(\hat{\mathbf{n}})$.

The above is a useful representation for the data pre-processing in § 3.4.3. However, the novelty of Spin-SILC is to develop a component separation algorithm that directly makes use of the spin properties of the CMB polarisation field. The Stokes Q and U parameters are defined with respect to a fixed coordinate system on the sky. They can be identified by their respective '+' and '$\times$' patterns, as seen in the top row of Fig. 3.3. However, there is no rotationally-invariant measure of the power as a function of scale. In order to address this, we can first form the complex-valued, spin ±2 polarisation field

$$\begin{aligned} {}_{\pm2}P(\hat{\mathbf{n}}) &= Q(\hat{\mathbf{n}}) \pm iU(\hat{\mathbf{n}}) \\ &= \sum_{\ell m} {}_{\pm2}a_{\ell m} \, {}_{\pm2}Y_{\ell m}(\hat{\mathbf{n}}). \end{aligned} \tag{3.6}$$

The spin ±2 property implies that ${}_{\pm2}P(\hat{\mathbf{n}})$ transforms under local rotations of angle ψ via ${}_{\pm2}P \rightarrow e^{-is\psi} \, {}_{\pm2}P$, where spin number $s = \pm2$. The second equality therefore follows by expanding the spin field ${}_{\pm2}P(\hat{\mathbf{n}})$ in the basis of spin spherical harmonics ${}_{\pm2}Y_{\ell m}(\hat{\mathbf{n}})$ with spin spherical harmonic coefficients ${}_{\pm2}a_{\ell m}$.

We can then define the scalar E and pseudo-scalar B fields in harmonic space as

$$\begin{aligned} {}_0E_{\ell m} &= -\frac{1}{2}({}_2a_{\ell m} + {}_{-2}a_{\ell m}) \\ {}_0B_{\ell m} &= \frac{i}{2}({}_2a_{\ell m} - {}_{-2}a_{\ell m}). \end{aligned} \tag{3.7}$$

By construction, the E and B fields are real-valued and allow the calculation of rotationally-invariant angular power spectra. E modes are identified by the polarisation strength increasing in a direction parallel or perpendicular to the sense of the polarisation; it is the curl-free component of the spin-2 signal. B modes are identified by the polarisation strength increasing in a direction unaligned to the sense of the polarisation; it is the divergence-free component of the spin-2 signal. E and B modes

therefore respectively separate the underlying field into parity-even and parity-odd components.

3.4.3 Data pre-processing

The input frequency Q and U maps are diffusively inpainted in a small point source mask following the method employed by Planck Collaboration et al. (2015g). This recognises that the ILC fails when the CMB is obscured by bright extragalactic polarised sources. The inpainting removes these sources and replaces them with an extrapolation of the surrounding signal. The mask is the union of the *Planck* LFI and HFI point source masks, which are constructed from the Second Planck Catalogue of Compact Sources (PCCS2) (Planck Collaboration et al., 2015f)[3]. It masks about 0.6% of the whole sky, predominantly along the Galactic equator.

After the inpainting, we convert all the input frequency maps to the same resolution by performing a deconvolution/convolution procedure that gives scalar spherical harmonic coefficients

$$ {}_0a^c_{\ell m} = \frac{B^{\rm EFF}_\ell}{B^c_\ell}\, {}_0a^{{\rm OBS},c}_{\ell m}\,, \tag{3.8}$$

where $B^{\rm EFF}_\ell$ is the beam transfer function giving the resolution at which we perform the ILC. For *Planck* data, we use a Gaussian beam with a FWHM of 5' as our input beam. Our final map products are re-convolved to a 10' beam in order to suppress residual noise. The beams we deconvolve B^c_ℓ are taken from the Reduced Instrument Model (RIMO)[4]. For the LFI beams, we use Gaussian approximations with FWHM 32.33', 27.01' and 13.25' for 30, 44 and 70 GHz respectively. Following Planck Collaboration et al. (2014c); Rogers et al. (2016), the deconvolved beams are thresholded such that the B^c_ℓ is set to the value given in the RIMO or 0.001, whichever is larger. This suppresses noise within the ILC method.

3.4.4 Spin wavelet analysis

The spin wavelet ILC method requires the decomposition of each band-limited, complex-valued polarisation map ${}_2P^c(\hat{\mathbf{n}})$ (as formed by Eq. (3.6)) into a set of spin wavelet coefficient maps $W^{\Psi^j}_P$. We use the spin, directional, scale-discretised wavelets of Leistedt et al. (2015); McEwen et al. (2014); McEwen et al. (2015b). Following an introductory summary in § 3.3, we now discuss some of the technical details of our wavelet implementation. We drop the c superscript on ${}_2P(\hat{\mathbf{n}})$

[3]The details of their construction are given within the FITS files. They can be downloaded from `http://pla.esac.esa.int/pla`.

[4]*Planck* 2015 Release Explanatory Supplement: The 2015 instrument model (`http://wiki.cosmos.esa.int/planckpla2015/index.php/The_RIMO`).

for the rest of this subsection since each frequency map is analysed using the same wavelets. We concentrate on the spin-2 wavelet transforms we use in Spin-SILC, but the wavelets we use can be generalised to arbitrary spin.

The spin wavelet coefficients are defined as the directional convolution of ${}_2P$ with spin wavelets ${}_2\Psi^j$ defined on the sphere $\mathbb{S}^2$ (specifically those shown in Fig. 3.2), where index j denotes the wavelet scale. Like the scalar case (McEwen et al., 2013; Wiaux et al., 2008), spin, directional wavelets yield coefficients $W_P^{{}_2\Psi^j}(\hat{\rho})$ that live on the space of three-dimensional rotations, *i. e.*, the rotation group SO(3):

$$W_P^{{}_2\Psi^j}(\hat{\rho}) \equiv \langle {}_2P \,|\, \mathcal{R}_{\hat{\rho}}\, {}_2\Psi^j \rangle = \int_{\mathbb{S}^2} \mathrm{d}\hat{\mathbf{n}}\, {}_2P(\hat{\mathbf{n}}) (\mathcal{R}_{\hat{\rho}}\, {}_2\Psi^j)^*(\hat{\mathbf{n}}), \tag{3.9}$$

where $\mathrm{d}\hat{\mathbf{n}}$ is the usual invariant measure on the sphere and $\cdot^*$ denotes complex conjugation. The rotation operator is defined by

$$(\mathcal{R}_{\hat{\rho}}\, {}_2\Psi^j)(\hat{\mathbf{n}}) \equiv {}_2\Psi^j(\mathbf{R}_{\hat{\rho}}^{-1}\hat{\mathbf{n}}), \tag{3.10}$$

where $\mathbf{R}_{\hat{\rho}}$ is the three-dimensional rotation matrix corresponding to $\mathcal{R}_{\hat{\rho}}$. In Eqs. (3.9) and (3.10), $\hat{\rho} = (\theta, \phi, \chi) \in \mathrm{SO}(3)$ denotes the Euler angles (in the zyz convention) with colatitude $\theta \in [0, \pi]$, longitude $\phi \in [0, 2\pi)$ and direction $\chi \in [0, 2\pi)$. In other words, the wavelet coefficients probe directional structure in ${}_2P$ with χ corresponding to the orientation about each point (θ, ϕ) on the sphere.

Spin, directional wavelets are defined by their spin spherical harmonic coefficients in factorised form:

$${}_2\Psi_{\ell n}{}^j \equiv \sqrt{\frac{2\ell+1}{8\pi^2}} \kappa_{\ell}^j {}_2\zeta_{\ell n}, \tag{3.11}$$

where κ_ℓ^j sets the harmonic localisation (Fig. 3.2) and ${}_2\zeta_{\ell n}$ sets the directional localisation. Their spin properties are maintained by being built on the basis of spin spherical harmonics. Full details of their construction are given in McEwen et al. (2015b). As in Rogers et al. (2016), we flexibly control the harmonic localisation by using different values of the wavelet dilation parameter λ in different multipole regions. For the wavelets we use in Fig. 3.2, we use $\lambda = 2, 1.3, 1.1$ with transitions at multipoles $\ell = 512, 2015$. The harmonic bounds of each wavelet for scale j are given by $(\ell_{\min}^j, \ell_{\max}^j) = (\lambda^{j'-1}, \lambda^{j'+1})$, taking account of the different values of λ we use and the stitching-together of wavelets at λ transitions. The index j' refers to the original index of the wavelet scale as if only that single value of λ was used. Their peak response is at $\lambda^{j'}$. The details of our harmonic tiling are given in Table 3.1. A single parameter N (at all scales) defines the number

of directions into which each wavelet scale is localised.

We also use an axisymmetric scaling function ${}_2\Phi$ to form scaling coefficients $W_P^{2\Phi}$ which characterise the largest-scale information (in this work for $\ell < 64$) and live on the sphere. This is motivated by testing in Rogers et al. (2016) that showed that the use of directionality on large scales in the ILC is not effective for CMB reconstruction. These wavelets (and the scaling function) satisfy the standard admissibility criterion for exact reconstruction, *i. e.*, no information is lost in the wavelet and inverse wavelet transforms of a band-limited spin signal. For the chosen band-limit, the smallest wavelet is harmonically-truncated. We choose not to use this wavelet, which means that exact reconstruction is only satisfied for $\ell \leq 2048$. This allows the tapering of the smallest remaining wavelet (for $2048 < \ell \leq 2253$) to suppress the smallest-scale power in the algorithm.

In order to apply the ILC algorithm, the above continuous wavelet coefficients must be discretised. Since they live on the rotation group SO(3), we represent them using the sampling scheme of McEwen et al. (2015), which is itself a generalisation of the sampling scheme of McEwen and Wiaux (2011). Since the wavelets are band-limited, we use a multi-resolution scheme where each wavelet scale j is pixellated with a minimal number of samples. This means that each wavelet coefficient map $W_P^{2\Psi^j}$ (band-limited at $\ell^j_{\rm max}$) is only evaluated at samples $(\theta^j_t, \phi^j_p, \chi_n)$, where $t \in \{0, 1, \ldots, \ell^j_{\rm max}\}$, $p \in \{0, 1, \ldots, 2\ell^j_{\rm max}\}$ and $n \in \{0, 1, \ldots, N-1\}$. In this way, each spin wavelet coefficient map can be separated into N spin, directional wavelet coefficient maps W^{2P}_{jnk} according to the value of n, where k indexes pixel number according to samples (θ^j_t, ϕ^j_p) on the sphere. It follows that each input frequency P map has been decomposed into wavelet coefficient maps, each localised according to harmonic scale j and orientation of structure n, while maintaining spatial localisation (pixel number k). (See Fig. 3.3 for a demonstration of this decomposition.)

3.4.5 ILC method

Following the spin, directional wavelet analysis of the input P maps (see § 3.4.4), there is a spin, directional wavelet coefficient map $W^{2P,c}_{jnk}$ for each channel c, scale j and orientation n with a pixel index k. With this compact notation, we conflate the scaling coefficient maps with the wavelet coefficient maps as the ILC method applies in exactly the same way. We develop the spin wavelet ILC method by an extension of the scalar wavelet ILC method we developed in Rogers et al. (2016) to operate on the complex-valued wavelet coefficient maps we now have. Similar to Fernández-Cobos et al. (2016), we consider the complex spin signal ${}_{\pm2}P$ rather than considering scalar fields independently (*e. g.*, Q and U independently or E and B independently). However,

unlike Fernández-Cobos et al. (2016), who work jointly on Q and U maps in real space, we work in wavelet space, where spatial, scale and directional localisation is possible. The most general extension of the scalar ILC is to estimate the CMB at each wavelet scale and orientation as a sum of wavelet coefficient maps for each frequency with complex-valued weights ω^c_{jnk}:

$$W_{jnk}^{2\,P,\mathrm{ILC}} \equiv \sum_{c=1}^{N_c} \omega^c_{jnk} W_{jnk}^{2\,P,c}\,, \tag{3.12}$$

where N_c is the number of input channels.

In order to recover an unbiased estimate of the CMB, we impose a constraint on the weights such that

$$\sum_{c=1}^{N_c} a^c \omega^c_{jnk} = 1 + 0i, \tag{3.13}$$

where we remind the reader that a^c is the real-valued set of calibration coefficients for the CMB Q and U maps introduced in Eq. (3.3). In order to calculate the weights at each pixel k, we choose to minimise the covariant quantity $\left\langle \left| W_{jnk}^{2\,P,\mathrm{ILC}} \right|^2 \right\rangle$ with respect to the complex-valued weights ω^c_{jnk} themselves, under the constraint in Eq. (3.13). This minimisation can be carried out with complex Lagrange multipliers (similarly to the scalar case) giving complex-valued weights

$$\omega^c_{jnk} = \frac{\sum_{c'=1}^{N_c} (R^{-1}_{jnk})^{cc'} a^{c'}}{\sum_{c=1}^{N_c} \sum_{c'=1}^{N_c} a^c (R^{-1}_{jnk})^{cc'} a^{c'}}, \tag{3.14}$$

where the true covariance matrices at scale j, orientation n and pixel k, $(R_{jnk})^{cc'} = \left\langle W_{jnk}^{2\,P,c^*} W_{jnk}^{2\,P,c'} \right\rangle$ (where the angled brackets indicate an ensemble average, although in practice we empirically estimate these covariances as explained below). In this work, we assume $a^c = 1,\ \forall\, c$, *i. e.*, that the CMB is perfectly calibrated in the data we use.

There are two main consequences from minimising the quantity we choose. First, as in scalar SILC, we assume that the CMB and foregrounds and the CMB and noise are respectively uncorrelated. It follows that the ensemble cross-term between CMB and residual contamination is zero and, since the CMB is conserved by the constraint in Eq. (3.13), we are minimising only the variance of the error in CMB reconstruction. Second, by minimising the variance of the full complex-valued spin-2 wavelet coefficients using complex weights, we are in turn jointly minimising the variance of the E and B modes. This is thanks to the construction of the spin wavelets we use, as discussed in § 3.4.6. Unlike a foreground cleaning algorithm acting on the Q and U or E and B maps separately, this approach ensures that all the information (*i. e.*, from the multiple polarisation channels and the Q

and U cross terms) is used to jointly construct clean CMB polarisation P and its E and B modes.

As in scalar SILC, we estimate the covariance matrices $(R_{jnk})^{cc'}$ empirically on the data. We achieve this by replacing the appropriate ensemble average with a weighted average of the surrounding pixels. Specifically, we smooth the maps of covariance matrix elements with a Gaussian kernel in harmonic space. The size of this kernel is proportional to the size of the wavelet used at each scale[5]. Full details of this empirical estimation of covariances and possible optimisations to the method are given in Rogers et al. (2016).

3.4.6 Spin wavelet synthesis to Stokes Q and U modes

Although a novelty of Spin-SILC is to simultaneously perform E-B decomposition and component separation (*i. e.*, to synthesise the ILC results directly to E and B maps as explained in § 3.4.7), one can also form Q and U maps. This is carried out by a single spin-2 inverse wavelet transform. After the ILC method (see § 3.4.5) has been applied to the frequency wavelet coefficient maps, there is one ILC estimate of the CMB P field (with some residual foreground and noise) at each wavelet scale and orientation $W_{jnk}^{2\,P,\mathrm{ILC}}$. Multiple orientations $\chi_0, \chi_1, \ldots, \chi_{N-1}$ are combined at each scale to form wavelet coefficient maps $W_P^{2\,\Psi^j,\mathrm{ILC}}(\hat{\rho})$ that live on SO(3). We also have the scaling coefficient map $W_P^{2\,\Phi,\mathrm{ILC}}(\hat{\mathbf{n}})$ that lives on the sphere (and characterises the largest scales). In order to calculate our real space estimate of the CMB polarisation spin field ${}_2P^{\mathrm{ILC}}(\hat{\mathbf{n}})$ (and hence the Stokes parameters $Q^{\mathrm{ILC}}(\hat{\mathbf{n}})$ and $U^{\mathrm{ILC}}(\hat{\mathbf{n}})$), we perform the following spin-2 inverse wavelet transform:

$$\begin{aligned} {}_2P^{\mathrm{ILC}}(\hat{\mathbf{n}}) &= Q^{\mathrm{ILC}}(\hat{\mathbf{n}}) + iU^{\mathrm{ILC}}(\hat{\mathbf{n}}) \\ &= \int_{\mathbb{S}^2} \mathrm{d}\hat{\mathbf{n}}' W_P^{2\,\Phi,\mathrm{ILC}}(\hat{\mathbf{n}}')(\mathcal{R}_{\hat{\mathbf{n}}'}\,{}_2\Phi)(\hat{\mathbf{n}}) \\ &\quad + \sum_j \int_{\mathrm{SO}(3)} \mathrm{d}\hat{\rho} W_P^{2\,\Psi^j,\mathrm{ILC}}(\hat{\rho})(\mathcal{R}_{\hat{\rho}}\,{}_2\Psi^j)(\hat{\mathbf{n}}), \end{aligned} \tag{3.15}$$

where $\mathrm{d}\hat{\rho}$ is the usual invariant measure on the rotation group. We have used the same spin wavelets as in the wavelet analysis in § 3.4.4. The final ILC Q and U maps are pixellated in the HEALPix format.

[5] $\mathrm{FWHM}^j = 50\sqrt{\frac{1200}{N^j_{\mathrm{samp}}}}$. This value is the same as used in the NILC implementation on *Planck* data.

3.4.7 Scalar wavelet synthesis to E and B modes

A considerable advantage of the Spin-SILC method is that it simultaneously removes foreground and noise contamination from the cosmological signal and carries out the E-B decomposition discussed in § 3.4.2. It achieves the latter by using a property of the spin scale-discretised wavelets that relates ${W_P^2}^{\Psi^j}(\hat{\rho})$, the spin-2 wavelet transform of P to ${W_{\tilde{E}}^0}^{\Psi^j}(\hat{\rho})$ and ${W_{\tilde{B}}^0}^{\Psi^j}(\hat{\rho})$, the scalar wavelet transforms of $\tilde{E}$ and $\tilde{B}$:

$$\begin{aligned} {W_{\tilde{E}}^0}^{\Psi^j}(\hat{\rho}) &= -\mathrm{Re}[{W_P^2}^{\Psi^j}(\hat{\rho})] \\ {W_{\tilde{B}}^0}^{\Psi^j}(\hat{\rho}) &= -\mathrm{Im}[{W_P^2}^{\Psi^j}(\hat{\rho})]. \end{aligned} \tag{3.16}$$

The intermediate fields $\tilde{E}$ and $\tilde{B}$ are respectively related to E and B by a harmonic normalisation of their scalar spherical harmonic coefficients:

$$\begin{aligned} {}_0E_{\ell m} &= \frac{1}{N_{\ell,2}} {}_0\tilde{E}_{\ell m} \\ {}_0B_{\ell m} &= \frac{1}{N_{\ell,2}} {}_0\tilde{B}_{\ell m}, \end{aligned} \tag{3.17}$$

where $N_{\ell,s} = \sqrt{\frac{(\ell+s)!}{(\ell-s)!}}$.

The straightforward E-B decomposition is achieved by the construction of the wavelets and is discussed in detail in McEwen et al. (2015b) and Leistedt et al. (2016). In Eq. (3.16) the scalar wavelets ${}_0\Psi^j$ are spin-lowered versions of the spin-2 wavelets ${}_2\Psi^j$:

$${}_0\Psi^j(\hat{\mathbf{n}}) = \overline{\eth}^2 {}_2\Psi^j(\hat{\mathbf{n}}). \tag{3.18}$$

(An equivalent equation links the scalar and spin scaling functions.) $\overline{\eth}$ is a first-order differential operator known as the spin-lowering operator since it lowers the spin of spherical harmonic functions: $\overline{\eth}\, {}_sY_{\ell m}(\hat{\mathbf{n}}) = \frac{1}{N_{\ell,s}}\, {}_{s-1}Y_{\ell m}(\hat{\mathbf{n}})$.

By applying Eq. (3.16), we can separate the spin wavelet coefficient maps ${W_P^2}^{\Psi^j,\mathrm{ILC}}(\hat{\rho})$ (defined in § 3.4.6) into scalar wavelet coefficient maps of the intermediate $\tilde{E}$ and $\tilde{B}$ modes ${W_Y^0}^{\Psi^j,\mathrm{ILC}}(\hat{\rho})$, where $Y = \tilde{E}, \tilde{B}$. An equivalent separation forms the scaling coefficient maps ${W_Y^0}^{\Phi,\mathrm{ILC}}(\hat{\mathbf{n}})$. In order to calculate our real space estimates of the CMB $\tilde{E}$ and $\tilde{B}$ modes ${}_0Y^{\mathrm{ILC}}(\hat{\mathbf{n}})$, we perform inverse scalar wavelet transforms with the (spin-lowered) scalar scaling function ${}_0\Phi$ and scalar wavelets

${}_0\Psi^j$ (as defined in Eq. (3.18)):

$$
\begin{aligned}
{}_0Y^{\mathrm{ILC}}(\hat{\mathbf{n}}) = &\int_{\mathbb{S}^2} \mathrm{d}\hat{\mathbf{n}}' W_Y^{0^{\Phi,\mathrm{ILC}}}(\hat{\mathbf{n}}')(\mathcal{R}_{\hat{\mathbf{n}}'}\,{}_0\Phi)(\hat{\mathbf{n}}) \\
&+ \sum_j \int_{\mathrm{SO}(3)} \mathrm{d}\hat{\rho} W_Y^{0^{\Psi^j,\mathrm{ILC}}}(\hat{\rho})(\mathcal{R}_{\hat{\rho}}\,{}_0\Psi^j)(\hat{\mathbf{n}}).
\end{aligned}
\tag{3.19}
$$

The output scalar spherical harmonic coefficients can be renormalised to the usual E and B fields by applying Eq. (3.17). The final ILC E and B maps are pixellated in the HEALPix format.

3.4.8 Spin-SILC on partial sky observations

We have outlined above the Spin-SILC method specifically as it applies on the full sky. However, Spin-SILC is being primarily developed for application to future CMB polarisation experiments, which will have greater signal-to-noise and/or resolution, but will typically observe only part of the sky. The decomposition of Stokes Q and U measurements into E and B modes is essential for cosmological analyses, in particular for a measurement of the BB angular power spectrum. This is strictly well-defined only on the whole sky, as in Eq. (3.7).

This decomposition is not well-defined if the input measurements only cover a part of the sky. However, following Bunn et al. (2003), a polarisation field on the cut-sky can be decomposed into a complete orthonormal basis defined by "pure E", "pure B" and "ambiguous" modes. Pure E modes have vanishing curl and are orthogonal to all B modes on the partial sky. Pure B modes have vanishing divergence and are orthogonal to all E modes on the partial sky. Ambiguous modes are all other modes, which will have both vanishing divergence and curl. It is the inability to distinguish ambiguous modes which leads to the problem of E-B leakage where ambiguous modes are erroneously counted as E or B. However, following, *e. g.*, Bunn et al. (2003); Smith and Zaldarriaga (2007), if the pure B modes can be isolated, they will form an estimate of the cosmological B power, unbiased by E-B leaking.

Leistedt et al. (2016) have shown how the spin wavelets we use can be employed to construct estimates of the pure modes defined above on a masked sky. This builds on the work of, *e. g.*, Grain et al. (2012); Lewis et al. (2002); Smith and Zaldarriaga (2007). In particular, the spin wavelet pure-mode estimation of Leistedt et al. (2016) requires only two additional wavelet transforms of the input Q and U data and a suitably apodised mask. The Spin-SILC method can be applied on partial sky observations by coherently combining the full-sky method with the pure mode estimation. The application of Spin-SILC to partial sky observations will be investigated in future

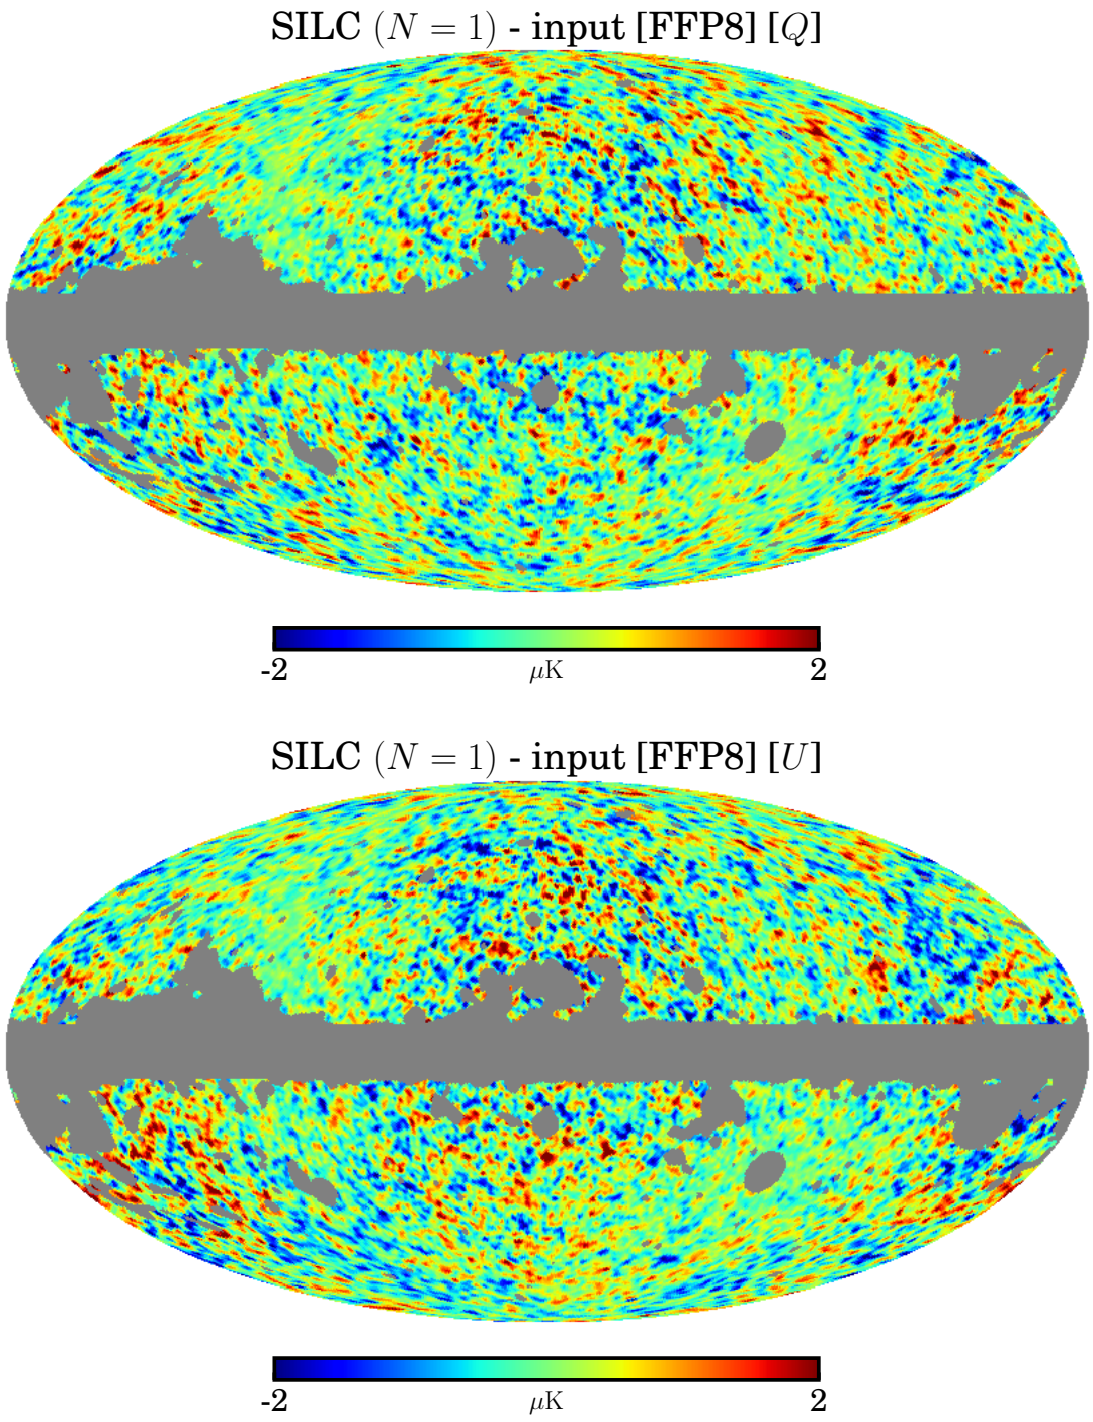

Figure 3.4: Planck *simulations.* Differences between output ILC and input CMB maps from FFP8 simulations with lensed scalar perturbations. The maps have been smoothed to $\mathrm{FWHM} = 80'$ and downgraded to $N_{\mathrm{side}} = 128$. The grey pixels are the UPB77 confidence mask from Planck Collaboration et al. (2015c), which masks the Galactic region in FFP8 simulations where foreground emission is strongest. *From top to bottom,* we show differences in (a) Stokes Q and (b) Stokes U maps.

work, providing the first integrated pipeline to simultaneously carry out E-B decomposition and foreground component separation for future CMB polarisation experiments.

3.4.9 Numerical implementation

Spin-SILC is implemented in `Python` and is parallelised. At full *Planck* resolution ($N_{\rm side} = 2048$, $\ell_{\rm max} = 2253$), when run on a 60-core symmetric multiprocessor (SMP) with 1.5 TB RAM and a 24-core cluster node with 256 GB RAM[6], the pipeline takes approximately 1.5 hours per direction. The wavelet transforms in Spin-SILC are carried out using the latest version of the `S2LET`[7] code (Leistedt et al., 2013; McEwen et al., 2015), written in `C` with `Python` wrappers. This employs `SSHT`[8] (McEwen and Wiaux, 2011) and `SO3`[9] (McEwen et al., 2015) to compute spin spherical harmonics and Wigner transforms exactly and efficiently. Spin-SILC is developed from the scalar SILC[10] code (Rogers et al., 2016) (which performs component separation on the temperature anisotropies of the CMB).

3.5 Application to *Planck* simulations

We tested Spin-SILC on the fiducial full-mission *Planck* FFP8 simulated Stokes Q and U sky maps. We use simulations with lensed scalar perturbations. Figure 3.4 shows the differences between the reconstructed CMB (using $N = 1$) and the input simulated CMB. The two panels show the differences in Q and U maps, as this most directly compares to the input data. The most striking features are the reductions in residuals in the top left and bottom right corners, aligning with the Ecliptic poles. This reflects reduced noise residuals because there is less noise in the input data due to the scanning strategy of the *Planck* satellite, which integrated for longer in those directions. In general, the difference maps are consistent with noise residuals. This is a consistent attribute of the *Planck* polarisation datasets. Figure 3.5 compares the full-sky angular power spectra ($D_\ell = \ell(\ell+1)C_\ell/2\pi$) of the same reconstructed CMB and the input signal. The three panels respectively compare the EE, BB and EB spectra, as these are the cosmologically-interesting observables. The EE and BB spectra are consistent with significant residual noise power due to the noisiness of the input maps, although the first four acoustic peaks of the EE spectrum are

[6]The exact specification for our infrastructure is an Intel Xeon E7-4890 2.8 GHz SMP with 4 $\times$ 15-core CPUs with 25.6 GB RAM per core, and an Intel Xeon E5-2697 2.7 GHz node with 2 $\times$ 12-core CPUs with 10.7 GB RAM per core.
[7]`http://www.s2let.org`
[8]`http://www.spinsht.org`
[9]`http://www.sothree.org`
[10]`http://www.silc-cmb.org`

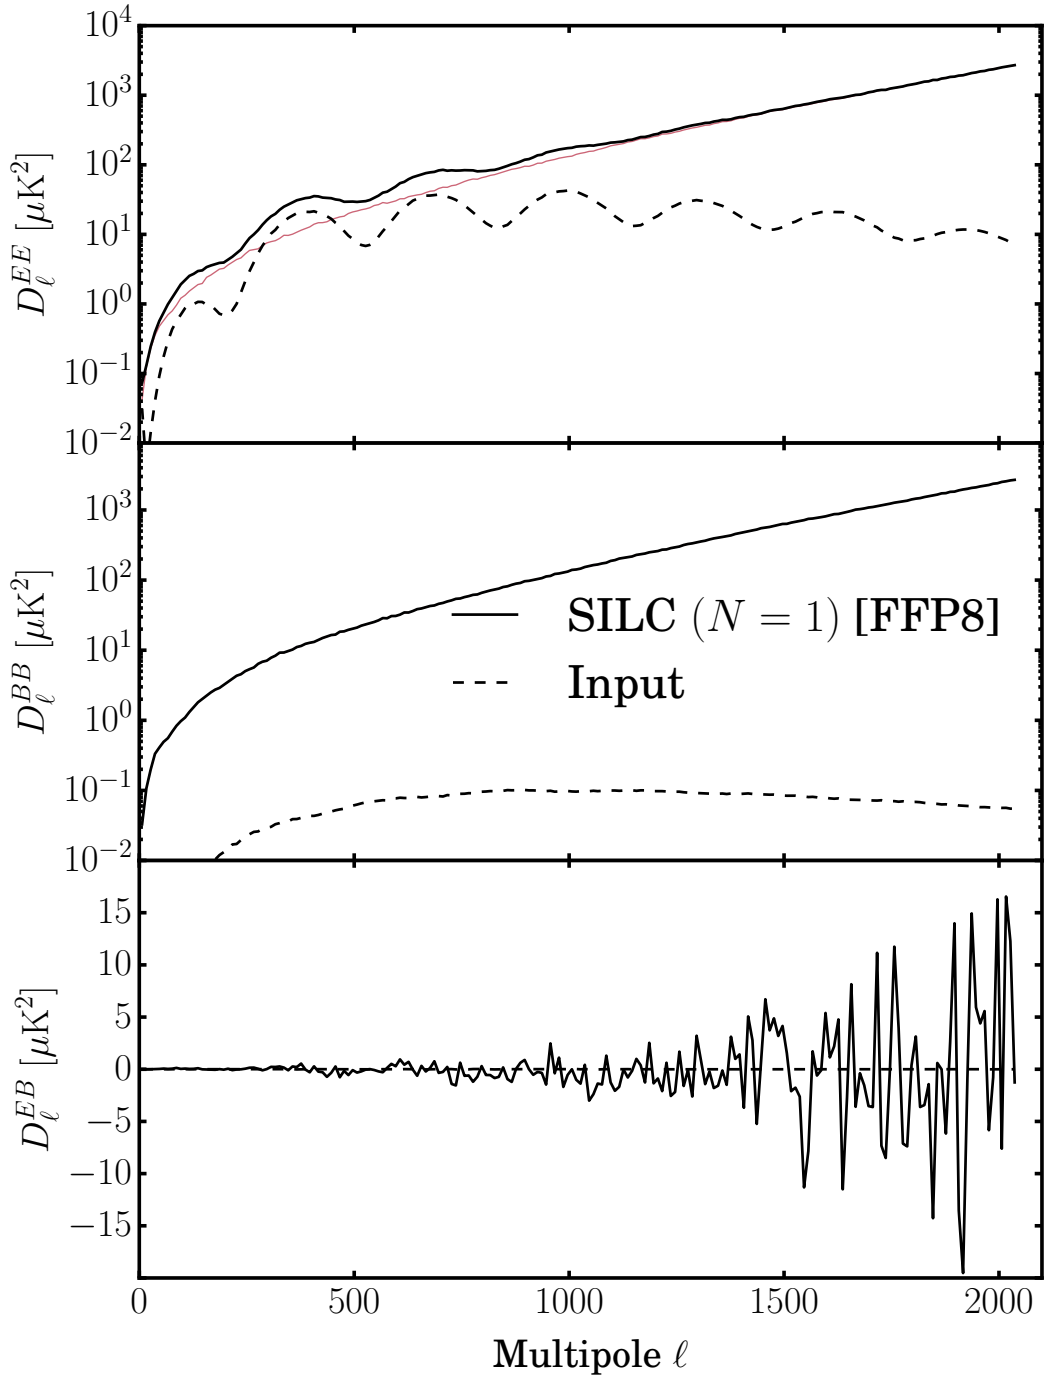

Figure 3.5: Planck *simulations. From top to bottom,* (a) EE, (b) BB and (c) EB angular power spectra comparing output ILC in the axisymmetric limit $(N = 1)$ to input CMB from FFP8 simulations with lensed scalar perturbations. *In the top panel* (a), the thin red line shows residuals after subtracting the input CMB spectrum.

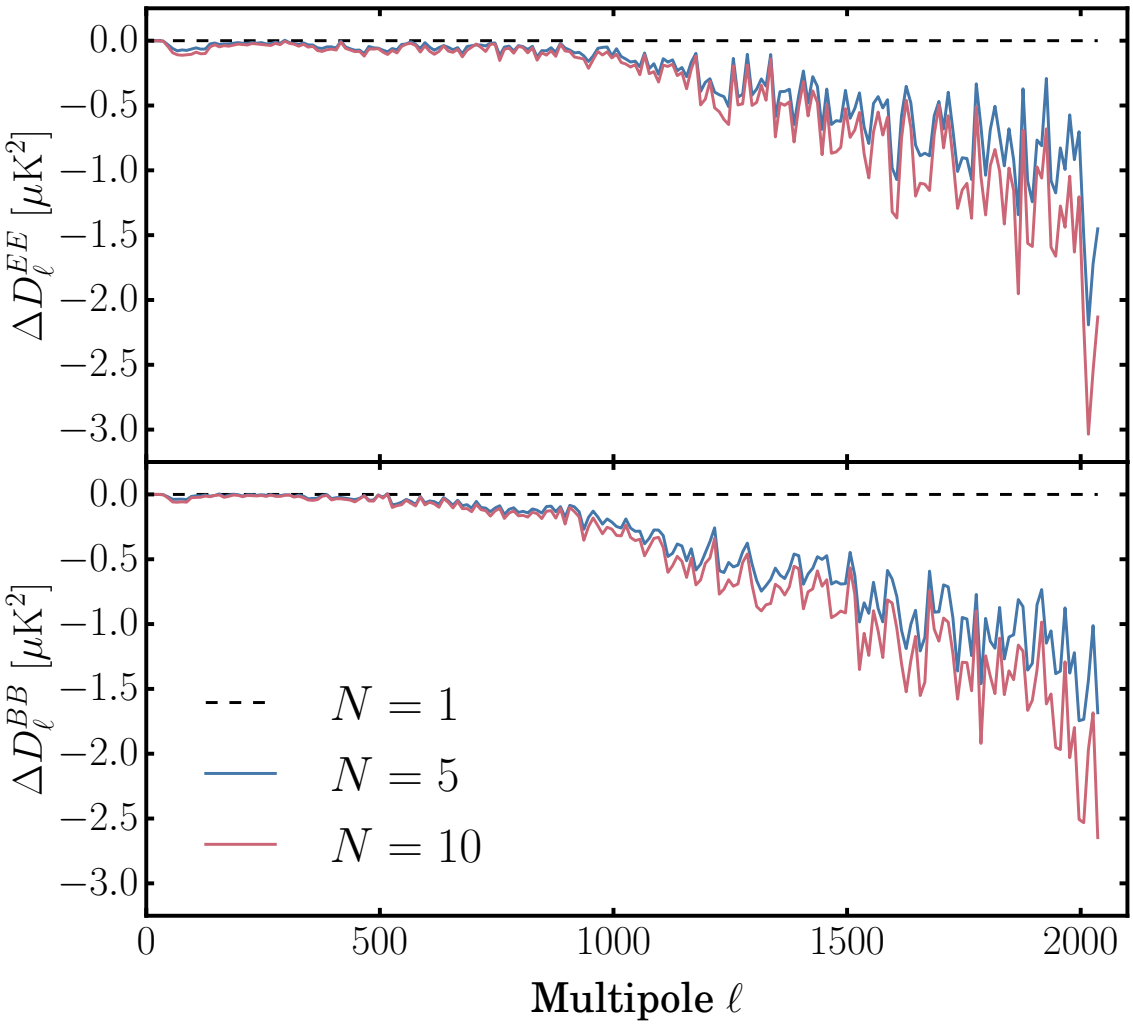

Figure 3.6: Planck *simulations.* Differences between angular power spectra of different values of N minus the axisymmetric limit ($N = 1$). The input data are FFP8 simulations with lensed scalar perturbations. *From top to bottom,* we show differences in (a) EE and (b) BB spectra. We note the small amplitude of the reductions in reconstruction residuals from increasing N.

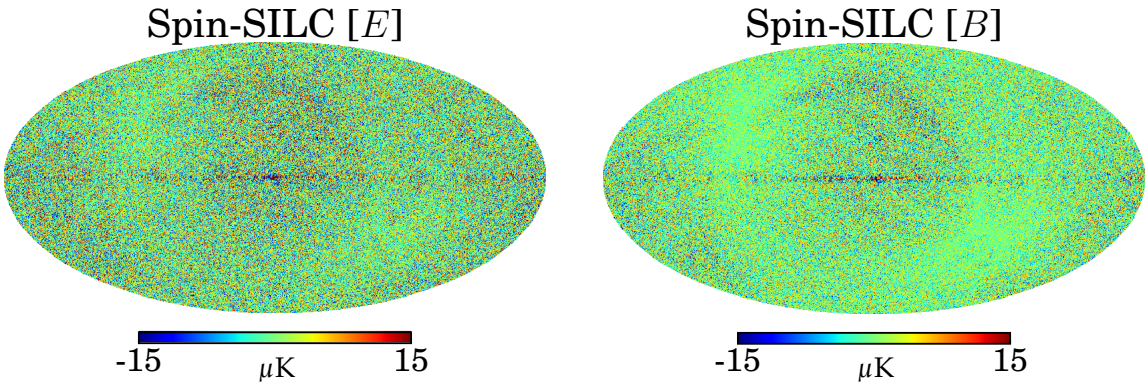

Figure 3.7: Planck *data. From left to right,* (a) the CMB polarisation E map and (b) the CMB polarisation B map reconstructed using Spin-SILC in the axisymmetric limit ($N = 1, \mathrm{FWHM} = 10', N_{\mathrm{side}} = 1024$).

discernible nonetheless. The reconstructed EB spectrum is consistent with the zero input value.

We tested the impact of using directional spin wavelets in the Spin-SILC method with the simulated dataset. Figure 3.6 compares the differences in full-sky power spectra between the directional case (for $N = 5, 10$) minus the axisymmetric limit ($N = 1$). The two panels compare EE and BB spectra. It can be seen that using more directional wavelets per wavelet scale reduces power spectrum reconstruction residuals with respect to the axisymmetric limit, very modestly on large scales and more so on small scales. However, the magnitude of these reductions is very small compared to the total power in the output ILC maps, which are dominated by residual instrumental noise; in the BB spectrum, the reduction is comparable to the magnitude of the input lensing signal. These results are fully expected following Rogers et al. (2016), where it was found that the gains in component separation efficacy from employing directionality was marginal in the low signal-to-noise (S/N) regime.

3.6 Application to *Planck* data

After testing Spin-SILC on the simulated dataset, we apply as input data the real full-mission *Planck* Stokes Q and U maps. Figure 3.7 shows our main output data products: full-sky ILC estimates of the CMB polarisation E and B modes using Spin-SILC in the axisymmetric limit ($N = 1$). We show E and B maps in order to highlight the E-B decomposition from input Q and U maps that Spin-SILC automatically carries out thanks to the construction of the spin wavelets that we use (see § 3.4.4). We reiterate that these maps have been high-pass filtered (for $\ell < 40$) in order to mitigate

for residual systematics in the *Planck* polarisation data (see § 3.4.1). The maps are consistent with large levels of residual instrumental noise, with the scanning pattern of the *Planck* satellite clearly visible. We also note the poor reconstruction in the Galactic plane, particularly towards the Galactic centre, where foreground emission is strongest and most complex.

3.7 Comparison to previous work

Having presented the main results of applying Spin-SILC to *Planck* data, we can perform a validation check by comparing to other component separation reconstructions of the CMB from the same dataset. For this purpose, we concentrate on the methods NILC (Delabrouille et al., 2009) and SMICA (Cardoso et al., 2008), which are two of the four methods used internally by the Planck Collaboration (Planck Collaboration et al., 2015c): the former because it is the most similar method to Spin-SILC and the latter because it is the baseline method adopted by the Planck Collaboration for high-resolution analyses. Like Spin-SILC, NILC is an internal linear combination (ILC) method performed in wavelet space. Unlike Spin-SILC, NILC uses scalar axisymmetric wavelets, specifically scalar needlets (Baldi et al., 2009; Marinucci et al., 2008; Narcowich et al., 2006), rather than the spin directional wavelets we use (Leistedt et al., 2015; McEwen et al., 2014; McEwen et al., 2015b) (although spin needlets (Geller et al., 2008) and mixed needlets (Geller and Marinucci, 2011) have also been developed). This means that in its extension to polarisation (Basak and Delabrouille, 2013; Planck Collaboration et al., 2015c), NILC acts independently on input E and B maps, having been previously decomposed from the original Stokes Q and U data. Similarly to ILC methods, SMICA forms a linear combination of multifrequency data, but in harmonic space. Unlike blind ILC methods which require no physical modelling of the sky components, SMICA is only semi-blind in that on large scales, rather than empirically estimating covariances on the data (as on small scales), a fit is performed to a model of the component covariances, with the option to constrain these covariances. This is extended to polarisation by performing a joint processing of the E and B modes in harmonic space. A further difference between ILC methods and SMICA is that SMICA has no spatial localisation in its component separation, although a wavelet implementation of SMICA does exist (Moudden et al., 2005). Indeed, Spin-SILC localises with regard to the greatest number of domains of information, allowing spatial, harmonic and morphological localisation through the use of directional wavelets (see Rogers et al. 2016 for a discussion of the morphological localisation properties of directional wavelets in the SILC method). A significant advantage of Spin-SILC over existing component separation methods (including NILC and SMICA) is the use of spin wavelets,

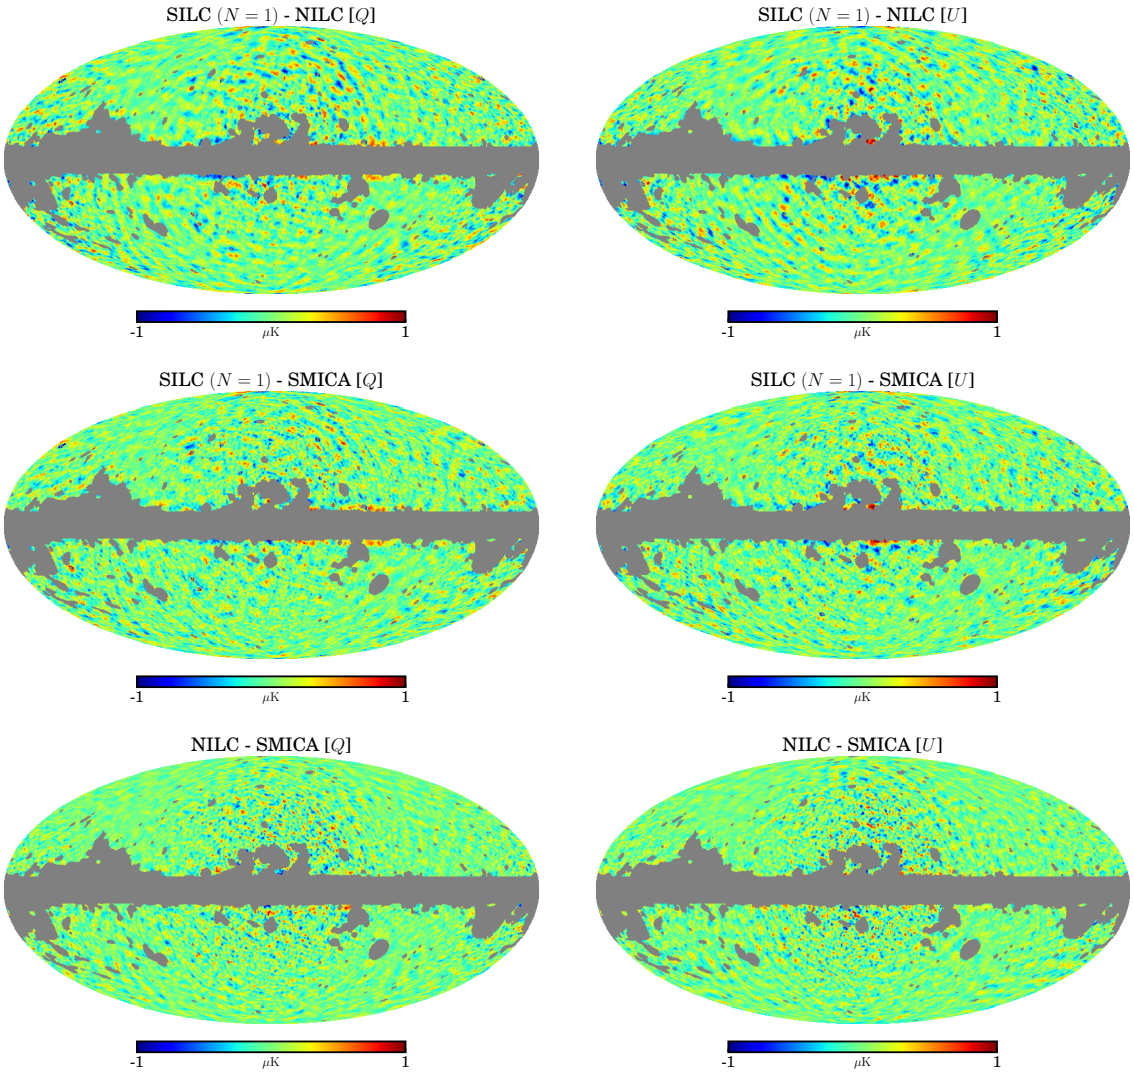

Figure 3.8: Planck *data*. Differences between the axisymmetric limit $(N = 1)$ of Spin-SILC, NILC and SMICA. The maps have been smoothed to $\mathrm{FWHM} = 80'$ and downgraded to $N_{\mathrm{side}} = 128$. The grey pixels are the UPB77 confidence mask from Planck Collaboration et al. (2015c), which masks the regions of the NILC and SMICA maps not recommended for cosmological analysis. The differences are (*from top to bottom*) (a) SILC $(N = 1)$ - NILC, (b) SILC $(N = 1)$ - SMICA and (c) NILC - SMICA; and in (*from left to right*) (i) Stokes Q and (ii) Stokes U maps.

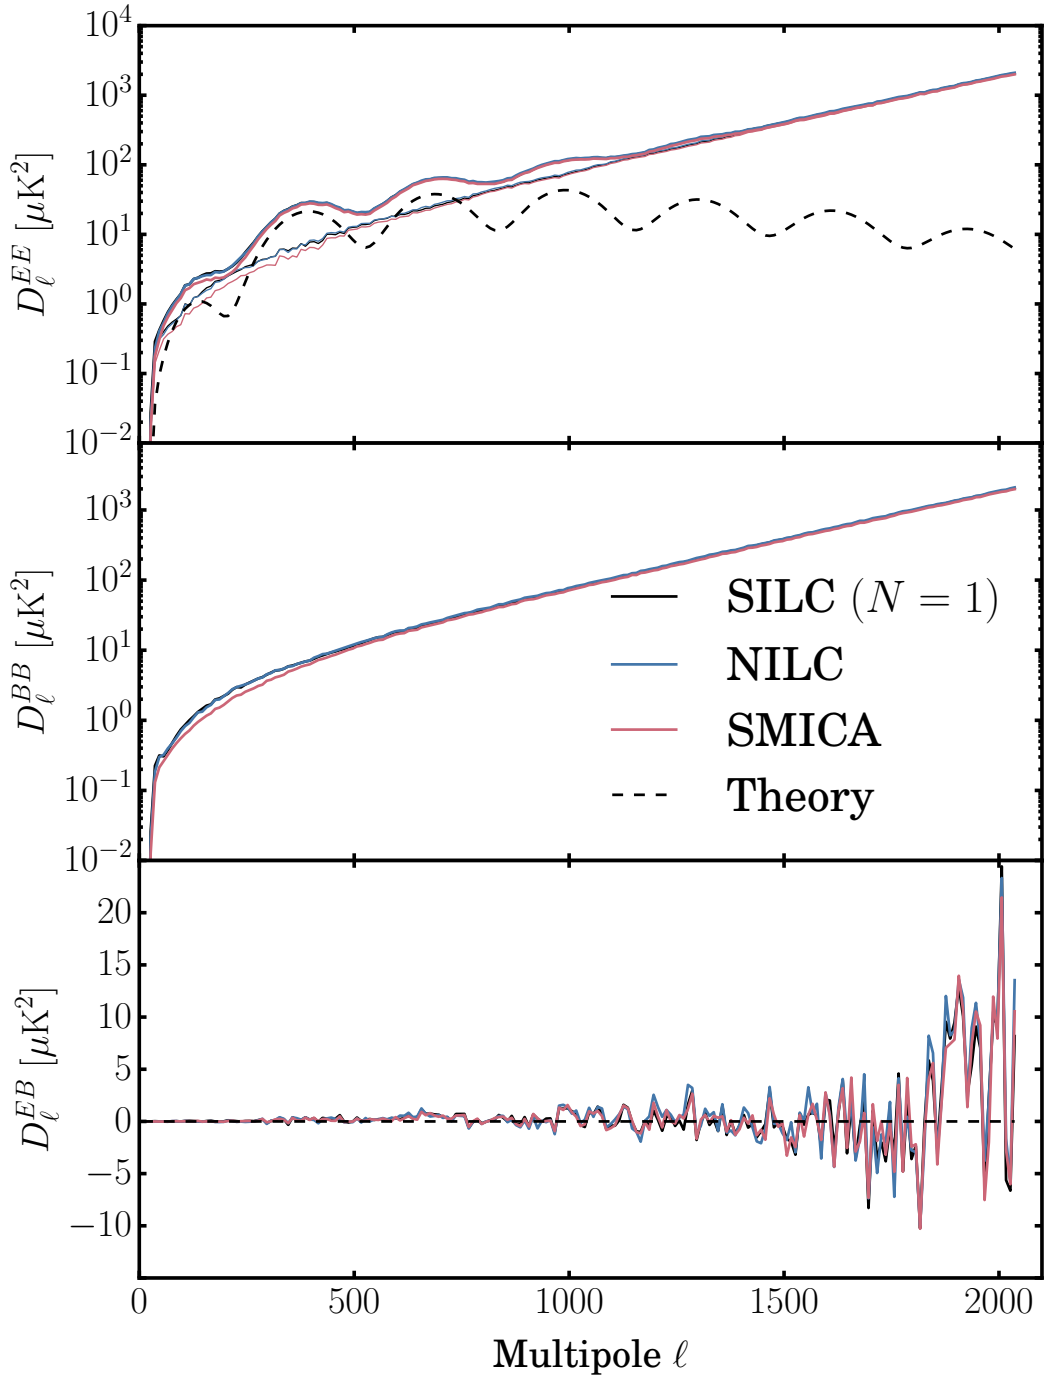

Figure 3.9: Planck *data. From top to bottom,* (a) EE, (b) BB and (c) EB angular power spectra comparing the axisymmetric limit $(N = 1)$ of Spin-SILC to NILC and SMICA. *In the top panel* (a), the thin lines show residuals after subtracting the best-fit ΛCDM model from the *Planck* 2015 likelihood.

which allows simultaneous component separation and E-B decomposition.

We can empirically compare the three methods with an analysis of the CMB maps reconstructed from the (full-mission 2015 release) *Planck* data and full-sky power spectra measured from those maps. Figure 3.8 shows the differences between the CMB reconstructed by Spin-SILC (in the axisymmetric limit $N = 1$), NILC and SMICA. We show differences in Q and U maps as this most directly compares to the map products provided by the Planck Collaboration. The differences between the three methods are small in magnitude in both Q and U and mostly concentrated at the edges of the Galactic mask towards the Galactic centre, where foreground emission is most intense and complex. Quantitatively, we can compare the mean values and standard deviations of the full-sky difference maps. The mean values of the Q difference maps in Figs. 3.8 (i) (a), (b) and (c) (*from top to bottom on the left-hand side*) are respectively 5.2×10^{-5}, 4.3×10^{-5} and -9.3×10^{-6} μK, while the standard deviations are 0.34, 0.37 and 0.34 $\mu\mathrm{K}^2$. The mean values of the U difference maps in Figs. 3.8 (ii) (a), (b) and (c) (*from top to bottom on the right-hand side*) are respectively 6.4×10^{-5}, -6.9×10^{-5} and -1.3×10^{-4} μK, while the standard deviations are 0.33, 0.37 and 0.31 $\mu\mathrm{K}^2$. These values are small and similar, suggesting a strong consistency between the three methods. As discussed in § 3.4.1, some of the input *Planck* data are high-pass filtered and so are also the output results of all three methods (for $\ell < 40$). This means that any comparison can only be carried out for $\ell \geq 40$.

Figure 3.9 compares full-sky power spectra measured from component separation maps with CMB spectra derived from the *Planck* 2015 TT and low TEB likelihood[11]. The three panels respectively compare EE, BB and EB spectra. As with the simulated results in § 3.5, the EE and BB spectra from all three methods are consistent with significant residual noise power due to the noisiness of the input maps. The only discernible difference is marginally less power in the SMICA maps at multipoles around $\ell = 250$. This could be attributed to the semi-blindness of SMICA better characterising the noise properties of the *Planck* data. The two blind methods, Spin-SILC and NILC have near-identical spectra at all multipoles. The EB spectra of all three methods are consistent with zero.

The comparison of Spin-SILC to existing methods NILC and SMICA has strongly validated the results we showed in § 3.6. An analysis of maps and power spectra shows an internal consistency between the three algorithms. It also shows that the power of Spin-SILC and other component separation methods is limited by the low S/N of the *Planck* polarisation data, with large amounts of

[11]The parameters come from the `base_plikHM_TT_lowTEB` likelihood. The values are available in the *Planck* 2015 Release Explanatory Supplement: 2015 Cosmological parameters and MC chains (`http://wiki.cosmos.esa.int/planckpla2015/images/f/f7/Baseline_params_table_2015_limit68.pdf`).

residual noise in the reconstructed CMB. The full potential of the Spin-SILC method thus awaits the input of higher-S/N polarisation data available from upcoming CMB observations.

3.8 Discussion

The testing of Spin-SILC on *Planck* simulations in § 3.5 and data in § 3.6 shows that the use of spin wavelets in CMB polarisation component separation can successfully reconstruct the cosmological background. This is particularly true of the EE power spectrum with the clear detection of the first four acoustic peaks in both simulations and real data. The residual maps to the input simulated CMB (Fig. 3.4) and power spectra estimated from the SILC maps (Fig. 3.5) show high levels of residual noise, reflecting the relatively low S/N of the *Planck* data we used. In § 3.7, we carried out a comparison of the Spin-SILC method with two of the most accurate existing component separation algorithms, NILC and SMICA. We validated our main results by showing a strong internal consistency in reconstructed CMB maps (Fig. 3.8) and power spectra measured from those maps (Fig. 3.9). However, this comparison also revealed high levels of residual noise in the CMB estimated by all three methods, due to the S/N limitations of the *Planck* polarisation data.

We also tested the use of directional spin wavelets in the Spin-SILC method on the simulations in § 3.5. We found very modest reductions in reconstructed residual power (Fig. 3.6) as the amount of directionality was increased. However, given the low S/N input data, the magnitude of these reductions is much smaller than the overall amount of residual power, though at the accuracy required to reconstruct the cosmological BB signal, this level of power reduction may become relevant in the high S/N regime. As discussed in Rogers et al. (2016) in the scalar SILC method, instrumental noise has no particular directional structure and thus in the low S/N regime the use of directionality is expected to only have a small effect on the estimate of the reconstructed CMB. The community is only just beginning to accurately characterise polarised foregrounds at high resolution at a range of frequencies. If the foregrounds are complex in high S/N observations, the ability to use directional wavelets may prove useful in localising component separation according to the morphology of the CMB and foregrounds.

Spin-SILC introduces a number of novelties into CMB polarisation component separation. Most notably, the use of spin wavelets allows simultaneous E-B decomposition and the joint minimisation of E and B auto-correlations in residual contamination. Moreover, the use of directionality allows the fine-tuning of the cleaning algorithm according to the morphology of the local signal. As discussed in Rogers et al. (2016), there are various sources of error in the ILC method, which will

affect the spin-ILC in an equivalent fashion. Of particular note is the ILC bias, corresponding to the empirical cancellation of CMB modes due to chance correlations with foregrounds and noise (see Delabrouille et al. 2009 for a fuller discussion of this effect), which will also affect the reconstruction of the CMB polarisation. The amount of cancellation may increase with the amount of directionality used within the method. In Rogers et al. (2016), we showed that this can be mitigated either directly from the data or through suites of simulations.

3.9 Conclusions

We have presented Spin-SILC, a foreground component separation method specifically developed for the analysis of CMB polarisation data. The use of spin wavelets allows the full analysis of the spin-2 polarisation signal $P = Q + iU$, formed by the Stokes Q and U parameters. By the particular construction of the spin wavelets we use, Spin-SILC carries out the decomposition of the polarisation signal into E and B modes by separating the real and imaginary parts of the complex spin-2 wavelet coefficients. This occurs simultaneously to the component separation, where the auto-correlations of E and B modes are jointly minimised in residual contamination to the reconstructed CMB. Moreover, the wavelets we use are directional. This allows different directional morphologies of CMB and polarised foreground to be separated. This extra information can then be used to better localise the Spin-SILC cleaning algorithm.

We have tested Spin-SILC on full-mission *Planck* simulations and data. We showed that the method can accurately extract cosmological information from input Q and U maps. We also validated our main results with a comparison to the internal *Planck* methods, NILC and SMICA, showing a strong consistency in both CMB maps and power spectra, with small residuals compared to the two. However, we note that the analysis in this paper is limited by the low S/N of the *Planck* polarisation data. Our final E and B maps (as well as those of NILC and SMICA) are dominated by residual instrumental noise. Moreover, the full power of the use of directionality in Spin-SILC cannot be fully explored due to the high level of noise in the *Planck* input data. If polarised foregrounds have complex morphology in the high S/N regime, then the use of directionality may prove a useful extra tool in extracting the CMB. In general, it will be interesting to test Spin-SILC further with the high S/N data of upcoming CMB polarisation observations. We make our Q, U, E and B maps available at `http://www.silc-cmb.org`[12].

Furthermore, Spin-SILC can be combined with the estimators of Leistedt et al. (2016) to perform

[12]The DOI for our data release is 10.5281/zenodo.50579.

component separation on the cut-sky and give accurate estimates of pure E and B modes (pure E (B) modes are orthogonal to all B (E) modes on the cut-sky, respectively). It achieves this in a straightforward fashion (with only two additional wavelet transforms of the input data) due to the construction of the spin wavelets (see Leistedt et al. 2016 for more details about pure mode estimation on the cut-sky using spin wavelets). This is of particular importance for the upcoming high resolution, high S/N CMB polarisation experiments, which will typically make partial-sky observations. Spin-SILC will provide a computationally-efficient algorithm to perform simultaneous E-B decomposition and accurate foreground component separation for these next-generation experiments.

Acknowledgements

KKR thanks Franz Elsner and Stephen Feeney for valuable discussions. 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.

4

Simulating the effect of high column density absorbers on the one-dimensional Lyman-alpha forest flux power spectrum

4.1 Abstract

We measure the effect of high column density absorbing systems of neutral hydrogen (HI) on the one-dimensional (1D) Lyman-alpha forest flux power spectrum using cosmological hydrodynamical simulations from the Illustris project. High column density absorbers (which we define to be those with HI column densities $N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$) cause broadened absorption lines with characteristic damping wings. These damping wings bias the 1D Lyman-alpha forest flux power spectrum by causing absorption in quasar spectra away from the location of the absorber itself. We investigate the effect of high column density absorbers on the Lyman-alpha forest using hydrodynamical simulations for the first time. We provide templates as a function of column density and redshift, allowing the flexibility to accurately model residual contamination, *i. e.*, if an analysis selectively clips out the largest damping wings. This flexibility will improve cosmological parameter estimation, *e. g.*, allowing more accurate measurement of the shape of the power spectrum, with implications for cosmological models containing massive neutrinos or a running of the spectral index. We provide fitting functions to reproduce these results so that they can be incorporated straightforwardly into a data analysis pipeline.

4.2 Introduction

The Lyman-alpha forest (a series of neutral hydrogen absorption lines in the spectra of quasars) is a uniquely powerful probe of the clustering of matter at redshifts from about $z = 2$ to $z = 6$ (Croft et al., 1998; Iršič et al., 2016; McDonald et al., 2000, 2005b; Palanque-Delabrouille et al., 2013; Viel et al., 2013) and from sub-Mpc to hundreds of Mpc scales. The one-dimensional (1D) Lyman-alpha forest flux power spectrum (along the line of sight) is particularly sensitive to small-scale clustering in the quasi-linear regime and provides important constraints on extended cosmological models that suppress small-scale power (Armengaud et al., 2017; Iršič et al., 2017a,b; Palanque-Delabrouille et al., 2015; Seljak et al., 2005; Yeche et al., 2017), notably those containing massive neutrinos and warm dark matter. This small-scale information complements the larger scales probed by the angular power spectrum of the cosmic microwave background (CMB). For example, the best upper limit on the sum of neutrino masses (Palanque-Delabrouille et al., 2015) comes from combining CMB data from the Planck Collaboration (Planck Collaboration et al., 2016f) with the 1D Lyman-alpha forest power spectrum as measured from Sloan Digital Sky Survey (SDSS)-III/Baryon Oscillation Spectroscopic Survey (BOSS) Data Release 9 (DR9) quasar spectra (Dawson et al., 2013; Eisenstein et al., 2011; Palanque-Delabrouille et al., 2013).

Future surveys like the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al., 2016a,b) will further improve constraints on extended cosmological models. Font-Ribera et al. (2014b) forecast one-sigma errors on a DESI measurement of the sum of neutrino masses to be 0.017 eV[1]. Considering that the lower limit on the sum of neutrino masses from neutrino oscillation experiments is 0.06 eV (Esteban et al., 2017; Forero et al., 2014; Gonzalez-Garcia et al., 2014), this would constitute at least a three-sigma detection. Furthermore, the 1D Lyman-alpha forest flux power spectrum probes the primordial power spectrum on the smallest currently accessible scales, $k \sim 4$ Mpc^{-1}. Including Lyman-alpha forest data will improve constraints on the running of the spectral index (which quantifies deviations from a pure power-law spectrum) by a factor of two, reaching one-sigma errors of ± 0.002 (Font-Ribera et al., 2014b). This would provide new insights into early universe physics, potentially ruling out classes of models of inflation. Importantly, it will also provide a unique independent cross-check at small scales of the primordial power spectrum shape inferred from CMB measurements at large scales.

Achieving these limits requires marginalisation over the uncertain impact of a number of

[1]This is the full forecasted constraint considering a combination of *Planck* CMB data, DESI broadband galaxy power spectrum, DESI broadband Lyman-alpha forest flux power spectrum and ~ 100 high-resolution Lyman-alpha forest quasar spectra.

astrophysical effects on the 1D Lyman-alpha forest power spectrum. In particular, this includes broadened absorption features from high column density absorbers. High column density absorbers are usually classified as either damped Lyman-alpha absorbers (DLAs), with column densities $N(\mathrm{HI})$ exceeding $2 \times 10^{20}\,\mathrm{atoms\,cm^{-2}}$ (Wolfe et al., 1986), or Lyman-limit systems (LLS), which correspond to $2 \times 10^{20}\,\mathrm{atoms\,cm^{-2}} > N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$. Both types of system produce broad damping wings which extend to large distances in redshift space. If not accounted for, they will bias cosmological parameter estimation from the Lyman-alpha forest. The systems are formed at peaks of the underlying density distribution; consequently, they cluster more strongly than the forest itself (Font-Ribera et al., 2012b).

To remove the bias induced by damped absorbers, one can fit a model for their effect on power spectra. The most widely used approach (McDonald et al., 2005a) is now more than a decade old. Although this model was adequate for the data available at the time, future surveys will be substantially more constraining and therefore demand tighter control over systematics. Furthermore, there have been significant improvements in theoretical modelling of these systems (*e. g.*, Bird et al., 2015; Pontzen et al., 2008). An updated model for the effects of high column density absorbers is therefore both timely and essential in order to achieve the forecasted cosmological constraints from future surveys.

Different column densities correspond to gas at different physical densities, so that simulations suitable for modelling the forest are often not suited to reproducing high column density systems. The Lyman-alpha forest is largely insensitive to the physics of galaxy formation since it is sourced by gas at below mean density; the primary uncertainties arise from cosmological parameters and the thermal history of the intergalactic medium. Conversely, high column density absorbers arise largely from regions within or around galaxies and are thus very sensitive to the physics of galaxy formation and less sensitive to large-scale cosmology. It is consequently essential to model the effect of high column density absorbers using simulations which include detailed galaxy formation physics and can thus reproduce their characteristics and statistics.

In Lyman-alpha forest studies, damping wings are sometimes "clipped" (*i. e.*, removed or masked) from quasar spectra (*e. g.*, see Lee et al., 2013, for details of the process for BOSS DR9 spectra). However, not all damping wings are identified and many will remain in the spectra, especially in noisier spectra where they are harder to spot and for lower-density absorbers (*i. e.*, LLS) which have narrower wings. Therefore, in the final cosmological parameter estimation from the 1D Lyman-alpha forest power spectrum, the effect of residual high column density absorbers is modelled as a multiplicative scale-dependent bias of the power spectrum with an amplitude (reflecting the level of

residual contamination) that is fitted and marginalised (Palanque-Delabrouille et al., 2015). The functional form of this model (*i. e.*, its scale and redshift dependence) is based on the measurements made in McDonald et al. (2005a).

McDonald et al. (2005a) investigated the effect with lognormal model mock quasar spectra (*i. e.*, generated without hydrodynamical simulations; details of their generation are given in McDonald et al., 2006), since the numerical simulations available at the time were not large enough to generate spectra encompassing the full width of damping wings. They then probe the effect of high column density absorbers on the Lyman-alpha forest by inserting damping wings in mock spectra at the peaks of the lognormal field, based on the observationally-determined column density distribution function (CDDF). They find a systematic effect on the observed 1D Lyman-alpha forest power spectrum that is maximised on scales corresponding to the width of a damped system and which has negligible redshift evolution (considering three redshift slices at $z = [2.2, 3.2, 4.2]$). They provide a single template to fit their bias measurement, including the effect of all LLS and DLAs together. However, as discussed above, in current data analysis pipelines, damping wings are removed from quasar spectra in a way that preferentially removes higher density systems. Therefore, when the template is used in parameter inference, it may not correctly model the bias of the *residual* contamination, which will have a different CDDF to the total — the clipping of the survey spectra changes the survey CDDF. The bias will have a different scale-dependence (not just amplitude), since this is driven by the distribution of the widths of damping wings remaining in quasar spectra.

In this work, we investigate the effect of high column density absorbers on the 1D Lyman-alpha forest power spectrum as a function of their column density and redshift using hydrodynamical simulations of galaxy formation from the Illustris project (Nelson et al., 2015; Vogelsberger et al., 2014a). Comparison to relevant observations has shown that Illustris reproduces the observed CDDF and spatial clustering of high-density systems (Bird et al., 2014; Vogelsberger et al., 2014a, see § 4.4.1 for more details) at the 95% confidence level. Spectra are generated from this simulation, then separated into categories according to the maximum column density within each spectrum (see § 4.3 for more details). We measure the 1D flux power spectrum of each of these types of spectrum and measure the (multiplicative) bias of each type compared to the power spectrum of the Lyman-alpha forest alone. We make this measurement at multiple redshifts and so probe the redshift evolution of this effect.

We discuss high column density absorbers in more detail in § 4.3. In § 4.4, our methodology in going from hydrodynamical simulations to measurements of the 1D flux power spectrum is explained. We present our main results in § 4.5. These results are discussed in § 4.6 and in § 4.7,

Table 4.1: The neutral hydrogen (HI) column density limits $[N(\mathrm{HI})_{\mathrm{min}}, N(\mathrm{HI})_{\mathrm{max}}]$ that define the categories of absorbing systems used in this work. The columns on the right show the percentage of spectra (at each redshift z that is considered) in our $(106.5\,\mathrm{Mpc})^3$ simulation box (Nelson et al., 2015; Vogelsberger et al., 2014a, Illustris-1) where the highest-density system belongs to a given category.

Category	$N(\mathrm{HI})_{\mathrm{min}}$	$N(\mathrm{HI})_{\mathrm{max}}$	% of spectra in $(106.5\,\mathrm{Mpc})^3$ simulation at				
	$[\mathrm{atoms\,cm^{-2}}]$		$z = 2.00$	$z = 2.44$	$z = 3.01$	$z = 3.49$	$z = 4.43$
Forest	0	1.6×10^{17}	77.7	69.6	57.4	45.7	22.0
LLS	1.6×10^{17}	1×10^{19}	10.6	14.9	21.8	27.0	36.6
Sub-DLA	1×10^{19}	2×10^{20}	5.9	8.1	11.4	14.3	20.1
Small DLA	2×10^{20}	1×10^{21}	3.1	4.1	5.5	7.8	12.8
Large DLA	1×10^{21}	∞	2.7	3.3	3.9	5.2	8.5

we present the templates that we have fitted to our measurements. Finally, conclusions are drawn in § 4.8.

4.3 Damped Lyman-alpha absorbers and Lyman-limit systems

High column density absorbers are regions of neutral hydrogen (HI) gas that are above a column density threshold of $N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$. By contrast, lower column density absorbers form the Lyman-alpha forest. The absorption lines formed by high column density absorbers are broadened, forming damping wings and hence absorption in the spectrum away from the location of the absorbing gas. The damping wings have a characteristic Voigt profile, which is a convolution of a Gaussian profile (caused by Doppler broadening) and a Lorentzian profile (caused by natural or collision broadening). The width of these wings in velocity space increases with the column density of the absorbing system. High column density absorbers are then usually classified as either damped Lyman-alpha absorbers (DLAs), whose damping wings are considered significantly broadened and which correspond to $N(\mathrm{HI}) > 2 \times 10^{20}\,\mathrm{atoms\,cm^{-2}}$ (Wolfe et al., 1986); or Lyman-limit systems (LLS), which correspond to column densities in the range $2 \times 10^{20}\,\mathrm{atoms\,cm^{-2}} > N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$.

In this work, we aim to investigate the effect of high column density absorbers (and especially their damping wings) on the one-dimensional Lyman-alpha forest flux power spectrum, as a function of their column density (and redshift). We therefore use a more refined classification of high column density absorbers based on their column densities, in particular accounting for the fact that higher density LLS do have wide damping wings. Table 4.1 shows the column density limits that define our

categories, as well as the percentage of simulated spectra (see § 4.4.1) where the highest-density system is a given type and hence is the main contaminant. The overall percentage of spectra contaminated by high column density absorbers (LLS, sub-DLAs, small and large DLAs) increases with redshift because the HI CDDF increases at higher densities at higher redshifts, but always there are more LLS than DLAs.

4.4 Method

We first outline the method we have used and then explain the steps in more detail in the following subsections (§ 4.4.1 to 4.4.3).

(1). We use a cosmological hydrodynamical simulation from the Illustris project (Nelson et al., 2015; Vogelsberger et al., 2014a) and generate mock spectra on a grid (562 500 in total, each at a velocity resolution of $25\,\mathrm{km\,s^{-1}}$ and with a typical length of $\simeq 8\,000\,\mathrm{km\,s^{-1}}$). We repeat this for a number of redshift slices at which the Lyman-alpha forest is observed ($z = [2.00, 2.44, 3.01, 3.49, 4.43]$). (See § 4.4.1.)

(2). For each redshift slice, we separate the spectra according to the highest column density system within that spectrum using the absorber categories defined in Table 4.1. For each absorber category (and the total set of spectra), we measure the one-dimensional (1D) flux power spectrum (*i. e.*, along the line of sight, integrating over transverse directions) using a fast Fourier transform (FFT). (See § 4.4.2.)

(3). We then measure the (multiplicative) bias of the flux power spectra from each category relative to the 1D flux power spectrum of the Lyman-alpha forest, as a function of absorber type (*i. e.*, maximum column density) and redshift (see § 4.4.3). We fit parametric models to these bias measurements and provide these templates in § 4.7.

4.4.1 Hydrodynamical simulations and mock spectra

Our main results make use of snapshots from the highest-resolution (in terms of both dark matter particles and hydrodynamical cells) cosmological hydrodynamical simulation from the Illustris project (Nelson et al., 2015; Vogelsberger et al., 2014a, Illustris-1[2]). The simulation adopts the following cosmological parameters: $\Omega_\mathrm{m} = 0.2726$, $\Omega_\Lambda = 0.7274$, $\Omega_\mathrm{b} = 0.0456$, $\sigma_8 = 0.809$, $n_\mathrm{s} = 0.963$ and $H_0 = 100\,h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, where $h = 0.704$ (Vogelsberger et al., 2014b). The

[2]The simulation we use is publically available at `http://www.illustris-project.org/data`.

box has a volume in comoving units of $(106.5\,\mathrm{Mpc})^3$ and we consider snapshots at redshifts $z = [2.00, 2.44, 3.01, 3.49, 4.43]$.

The Illustris simulations use the moving mesh code AREPO (Springel, 2010). The galaxy formation physics implemented is of relevance to dense regions of neutral hydrogen gas, and therefore we describe it briefly here. The subgrid models include prescriptions for supernova (Springel and Hernquist, 2003; Vogelsberger et al., 2013) and active galactic nuclei (AGN) (Sijacki et al., 2007; Springel et al., 2005) feedback (Bird et al. 2014 showed that the properties of DLAs are quite insensitive to the details of AGN feedback); radiative cooling; star formation and metal enrichment of gas. Self-shielding is implemented as a correction to the photoionization rate, which is a function of hydrogen density and gas temperature. The potential ionising effect of local stellar radiation within the most dense absorbers (*i. e.*, large DLAs) (*e. g.*, Fumagalli et al., 2011) is neglected. Pontzen et al. (2010) found this effect to be negligible and accurate calculations in any case require physics on parsec scales, well below the resolution of the simulation (it can then be viewed as part of the unresolved physics included in the above feedback prescriptions). More details of these models are given in Bird et al. (2014); Vogelsberger et al. (2013). Gravitational interactions are computed using the TreePM approach (Springel, 2005).

We require that these simulations accurately reproduce the necessary statistics of high column density absorbers that are observed in surveys. As a means of quantifying this, we can first consider the CDDF of neutral hydrogen over relevant column densities ($N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$). Vogelsberger et al. (2014a) make a comparison of the CDDF as produced by Illustris centered at $z = 3$ to the distribution observed in a number of surveys [Prochaska et al. (2010) for LLS; Zafar et al. (2013) for sub-DLAs; Noterdaeme et al. (2009) for DLAs]. In particular, the distributions are consistent with the feature in the CDDF around the DLA threshold, where the distribution rises, being reproduced well (the results of Bird et al. 2017 from SDSS-III DR12 spectra are also consistent for DLAs). Bird et al. (2014) showed that the AREPO code with the above hydrodynamical models can produce values of the DLA halo bias (at $z = 2.3$) which are in agreement with measured values from real surveys (Font-Ribera et al., 2012b), indicating that the clustering of high column density absorbers is well reproduced. Bird et al. (2015) compared the distribution function of velocity widths of low ionization metal absorbers associated with DLAs as produced by the simulations at $z = 3$ to the distribution observed in Neeleman et al. (2013). The data points are within the 68% confidence interval of the simulated distribution. This suggests that the simulations are reproducing the kinematics, and thus the host halo distribution, of high column density absorbers. One potential caveat is that these simulations are found to produce too high a total incidence rate of DLAs when

compared to observations (Noterdaeme et al., 2012) at $z = 2$ (Bird et al., 2014). However, the overall incidence rate is absorbed into a normalisation that must in any case be allowed to float during analysis of clipped spectra (as discussed in § 4.6).

For each snapshot, we generate mock spectra containing only the Lyman-alpha absorption line (*i. e.*, with a rest wavelength of 1215.67Å) from neutral hydrogen. We do this on a square grid of 562 500 spectra, in the plane perpendicular to a direction that we define as the line of sight. Each spectrum extends the full length of the simulation box with periodic boundary conditions, giving a size in velocity (or "redshift") space of $\{7111, 7501, 8000, 8420, 9199\}\,\mathrm{km\,s^{-1}}$ respectively at $z = [2.00, 2.44, 3.01, 3.49, 4.43]$[3]. We first measure the optical depth τ in velocity bins of size $25\,\mathrm{km\,s^{-1}}$ along the spectrum[4]. We further convolve our spectra with a Gaussian kernel of $\mathrm{FWHM} = 8\,\mathrm{km\,s^{-1}}$, setting the simulated spectrographic resolution. We then calculate the transmitted flux $\mathcal{F} = e^{-\tau}$. In this way, the spectra we have constructed are insensitive to contamination from other absorption (or emission) lines, estimation of the emitted quasar continuum (which here is effectively set to unity) or instrumental noise. In each spectrum pixel, we are also able to measure the column density (integrated along the line of sight in each bin) of neutral hydrogen, which we use in measuring the maximum density systems in each spectrum (§ 4.4.2).

4.4.2 One-dimensional flux power spectrum

We separate our spectra into the absorber categories (Lyman-alpha forest, LLS, sub-DLAs, small and large DLAs) defined in Table 4.1 according to the maximum column density system within each spectrum. We search for the highest column density integrated over any four neighbouring velocity bins; this amounts to a comoving length along the line of sight of $\{1.50, 1.42, 1.33, 1.27, 1.16\}\,\mathrm{Mpc}$ respectively at $z = [2.00, 2.44, 3.01, 3.49, 4.43]$. The categorisation is insensitive to the number of neighbouring velocity bins that we use, as the boundaries between categories differ by orders of magnitude in column density. Moreover, the method is efficient in identifying high column density absorbers since they are vastly more dense than the surrounding gas forming the Lyman-alpha forest[5]. We have chosen a length that is much larger than the most extensive DLAs as found by recent studies (Krogager et al., 2012) and so we are sure to integrate over the full length of any high column density absorbers. Our definition of high column density absorbers includes blends,

[3] We convert the comoving length of the box to a proper velocity by the Hubble law.

[4] For comparison, BOSS DR9 spectra are binned at a velocity resolution of $69.02\,\mathrm{km\,s^{-1}}$ (Lee et al., 2013).

[5] We have explicitly tested the impact of doubling or halving the number of neighbouring velocity bins we use on the 1D flux power spectra we measure in each absorber category. We find that the maximum absolute difference in any power spectrum bin is a negligible 0.2%.

where a number of smaller, lower column density systems have been added together. In this way, we have associated with each spectrum the most dominant absorbing system and in the case where high column density absorbers are identified, these are the main contamination to the spectrum through their associated damping wings. The percentage of spectra in each absorber category at each redshift slice is given in Table 4.1.

We measure the 1D flux power spectrum of all the spectra and each absorber category at each redshift slice. The 1D power spectrum $P^{\rm 1D}(k_{||}, z)$ is defined as the integral of the three-dimensional (3D) power spectrum $P^{\rm 3D}(k_{||}, \mathbf{k}_\perp, z)$ over directions perpendicular to the line of sight:

$$P^{\rm 1D}(k_{||}, z) = \int P^{\rm 3D}(k_{||}, \mathbf{k}_\perp, z) \frac{{\rm d}\mathbf{k}_\perp}{(2\pi)^2}, \tag{4.1}$$

where the wavevector $\mathbf{k} = [k_{||}, \mathbf{k}_\perp]$ is conjugate to velocities in real space and so is measured in units of inverse velocity (*e. g.*, $\mathrm{s\,km^{-1}}$). We also use the convention of absorbing the 2π into the conjugate variable[6].

To measure $P^{\rm 1D}$ for an individual line of sight, we first calculate the fluctuation in each velocity $v_{||}$ bin $\delta_{\mathcal{F}}(v_{||}) = \frac{\mathcal{F}(v_{||})}{\langle\mathcal{F}\rangle} - 1$, where $\langle\mathcal{F}\rangle$ is the average flux over all spectra at each redshift (Croft et al., 1998). We calculate the 1D Fourier transform along the line of sight $\hat{\delta}_{\mathcal{F}}(k_{||})$ using a fast Fourier transform (FFT)-based method since we have evenly-spaced velocity bins. We then estimate the 1D flux power spectrum for each sightline $P^{\rm 1D}_{\rm Raw}(k_{||}) = |\hat{\delta}_{\mathcal{F}}(k_{||})|^2$. Finally, we estimate the 1D flux power spectrum in Eq. (4.1) for each absorber category i by (*e. g.*, Palanque-Delabrouille et al., 2013)

$$P^{\rm 1D}_i(k_{||}, z) = \left\langle \frac{P^{\rm 1D}_{\rm Raw}(k_{||}, z)}{W^2(k_{||}, \Delta v, R)} \right\rangle_i, \tag{4.2}$$

where we explicitly indicate that the raw 1D power spectra depend on redshift z. The average is taken over spectra of a given category (or all spectra for the total power spectrum) at each redshift slice. The window function $W(k_{||}, \Delta v, R)$ that is divided out arises from the binning in velocity space (Δv) and the simulated spectrographic resolution R:

$$W(k_{||}, \Delta v, R) = \exp\left(-\frac{1}{2}(k_{||}R)^2\right) \times \frac{\sin\left(k_{||}\Delta v/2\right)}{k_{||}\Delta v/2}, \tag{4.3}$$

where $\Delta v = 25\,\mathrm{km\,s^{-1}}$ and $R = 3.40\,\mathrm{km\,s^{-1}}$ (not to be confused with the spectrographic resolving power; see § 4.4.1). We then have an estimate of the 1D flux power spectrum for each absorber category of spectra at each redshift slice.

[6] *I. e.*, we define the Fourier transform as $\delta(k) = \int \delta(x) e^{-ikx} {\rm d}x$.

4.4.3 Modelling the effect of high column density absorbers

The total 1D flux power spectrum of a set of spectra $P^{\rm 1D}_{\rm Total}(k_{||}, z)$ can be expressed as a weighted sum of the 1D flux power spectra calculated in Eq. (4.2) for each absorber category i:

$$P^{\rm 1D}_{\rm Total}(k_{||}, z) = \sum_i \alpha_i(z)\, P^{\rm 1D}_i(k_{||}, z), \tag{4.4}$$

where $\alpha_i(z)$ are the fraction of spectra in each absorber category at each redshift (as given in Table 4.1 for our simulated ensemble of spectra). In a real survey, $\alpha_i(z)$ may change from their raw values due to the attempt to clip (*i. e.*, remove) high column density absorbers discussed in § 4.2. We can rearrange Eq. (4.4) to isolate the 1D flux power spectrum of the Lyman-alpha forest alone:

$$P^{\rm 1D}_{\rm Total}(k_{||}, z) = P^{\rm 1D}_{\rm Forest}(k_{||}, z) \left[\alpha_{\rm Forest}(z) + \sum_{i \neq {\rm Forest}} \alpha_i(z) \frac{P^{\rm 1D}_i(k_{||}, z)}{P^{\rm 1D}_{\rm Forest}(k_{||}, z)} \right]. \tag{4.5}$$

In this way, we have isolated the effect of spectra containing high column density absorbers on the 1D flux power spectrum of the Lyman-alpha forest as a multiplicative bias (*i. e.*, the terms in square brackets)[7]. This matches the general form of modelling this effect in previous studies, as explained in Palanque-Delabrouille et al. (2015) (based on the results in McDonald et al. 2005a), but now additionally probing the bias as a function of column density (*i. e.*, by using the different absorber categories). We discuss in more detail in § 4.6 our motivations for using this particular form of the bias (as opposed to *e. g.*, an additive bias). Using the 1D flux power spectra we have calculated in § 4.4.2, we are able to measure the fractions in Eq. (4.5) ($P^{\rm 1D}_i(k_{||}, z)/P^{\rm 1D}_{\rm Forest}(k_{||}, z)$) and we present the results in § 4.5.

4.5 Results

Figure 4.1 shows the 1D flux power spectra of different subsets of sightlines that we have measured from our simulations [see § 4.4.2 and in particular Eq. (4.2)] at redshift $z = 2.00$. The different subsets shown are: the total as would be measured if no distinction between different types of spectra was made; spectra containing only Lyman-alpha forest (*i. e.*, the ensemble that is uncontaminated by high column density absorbers); and spectra contaminated by different categories of high column density absorber, as defined in Table 4.1. We first note that the total 1D flux power spectrum at

[7] We could simplify this form further by asserting the fact that $\sum_i \alpha_i(z) = 1$ to remove the parameter $\alpha_{\rm Forest}(z)$, but it is useful to keep this form as we explain in § 4.7.

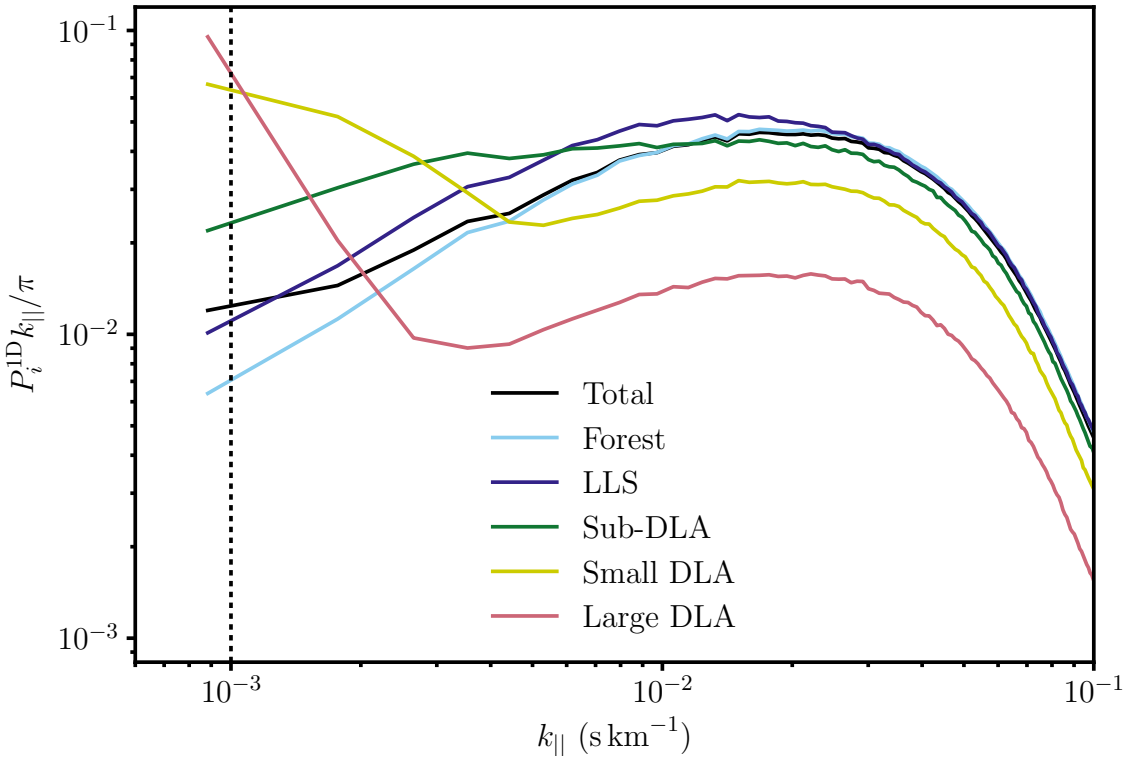

Figure 4.1: The one-dimensional flux power spectra of different categories of spectra, as a function of line-of-sight scale $k_{||}$ at redshift $z = 2.00$. The different categories are: the total from our full simulated sample of spectra; spectra containing only Lyman-alpha forest; and spectra contaminated by different types of high column density absorber [LLS, sub-DLAs, small and large DLAs]. The vertical dashed line shows the largest scale probed by the BOSS DR9 1D Lyman-alpha forest flux power spectrum; by comparison, the largest scale probed by our analysis at this redshift is larger at $9 \times 10^{-4}\,\mathrm{s\,km^{-1}}$. The definitions of the different categories of absorber are given in Table 4.1. (See § 4.7 for the full intermediate redshift evolution.)

any redshift slice can be reconstructed as a weighted sum of the other 1D flux power spectra for each absorber category at that redshift (see § 4.4.3 and in particular Eq. (4.4)). The weights are the fraction of spectra in each category (the values we measure for our simulated ensemble are given in Table 4.1). We can estimate the fractional (1σ) statistical error on each power spectrum data-point as $1/\sqrt{N_i}$, where N_i is the number of input modes (*i. e.*, simulated spectra) per data-point i. This assumes that data-points and input modes are independent. This is largest for the large DLA power spectrum at $z = 2.00$, which has 15,188 input simulated spectra giving an error of 0.81%, and smallest for the forest power spectrum at $z = 2.00$, which has 437,063 input simulated spectra, giving an error of 0.15%. All the other uncertainties for each measured power spectrum range in-between these values and can be computed from Table 4.1.

The total power spectrum deviates from the Lyman-alpha forest power spectrum at all redshifts, showing there is a bias from contamination of spectra by high column density absorbers. This bias can be deconstructed as a function of column density by looking at the power spectra of different absorber categories. The power spectra of high column density absorbers have a distinctive increase on large scales (small $k_{||}$). This is caused by self-correlations across the width of damping wings, which (as discussed in § 4.3) can be modelled by a Voigt profile (a convolution of a Gaussian and a Lorentzian). Therefore, the power spectrum of high column density absorbers (on large scales) is connected to the Fourier transform of a Voigt profile. This increases for higher column density systems since there is more line broadening, and starts on larger scales for higher column density systems since the damping wings are wider. (See Appendix A.1 for more analysis and discussion of the Voigt profile model.) On small scales, all the power spectra converge to a scaled version of the Lyman-alpha forest flux power spectrum. This reflects the fact that contaminated spectra do contain some uncontaminated spectral pixels. The amplitude of the small-scale power spectrum reflects the fraction of spectra that is uncontaminated, increasing for lower-column density systems since their damping wings are narrower. There is some sensitivity to the length of our simulated spectra, which primarily manifests in our results as the amplitude of the small-scale residual Lyman-alpha forest power in the contaminated power spectra. This is because longer simulated contaminated spectra would have a larger fraction with residual Lyman-alpha forest. This is discussed further and explicitly modelled such that this effect is removed in § 4.7.

Figure 4.2 shows 1D flux power spectra as in Fig. 4.1, but for more of the redshift slices that we consider ($z = [2.00, 2.44, 3.49, 4.43]$), for spectra containing only Lyman-alpha forest and spectra contaminated by large DLAs. The Lyman-alpha forest flux power spectrum has the expected shape, amplitude and redshift evolution, matching observations (*e. g.*, Palanque-Delabrouille et al., 2013)

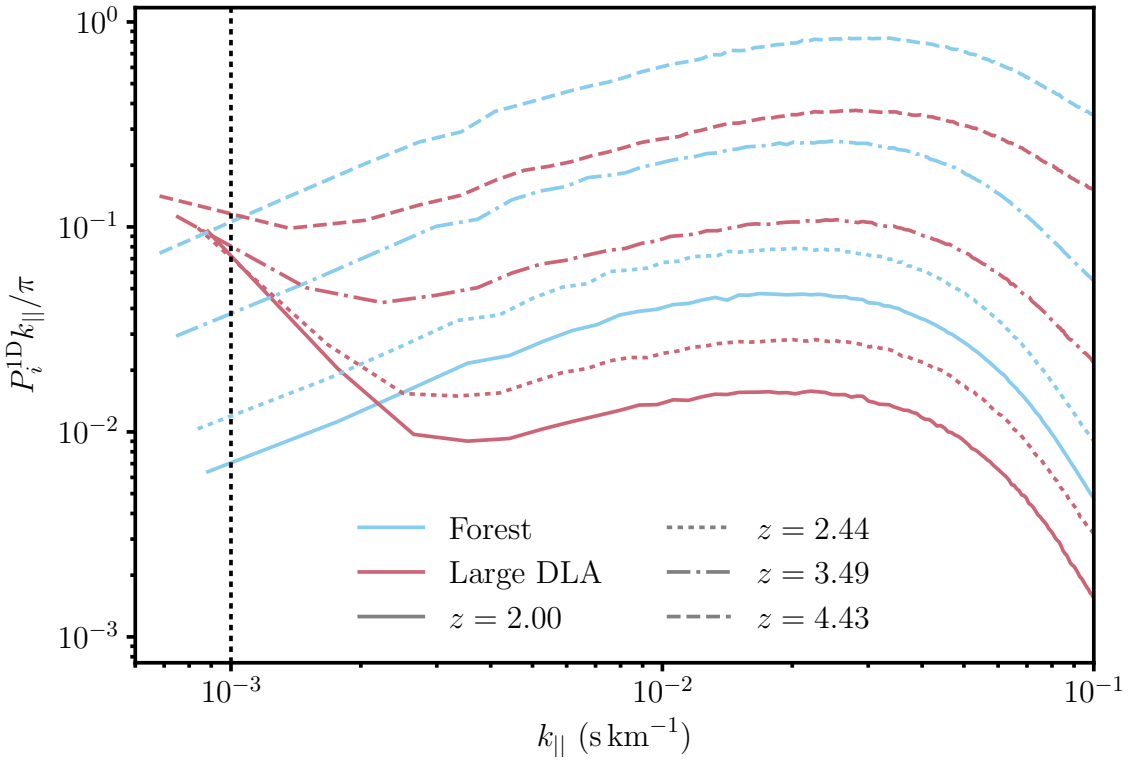

Figure 4.2: As Fig. 4.1, but showing more of the redshift slices that we consider (for $z = [2.00, 2.44, 3.49, 4.43]$), for spectra containing only Lyman-alpha forest and spectra contaminated by large DLAs.

and reflecting the fact that it is an integral of a (biased) matter power spectrum. A peculiarity of the Lyman-alpha forest flux power spectrum is that its amplitude *increases* with redshift (unlike the linear matter power spectrum); this is because neutral hydrogen is more abundant at higher redshift and so there is more absorption in quasar spectra (*i. e.*, the Lyman-alpha forest becomes a more negatively biased tracer of the matter distribution). By contrast, it can be seen that the large-scale correlations associated with the large DLAs are converging to a single point as redshift changes. This reflects the fact that these correlations arise from individual damping wings, which do not evolve with redshift.

Figure 4.3 shows the same 1D flux power spectra as in Figs. 4.1 and 4.2, but now as ratios between the flux power of spectra contaminated by high column density absorbers and the flux power of spectra containing only Lyman-alpha forest, for $z = 2.00$ and $z = 4.43$. These ratios are the quantities to which we fit our templates (see § 4.7) as part of our bias model (see § 4.4.3). Plotted in this form, it is clear that the large-scale corrections associated with damping wings increase with column density of the damped system. The corrections also decrease with increasing redshift because the Lyman-alpha forest flux power spectrum (on the denominator of the ratio) increases with redshift. On small scales, the ratios converge to a constant value, which reflects the fraction of a line of sight that is uncontaminated (see above). The redshift evolution of this constant value is driven by the transformation from comoving to velocity space: spectra are longer in velocity space at higher redshift (despite being drawn from the same comoving length of the simulation). Conversely, the width of damping wings does not change with redshift (for a given column density) because this width just arises from the physical processes within the hydrogen gas (rather than cosmological evolution). Therefore, the fraction of spectra uncontaminated by the damping wings increases with redshift.

4.6 Discussion

We first discuss and summarise the results we have presented in § 4.5. Using our measurements from cosmological hydrodynamical simulations, we have been able to confirm and characterise the effect of high column density absorbers on the 1D Lyman-alpha forest flux power spectrum as a function of column density, scale and redshift. There are distinctive large-scale correlations across the widths of individual damping wings (a “one-halo” term) arising from high column density absorbers that are seen to bias the 1D flux power spectrum of a set of spectra, relative to the power spectrum of the Lyman-alpha forest alone (Fig. 4.1). These correlations persist for all the high column density

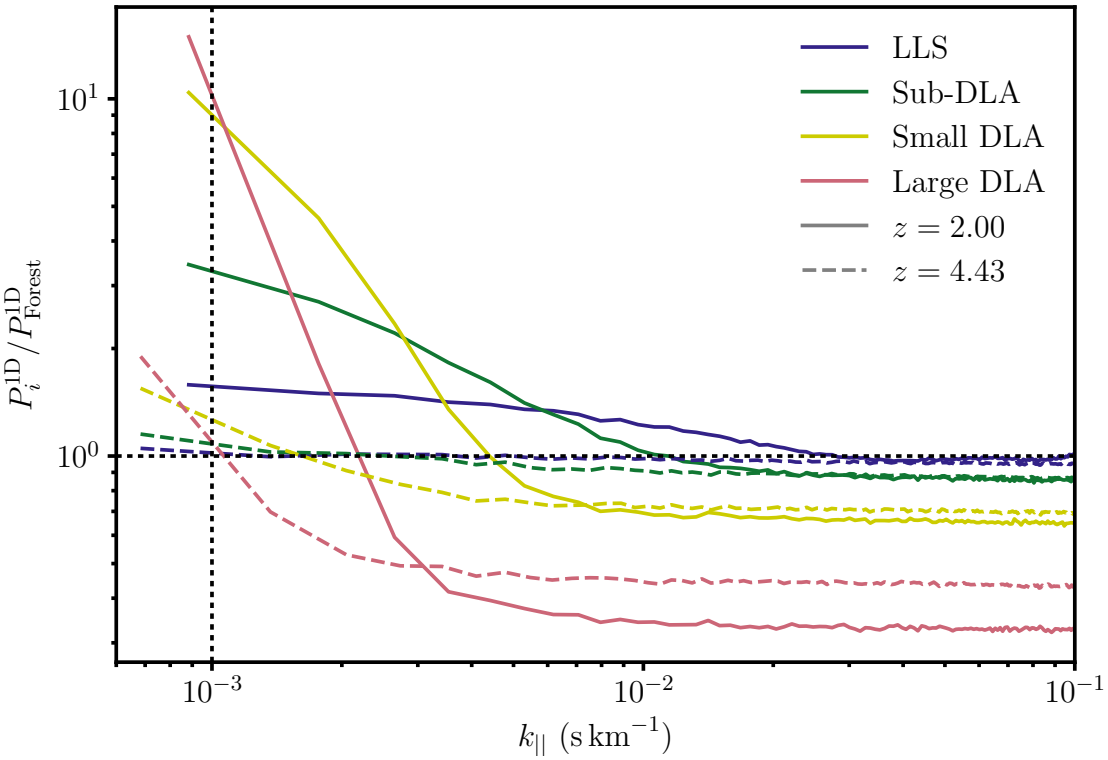

Figure 4.3: The multiplicative bias of high column density absorbers to the one-dimensional Lyman-alpha forest flux power spectrum, as a function of line-of-sight scale $k_{||}$ and redshift z, *i. e.*, the ratio of the 1D flux power spectrum of spectra contaminated by high column density absorbers [LLS, sub-DLAs, small and large DLAs] over spectra containing only Lyman-alpha forest. The vertical dashed line shows the largest scale probed by the BOSS DR9 1D Lyman-alpha forest flux power spectrum. The definitions of the different categories of high column density absorber are given in Table 4.1. The different line styles correspond to different redshift slices, showing the maximum and minimum redshifts that we consider. (See § 4.7 for the full intermediate redshift evolution.)

absorbing systems that we identify (*i. e.*, for all column densities $N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$). Our results can be further understood by relating the shape and amplitude of the large-scale power spectrum of spectra contaminated by high column density absorbers to the Fourier transform of the Voigt profile that is normally used to model damping wings (due to the combination of physical effects that broaden absorption lines; see Appendix A.1). We find evidence in our simulation results that the 1D flux power spectrum of high column density absorbers does not evolve with redshift (Fig. 4.2). This reflects the fact that the Voigt profiles of damping wings depend only on column density (*i. e.*, the physical processes within high column density absorbing regions) and not redshift (*i. e.*, cosmological evolution) (see Eq. (A.1)).

The most recent previous investigation into the effect of high column density absorbers on the Lyman-alpha forest was performed by McDonald et al. (2005a) (see also Croft et al., 1999; Viel et al., 2004a). These authors measured a single bias function for the 1D Lyman-alpha forest flux power spectrum (at each redshift they consider) that includes the combined effect of all high column density absorbers (*i. e.*, all LLS and DLAs). Our results are qualitatively similar to those of the previous study; however, by investigating different absorber categories based on column density ranges (Table 4.1), we have shown that the form of the bias as a function of wavenumber depends strongly on column density.

This will have implications for any parameter inference from the 1D flux power spectrum. For instance, the analysis by Palanque-Delabrouille et al. (2015) uses a single multiplicative bias model for the Lyman-alpha forest flux power spectrum based on the results in McDonald et al. (2005a)[8]. The model has a free amplitude that is allowed to vary (reflecting the level of contamination in a given survey) and is then marginalised. The shape of this model is therefore based on the observed CDDF of high column density absorbers. However, as discussed in § 4.2, in the measurement of the 1D flux power spectrum, high column density absorbers in the quasar spectra are clipped out in the hope of removing noise (Lee et al., 2013; Palanque-Delabrouille et al., 2013). This process changes the CDDF of high column density absorbers by preferentially removing higher column density systems which are easier to spot in the noisy spectra. Hence, the shape of the bias from residual high column density absorbers is different (as we have shown in § 4.5) and the model used by Palanque-Delabrouille et al. (2013) may not be flexible enough to account for this, especially at the level of accuracy required by future surveys. Our measurements provide a set of templates for the effect of different absorber categories as a function of column density. By using our templates

[8]The model is reported in Palanque-Delabrouille et al. (2015) as $1 - 0.2\,\alpha_{\mathrm{DLA}}\,[1/(15000.0\,k_{||} - 8.9) + 0.018]$, where α_{DLA} is the free amplitude.

as part of the model in Eq. (4.5), it is now possible to more accurately characterise the bias of the residual contamination. We also find evidence for redshift dependence of the fractional effect of high column density absorbers on the forest power spectrum (driven by the changing amplitude of the forest power spectrum), which is also not included in the current model, but is incorporated into our templates. Fits allowing incorporation of our new results into future pipelines are described in § 4.7.

We now discuss our motivations for some of the choices made in our analysis. We have chosen to present our main results as the 1D flux power spectra of different sets of simulated spectra, where we have categorised spectra according to the maximum column density system within each spectrum. This means that we are measuring the power spectra of ensembles of spectra that are contaminated to similar extents, rather than the flux power spectra of high column density absorbers alone. Furthermore, a consequence of this categorisation is that within the spectra of a given category, there may be high column density absorbers of a lower column density (*e. g.*, there may be LLS in the large-DLA category of spectra). In the first instance, this does not affect our results because the power spectrum measurements we have made (§ 4.5) and the templates that we construct (§ 4.7) include the effect of this possible additional lower column density contamination. A subtlety arises because the amount of additional lower column density contamination will be partly sensitive to the length of simulated spectra, since longer spectra have a greater chance of being contaminated. However, the damping wings of the highest column density systems already produce zero transmitted flux over a significant fraction of the length of our simulation box, so that the presence of possible additional high column density absorbers will make very little difference in any case. We tested this by inserting an LLS into a spectrum contaminated by a large DLA, which reduced the total transmitted flux by 0.07%. By carrying out this insertion test with a "control" scenario without the additional contamination, we are able to show that this subtlety will have negligible impact on our conclusions and the validity of our templates.

Finally, we comment on the particular form of our preferred model for the bias of high column density absorbers to the 1D Lyman-alpha forest flux power spectrum (as shown in Eq. 4.5). We model the bias as a multiplicative correction, rather than *e. g.*, an additive form. First, this matches the form of the currently-used model (as shown in Palanque-Delabrouille et al., 2015). Moreover, an additive form would require either the separation of high column density absorbers and the Lyman-alpha forest in the simulated spectra or a complete physical understanding of how the two components interact at the ensemble level. The former is not trivial for our analysis since we are not inserting high column density absorbers (as previous studies have done), but are simultaneously

Table 4.2: Best-fit values of the parameters in our templates for the bias of spectra contaminated by high column density absorbers on the one-dimensional Lyman-alpha forest flux power spectrum. The template parameters are defined in Eqs. (4.6) and (4.7). Values are shown for each high column density absorber category. The definitions of the different categories of high column density absorber are given in Table 4.1.

Absorber category	Template parameter values					
	a_0	a_1	b_0	b_1	c_0	c_1
LLS	2.2001	0.0134	36.449	-0.0674	0.9849	-0.0631
Sub-DLA	1.5083	0.0994	81.388	-0.2287	0.8667	0.0196
Small DLA	1.1415	0.0937	162.95	0.0126	0.6572	0.1169
Large DLA	0.8633	0.2943	429.58	-0.4964	0.3339	0.4653

simulating the low and high column density regions of gas. We avoid the latter due to any remaining physical uncertainties and instead form a parametric multiplicative model based on our simulated results (see § 4.7).

4.7 Templates for the effect of high column density absorbers

To aid incorporation in future pipelines, we have produced fits to the biases induced by contaminants in our different column density bins. The parametric form of our templates is

$$\frac{P_i^{\rm 1D}(k_{||},z)}{P_{\rm Forest}^{\rm 1D}(k_{||},z)} = \left(\frac{1+z}{1+z_0}\right)^{-3.55} \times \frac{1}{(a(z)e^{b(z)k_{||}}-1)^2} + c(z), \tag{4.6}$$

where

$$a(z) = a_0\left(\frac{1+z}{1+z_0}\right)^{a_1}; \quad b(z) = b_0\left(\frac{1+z}{1+z_0}\right)^{b_1}; \quad c(z) = c_0\left(\frac{1+z}{1+z_0}\right)^{c_1}; \tag{4.7}$$

and the pivot redshift $z_0 = 2.00$. $[a_0, a_1, b_0, b_1, c_0, c_1]$ are free parameters that we fit simultaneously in $k_{||}$ and z space for each absorber category $i \in \{$LLS, sub-DLA, small DLA, large DLA$\}$. We fit using the Levenberg-Marquardt algorithm (Levenberg, 1944; Marquardt, 1963)[9].

Figure 4.4 shows the result of these fits (dashed lines) with the raw ratios measured from the simulation (solid lines); the corresponding parameter values are given in Table 4.2. These can be used to reconstruct a final model for the bias of spectra containing high column density absorbers by using Eq. (4.5). The model described by Eq. (4.6) characterises the results we have measured in our simulations and through Eq. (4.7) allows interpolation of our results to intermediate redshifts

[9]We were able to further validate our modelling by initially fitting using a subset of the available redshift slices and using this preliminary fit to predict the results at $z = 3.01$. We found the model to accurately predict the results at this intermediate redshift, acting as a form of successful blind test for our model. Our final best-fit parameters use all available data.

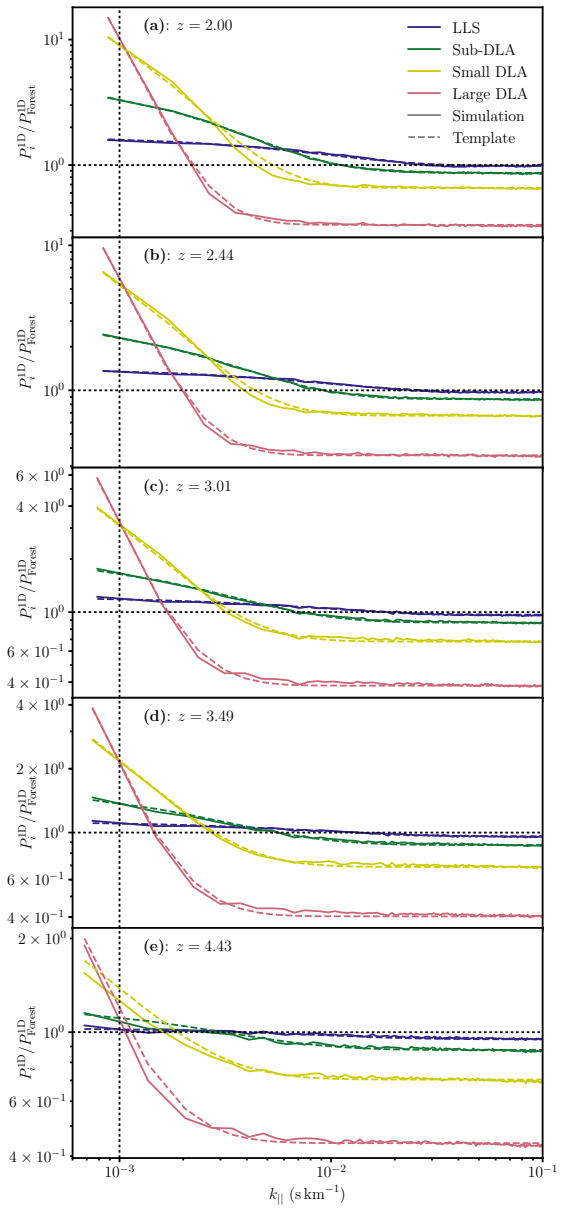

Figure 4.4: The multiplicative bias of high column density absorbers to the one-dimensional Lyman-alpha forest flux power spectrum, as a function of line-of-sight scale $k_{||}$ and redshift z, *i. e.*, the ratio of the 1D flux power spectrum of spectra contaminated by high column density absorbers [LLS, sub-DLAs, small and large DLAs] over spectra containing only Lyman-alpha forest. The solid lines are these ratios as measured in the hydrodynamical simulations; the dashed lines are our best-fitting templates to these measurements. The functional form of our templates is given in Eq. (4.6) and the best-fit values of the model parameters are given in Table 4.2. The vertical dashed lines show the largest scale probed by the BOSS DR9 1D Lyman-alpha forest flux power spectrum. The definitions of the different categories of high column density absorber are given in Table 4.1. *From top to bottom*, we show the templates for simulated results at increasing redshift [(a): $z = 2.00$; (b): $z = 2.44$; (c): $z = 3.01$; (d): $z = 3.49$; (e): $z = 4.43$].

that we have not explicitly probed. (Use of the model outside the limits of redshift and scale we have considered would constitute an extrapolation, but this should not be necessary since our measurements bracket the main redshifts and scales of interest to Lyman-alpha forest studies.) No strong physical meaning should be attached to its terms, although we can motivate the first term on the right-hand side of Eq. (4.6) as being the (reciprocal of the) main term of the redshift evolution of $P^{\rm 1D}_{\rm Forest}(k_{||}, z)$ as found by Palanque-Delabrouille et al. (2013) (using a maximum likelihood estimator). In this way, the parametric form isolates the redshift evolution from $P^{\rm 1D}_{\rm Forest}(k_{||}, z)$ and then fits the residual redshift evolution using the terms in Eq. (4.7). The best-fit values of the exponents in Eq. (4.7) (as given in Table 4.2) are small, indicating that most of the redshift evolution can indeed be ascribed to the expected cosmological evolution of $P^{\rm 1D}_{\rm Forest}(k_{||}, z)$.

Our results are dependent on the length of our simulated spectra. This manifests in the value of the constant that the ratios $P^{\rm 1D}_{i}(k_{||}, z)/P^{\rm 1D}_{\rm Forest}(k_{||}, z)$ have at high $k_{||}$, which is set by the fraction of the length of contaminated spectra which are unaffected by damping wings and contain only Lyman-alpha forest. Since the incidence rates of high column density absorbers are such that one per contaminated spectrum is most likely, a longer spectrum will have a larger fraction that is uncontaminated, causing the constant value at high $k_{||}$ to rise with spectrum length. However, in an analysis of observational data this will be absorbed into a free parameter. We have used a parametric form for our templates such that all this dependency is measured by the term $c(z)$[10]. By inserting Eq. (4.6) into Eq. (4.5), it can be seen that the term $c(z)$ is degenerate with $\alpha_{\rm Forest}(z)$ and hence these terms can be combined and allowed to vary. It follows that the full parametric form of our model for the effect of high column density absorbers on the 1D Lyman-alpha forest flux power spectrum is

$$P^{\rm 1D}_{\rm Total}(k_{||}, z) = P^{\rm 1D}_{\rm Forest}(k_{||}, z) \left[\alpha_0(z) + \sum_{i \neq \rm Forest} \alpha_i(z) \left(\frac{1+z}{1+z_0}\right)^{-3.55} \times \frac{1}{(a(z)e^{b(z)k_{||}} - 1)^2}\right]. \tag{4.8}$$

When using this model in inference from the 1D Lyman-alpha forest power spectrum $P^{\rm 1D}_{\rm Forest}(k_{||}, z)$, it will be necessary to vary five free parameters α_0 and α_i, where i indexes each high column density absorber category. In this way, the column density, scale and redshift dependence of the effect of high column density absorbers is fully determined by our templates, while the relative impact of each absorber category is fitted since this is specific to the survey at hand, as well as the details of

[10] It can then be understood why we do not factor out the redshift evolution of $P^{\rm 1D}_{\rm Forest}(k_{||}, z)$, as we do for the first term on the right-hand side of Eq. (4.6).

any clipping of damping wings that changes the survey CDDF. (See § 4.6 for more discussion of these details.) Note that the parameter α_0 is degenerate with factors that rescale the mean flux and could be omitted in an end-to-end analysis.

Figure 4.5 compares the model we have constructed to the existing model presented in Palanque-Delabrouille et al. (2015) and based on the results in McDonald et al. (2005a). There is broad agreement between the existing model and our model for the total contamination of high column density absorbers, although our model is less steep in its scale dependence. We also show our model applied to a possible "residual" contamination, *i. e.*, under the assumption that all DLAs are identified and clipped out in an analysis, leaving only contamination from LLS and sub-DLAs (*e. g.*, as assumed by Bautista et al. 2017). The model for this lower column density residual contamination has a shallow scale dependence that the model of McDonald et al. (2005a) is unable to characterise. The use of our more flexible model will avoid potential biases due to mischaracterisation of the scale dependence of the residual contamination, thus improving estimation of cosmological effects such as massive neutrinos or the tilt of the primordial power spectrum.

We now discuss the prior probability distributions that can be adopted for $\alpha_i(z)$ in any inference using the model we have presented. The $\alpha_i(z)$ are technically not independent parameters, but are each related to integrals of the HI CDDF for a particular survey over the appropriate column density ranges (and absorption distance per sightline). The effect of spectrum clipping which changes the survey CDDF can be modelled by applying a weighting function to the CDDF, which down-weights higher column densities, which are easier to spot and remove. If one wanted to reduce the dimensionality of these nuisance parameters, in particular in redshift space, they could be replaced by a parameterisation which quantifies deviations from the expected redshift evolution of the CDDF with only one or two parameters (rather than a parameter for each redshift bin considered). We leave the details of the construction of prior distributions to individual analyses, since the precise considerations will be survey-specific.

To conclude this section, we present a summary of the steps required to incorporate our final model for the effect of high column denisty absorbers into future 1D Lyman-alpha forest analyses:

- Our model describes the effect of quasar spectra contaminated by high column density absorbers as a multiplicative bias to the 1D Lyman-alpha forest flux power spectrum, as given by Eq. (4.8). It can therefore be incorporated into a pipeline at the stage of flux power spectrum interpretation to marginalise over effects of these absorbers.

- The free parameters are $\alpha_i(z)$, where i indexes different categories of high column density

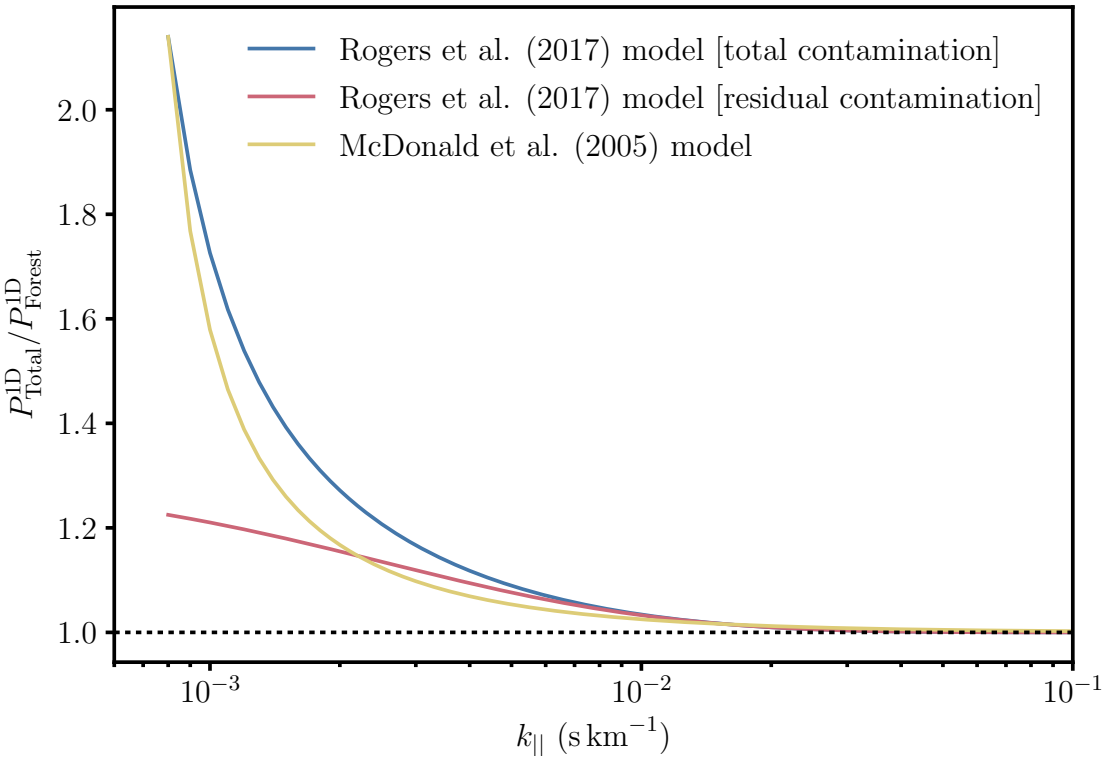

Figure 4.5: A comparison of the existing multiplicative bias model (McDonald et al., 2005a; Palanque-Delabrouille et al., 2015) for the effect of high column density absorbers on the 1D Lyman-alpha forest power spectrum and the model constructed in this paper using our results from hydrodynamical simulations. For our model, we show an example weighting of the different absorber categories for the full contamination from high column density absorbers on our simulated ensemble of spectra; and an example based on a possible "residual" contamination after the clipping of DLAs (*i. e.*, only LLS and sub-DLAs remaining). For comparison, the model of McDonald et al. (2005a) is rescaled to have the same amplitude on the largest and smallest scales considered.

absorber (as given in Table 4.1). Our model is of use to any Lyman-alpha forest survey that contains spectra which may be contaminated by high column density absorbers (both LLS and DLAs). The relative impacts of different categories of high column density absorbers will be determined in the estimation of posterior distributions of these nuisance parameters. While normalisation is necessarily floating, the model fully specifies the scale, column density and redshift dependence of the effect of high column density absorbers, using the results we have measured from hydrodynamical simulations.

- In a survey which does not clip its quasar spectra, strong priors can be given for the free parameters of our model, based on the expected or measured HI CDDF.
- In a survey which does clip its quasar spectra in an attempt to remove high column density absorbers (and therefore changes the survey CDDF), strong priors can still be given for our model parameters, assuming a model can be constructed for the effect of the clipping process on the CDDF. This will constitute some re-weighting of the CDDF.
- In order to reduce the dimensionality of our nuisance parameters, rather than having a separate parameter for each redshift bin in a given analysis, one could parameterise the redshift evolution by a simple deviation from the CDDF with only one or two numbers.

4.8 Conclusions

We have used a cosmological hydrodynamical simulation (Illustris; Nelson et al., 2015; Vogelsberger et al., 2014a) to investigate the effect of high column density absorbing systems of neutral hydrogen and their associated damping wings on the 1D Lyman-alpha forest flux power spectrum. We find that the effect of high column density absorbers on the Lyman-alpha forest flux power spectrum is a strong function of column density. Accounting for this change in scale-dependence with column density will remove a source of bias in cosmological inference from the Lyman-alpha forest. Previous models (Palanque-Delabrouille et al., 2015) combine the effect of all high column density absorbers together (*i. e.*, all neutral hydrogen column densities $N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$) based on the column density distribution function (CDDF) in the raw spectra (McDonald et al., 2006). However, the damping wings of some high column density absorbers are clipped out in the final analysis (Lee et al., 2013), which preferentially removes higher density systems (because they are easier to spot) and changes the column density distribution in the residual contamination. Our results apply for both clipped and unclipped survey spectra, since we separately model the effect for different

column densities of the dominant absorber, allowing us to accurately account for the contamination in the 1D flux power spectrum. We discuss in § 4.7 the practicalities of employing our model in future analyses.

The shape and amplitude of the distortions in the power spectrum due to a damped absorber depend on its column density because they are driven by the width of the damping wings; *i. e.*, the dominant effect is a "one-halo" term. We defer investigation of potential "two-halo" terms to future work where we measure the effect of high column density absorbers on the 3D Lyman-alpha forest flux power spectrum.

We anticipate that our model will help realise forecasted cosmological constraints from upcoming surveys like DESI. E. g., Font-Ribera et al. (2014b) forecast that DESI will have the constraining power to make a $\sim$ three-sigma detection of the sum of neutrino masses (in combination with *Planck* CMB data); and they show the power of the 1D Lyman-alpha forest power spectrum in probing the primordial power spectrum, *e. g.*, halving the one-sigma error on the running of the spectral index, with implications for inflationary models. It will be necessary to use the models we have presented here, alongside carefully constructed priors on the residual CDDF, to remove degeneracies between the effect of high column density absorbers and cosmological effects.

Acknowledgements

KKR thanks the organisers of the COSMO21 symposium in 2016, where this project was conceived.

5

Correlations in the three-dimensional Lyman-alpha forest contaminated by high column density absorbers

5.1 Abstract

Correlations measured in three dimensions (3D) in the Lyman-alpha forest are contaminated by the presence of the damping wings of high column density (HCD) absorbing systems of neutral hydrogen (HI; having column densities $N(\mathrm{HI}) > 1.6 \times 10^{17}\,\mathrm{atoms\,cm^{-2}}$), which extend significantly beyond the redshift-space location of the absorber. We measure this effect as a function of the column density of the HCD absorbers and redshift by measuring 3D flux power spectra in cosmological hydrodynamical simulations from the Illustris project. Survey pipelines exclude regions containing the largest damping wings. We find that, even after this procedure, there is a scale-dependent correction to the 3D Lyman-alpha forest flux power spectrum from residual contamination. We model this residual using a simple physical model of the HCD absorbers as linearly biased tracers of the matter density distribution, convolved with their Voigt profiles and integrated over the column density distribution function. We recommend the use of this model over existing models used in data analysis, which approximate the damping wings as top-hats and so miss shape information in the extended wings. The simple "linear Voigt model" is statistically consistent with our simulation results for a mock residual contamination up to small scales ($|\mathbf{k}| < 1\,h\,\mathrm{Mpc}^{-1}$), even though it cannot account for the effect of the highest column density absorbers (which are in any case preferentially removed from survey data) on the smallest scales (*e. g.*, $|\mathbf{k}| > 0.4\,h\,\mathrm{Mpc}^{-1}$ for small DLAs; $N(\mathrm{HI}) \sim 10^{21}\,\mathrm{atoms\,cm^{-2}}$). Our model is appropriate for an accurate analysis of the baryon acoustic oscillations (BAO) feature and it is additionally essential for reconstructing the full shape of the 3D flux power spectrum, assuming that the highest column density absorbers are removed.

5.2 Introduction

Absorption lines of the Lyman-alpha forest can be mapped in three dimensions (3D) (*i. e.*, line-of-sight direction along the lengths of quasar spectra and transverse direction in the angular positions of the spectra on the sky) to trace the fluctuations in the cosmological density field. Correlations in the Lyman-alpha forest are a powerful probe of high redshifts ($z > 2$), before dark energy came to dominate the evolution of the Universe. In particular, measurement of the 3D correlations on large scales (separations $r \sim 100\,\mathrm{Mpc}\,h^{-1}$) in the Lyman-alpha forest allows a measurement of the baryon acoustic oscillations (BAO) in the distribution of matter at $z \sim 2.3$ (Bautista et al., 2017; Busca et al., 2013; Delubac et al., 2015; Kirkby et al., 2013; Slosar et al., 2011, 2013). 3D correlations between the Lyman-alpha forest and the distribution of quasars have also been measured, including the detection of BAO (du Mas des Bourboux et al., 2017; Font-Ribera et al., 2013, 2014a). This has been achieved thanks to the large number of quasar spectra from the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al., 2013; Eisenstein et al., 2011) (157,783 were suitable for analysis in Data Release 12; DR12) and the large sky area they cover (the footprint in DR12 covers approximately one quarter of the sky). Consequently, Lyman-alpha forest analyses are no longer restricted to measurements of the one-dimensional flux power spectrum (along the line-of-sight only; Armengaud et al., 2017; Iršič et al., 2017a,b; Palanque-Delabrouille et al., 2015; Seljak et al., 2005; Yeche et al., 2017), which probes smaller-scale clustering ($k_{||} > 0.1\,h\,\mathrm{Mpc}^{-1}$) and constrains cosmological models that suppress small-scale power, *e. g.*, those containing massive neutrinos or warm dark matter (WDM).

Current measurements of the 3D correlations in the Lyman-alpha forest reconstruct the correlation function, where the BAO feature is most distinguishable. However, ongoing analyses in the extended Baryon Oscillation Sky Survey (eBOSS; Abolfathi et al., 2017) and future surveys like the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al., 2016a,b) will also measure its Fourier-space counterpart, the 3D flux power spectrum (Font-Ribera et al., 2017). A measurement of the 3D Lyman-alpha forest power spectrum will probe the full shape on a wide range of scales ($0.01\,h\,\mathrm{Mpc}^{-1} < k < 1\,h\,\mathrm{Mpc}^{-1}$). On large scales ($k < 0.1\,h\,\mathrm{Mpc}^{-1}$), *e. g.*, the 3D flux power spectrum can be used to determine the cosmological geometry through the Alcock-Paczyński test (Alcock and Paczynski, 1979; Hui et al., 1999; McDonald, 2003; McDonald and Miralda-Escudé, 1999). The 3D forest power spectrum on large scales can also be used to study fluctuations in the ultraviolet (UV) ionising background (Pontzen, 2014; Pontzen et al., 2014).

On smaller scales ($0.1\,h\,\mathrm{Mpc}^{-1} < k < 1\,h\,\mathrm{Mpc}^{-1}$), the 3D flux power spectrum adds comple-

mentary information to that from the 1D flux power spectrum. Font-Ribera et al. (2017) show that the 3D flux power spectrum for the BOSS survey is more constraining than the 1D counterpart up to a maximum $k = 1\,h\,\mathrm{Mpc}^{-1}$. Only for $k > 1\,h\,\mathrm{Mpc}^{-1}$ does the 1D power spectrum contain essentially all information. For future surveys such as DESI, where there will be a higher density of lines of sight, one may anticipate 3D information to even higher k, underscoring the importance of working with the 3D spectrum wherever possible. This will provide more power to constrain cosmological models with additional components (*e. g.*, massive neutrinos, WDM or fuzzy dark matter), or modifications to a simple power-law primordial power spectrum (*e. g.*, running of the primordial spectral index). In addition to providing greater statistical power, the 3D flux power spectrum is sensitive to different systematics than the 1D flux power spectrum (*e. g.*, in correlations with metal absorption lines). (See *e. g.*, Font-Ribera et al. 2014b for forecasts of the constraining power of the 3D flux power spectrum with DESI.)

As with the 1D Lyman-alpha forest flux power spectrum (McDonald et al., 2005a; Rogers et al., 2017), 3D correlations in the Lyman-alpha forest are biased by the presence in quasar spectra of high column density (HCD) absorbers and their associated broadened absorption lines (Font-Ribera and Miralda-Escudé, 2012; McQuinn and White, 2011; Slosar et al., 2011). HCD absorbers are defined as regions of neutral hydrogen (HI) gas with a column density $N(\mathrm{HI})$ exceeding $1.6 \times 10^{17}\,\mathrm{atoms\,cm}^{-2}$, and are usually identified with the gas in or around galaxies. They form at the peaks of the underlying density distribution and so cluster more strongly than the Lyman-alpha forest (Font-Ribera et al., 2012b; Pérez-Ràfols et al., 2017). The absorption lines of the highest column density systems are broadened, with large damping wings causing absorption in the spectrum away from the physical location of the absorber. These wings have a characteristic Voigt profile, a convolution of a Gaussian profile (caused by Doppler broadening) and a Lorentzian profile (caused by natural or collisional broadening). They are traditionally sub-classified as either damped Lyman-alpha absorbers (DLAs; $N(\mathrm{HI}) > 2 \times 10^{20}\,\mathrm{atoms\,cm}^{-2}$) or Lyman-limit systems (LLS; $1.6 \times 10^{17}\,\mathrm{atoms\,cm}^{-2} < N(\mathrm{HI}) < 2 \times 10^{20}\,\mathrm{atoms\,cm}^{-2}$), according to the width of their damping wings (Wolfe et al., 1986). However, as noted in *e. g.*, Font-Ribera and Miralda-Escudé (2012); McDonald et al. (2005a); Rogers et al. (2017), systems with $N(\mathrm{HI})$ exceeding $1 \times 10^{19}\,\mathrm{atoms\,cm}^{-2}$ have significant wings, which we classify as sub-DLAs. In Lyman-alpha forest analyses, it is usual to attempt to "clip" out HCD absorbers by identifying their damping wings in spectra, masking the central absorption region and then correcting the wings (*e. g.*, see Lee et al. 2013 for details of the process for BOSS DR9 spectra). Nonetheless, there is always a residual contamination of HCD absorbers, since the smallest damping wings are hard to identify amongst instrumental noise and

indeed the superposed Lyman-alpha forest itself. Estimates of the upper limit in column density for this residual contamination range from 10^{20} to $10^{21}\,\mathrm{atoms\,cm^{-2}}$ (*e. g.*, Bautista et al., 2017). It is therefore necessary to model the effect of this residual contamination to allow for robust cosmological inference from the Lyman-alpha forest (Bautista et al. 2017; du Mas des Bourboux et al. 2017 were the first to model this component in a 3D correlation analysis).

There is a small literature on modelling the effect of (residual) HCD absorbers on correlations in the 3D Lyman-alpha forest. In Appendix B of McQuinn and White (2011), a linear model for the 3D flux power spectrum of HCD absorbers convolved with their (Voigt) absorption profiles is considered, allowing for their auto-correlation and cross-correlation with the Lyman-alpha forest. They show that this "linear Voigt model" predicts that the cross-correlation is the dominant component of the HCD absorbers' correction to the 3D Lyman-alpha forest power spectrum. Font-Ribera and Miralda-Escudé (2012) measure the effect of HCD absorbers on the 3D Lyman-alpha forest correlation function using mock (quasar) spectra (details of their generation are given in Font-Ribera et al. 2012a). They find the cross-correlation of HCD absorbers and the Lyman-alpha forest to indeed be the dominant systematic error on the Lyman-alpha forest auto-correlation. They additionally identify as significant terms the HCD absorber auto-correlation and a three-point correlation between two Lyman-alpha forest modes and an HCD absorber mode.

An approximate model for HCD absorbers is used in the measurement of the 3D Lyman-alpha forest correlation function with BOSS DR12 spectra (Bautista et al., 2017) and the cross-correlation with the quasar distribution (du Mas des Bourboux et al., 2017). It is (the Fourier transform of) a biased linear power spectrum, with separate bias and redshift space distortion parameters for HCD absorbers, convolved with a top-hat filter in real space (*i. e.*, a sinc function in Fourier space) to approximate the profiles of HCD absorbers (hereafter, the "BOSS model"). The large-scale bias of dark matter halos hosting DLAs can be constrained through the cross-correlation of DLAs in spectra with the Lyman-alpha forest using BOSS spectra (Font-Ribera et al., 2012b; Pérez-Ràfols et al., 2017). This halo bias can then be related to the absorber flux transmission bias.

In this study, we measure the effect of HCD absorbers on correlations in the 3D Lyman-alpha forest using the 3D flux power spectrum in cosmological hydrodynamical simulations for the first time. We use simulation boxes from the Illustris project (Nelson et al., 2015; Vogelsberger et al., 2014a), which have been shown to reproduce the observed column density distribution function and spatial clustering of HCD absorbers at the 95% confidence level (Bird et al., 2014; Vogelsberger et al., 2014a). We measure the full anisotropic effect as a function of column density and redshift. We then consider how well the linear Voigt model characterises our results and identify the regimes

where this simple model breaks down. We also compare this model to the approximate BOSS model discussed above. Our results will improve the robustness of modelling HCD absorbers and hence improve cosmological inference for future Lyman-alpha forest surveys (*e. g.*, eBOSS/DESI).

We briefly explain the theory of the models that we consider for the Lyman-alpha forest and HCD absorbers in § 5.3. In § 5.4, our methodology in measuring the 3D flux power spectrum from hydrodynamical simulations and our modelling procedure are explained. We present our main results in § 5.5. These results are discussed and compared to previous work in § 5.6 and in § 5.7, we draw our conclusions.

5.3 Theory

5.3.1 Lyman-alpha forest

Fluctuations in the transmitted flux of the Lyman-alpha forest are given as $\delta_{\rm Forest}(\mathbf{x}) = \frac{\mathcal{F}_{\rm Forest}(\mathbf{x})}{\langle \mathcal{F}_{\rm Forest} \rangle} - 1$[1], where the transmitted flux $\mathcal{F}_{\rm Forest} = e^{-\tau_{\rm Forest}}$, $\tau_{\rm Forest}$ is the optical depth and $\langle \mathcal{F}_{\rm Forest} \rangle$ is the average flux over all spectral pixels. We then follow the standard treatment of the Lyman-alpha forest on large scales and model these fluctuations as a biased tracer of the underlying matter density fluctuation field with redshift-space distortions (by analogy with other tracers of the matter distribution like galaxies or galaxy clusters; Kaiser, 1984, 1987). It therefore follows that the 3D Lyman-alpha forest flux power spectrum can be modelled as

$$P^{\rm 3D}_{\rm Forest}(|\mathbf{k}|, \mu, z) = b^2_{\rm Forest}(1 + \beta_{\rm Forest}\mu^2)^2 P^{\rm 3D}_{\rm Linear}(|\mathbf{k}|, z) D_{\rm NL}(|\mathbf{k}|, \mu), \tag{5.1}$$

where (following Arinyo-i-Prats et al., 2015; McDonald, 2003) we introduce a parametric function $D_{\rm NL}(|\mathbf{k}|, \mu)$ to characterise deviations from linear theory due to non-linear effects; $P^{\rm 3D}_{\rm Linear}(|\mathbf{k}|, z)$ is the linear theory matter power spectrum; $b_{\rm Forest}$ is the (linear) bias parameter of the Lyman-alpha forest; and $\beta_{\rm Forest}$ is its redshift space distortion parameter. For the wavevector $\mathbf{k}$ (conjugate to $\mathbf{x}$), we use a spherical coordinate system with its zenith direction along the line-of-sight such that power spectra are functions of $|\mathbf{k}|$ and μ, which is the cosine of the angle between the wavevector and the line-of-sight. $P^{\rm 3D}_{\rm Forest}$ is a function of redshift z and in general, so are $b_{\rm Forest}$ and $\beta_{\rm Forest}$. Constraints on the redshift evolution of the Lyman-alpha forest power spectrum largely come at present from the 1D flux power spectrum (although there was some analysis of the redshift evolution of the

[1]For the 3D comoving spatial coordinate $\mathbf{x}$, the line-of-sight component $x_{||}$ is transformed from the line-of-sight velocity space of spectra by the Hubble law.

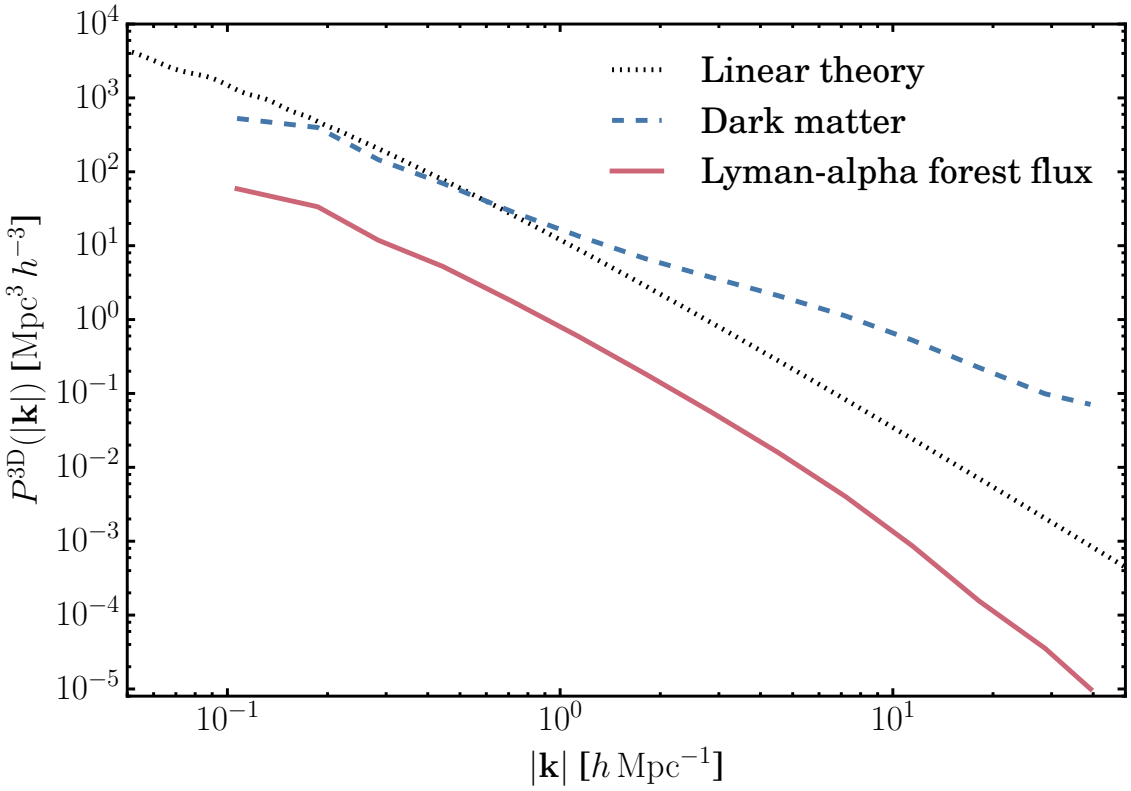

Figure 5.1: A comparison of three-dimensional power spectra (averaged over all angles) as predicted by linear theory, measured from dark matter particles in a hydrodynamical simulation and measured from the transmission flux of the Lyman-alpha forest in redshift space from mock spectra generated from the same simulation. Although the Lyman-alpha forest is a biased tracer of linear theory, it remains linear to much smaller scales than other probes, including the power spectrum of dark matter, which is affected by non-linear gravitational evolution from scales larger than $1\,h\,\mathrm{Mpc}^{-1}$. The simulation used is a $(75\,\mathrm{Mpc}\,h^{-1})^3$ box at redshift $z = 2.44$ from the Illustris project (Vogelsberger et al., 2014a).

3D power spectrum in Slosar et al. 2011). It is currently assumed (*e. g.*, Bautista et al., 2017) that $b_{\rm Forest} \propto (1+z)^{\gamma}$, where $\gamma = 2.9$, and that $\beta_{\rm Forest}$ does not depend on redshift such that $b^2_{\rm Forest} P^{\rm 3D}_{\rm Linear} \propto (1+z)^{3.8}$, roughly matching the evolution observed in the 1D flux power spectrum (*e. g.*, Palanque-Delabrouille et al., 2013). These assumptions are broadly supported by results from hydrodynamical simulations, although $\beta_{\rm Forest}$ is found to decrease with increasing redshift (*e. g.*, Arinyo-i-Prats et al., 2015).

The bias parameters of the Lyman-alpha forest differ from the biases of point objects like galaxies or halos in that denser regions of matter will have less transmitted flux (due to increased absorption by HI gas) and so $b_{\rm Forest}$ is negative. Figure 5.1 compares the 3D Lyman-alpha forest flux power spectrum (averaged over all angles) as measured in redshift space from our simulation (see § 5.4.1) to the linear theory matter power spectrum. The flux power spectrum appears to be a scaled version of the theory power spectrum, remaining so to much smaller scales than, *e. g.*, the dark matter power spectrum, which is strongly affected by non-linear gravitational evolution for $|\mathbf{k}| > 1\,h\,{\rm Mpc}^{-1}$. However, there is some small deviation from being a scaled version of the linear theory on small scales; in Eq. (5.1), this is characterised by the function $D_{\rm NL}$, which is calibrated from hydrodynamical simulations. This function allows for the isotropic growth in power due to non-linear growth, isotropic suppression by pressure on very small scales and suppression by non-linear peculiar velocities and temperature towards the line-of-sight. For consistency with Bautista et al. (2017), we use the fitting function of McDonald (2003) with the parameter values given in the first row of their Table 1. The parameters of this function have not been measured from data or previous simulations at the higher redshift that we consider ($z = 3.49$), and we are not able to do so with our simulations due to insufficient constraining power. Therefore for simplicity, we use the same parameter values in this function as at the lower redshift in the high redshift setting. Arinyo-i-Prats et al. (2015), in any case, found the shape of their simplest non-linear fitting function to evolve weakly in the redshifts they consider ($2.2 \leq z \leq 3$).

5.3.2 High column density absorbers

We follow Bautista et al. (2017); Font-Ribera and Miralda-Escudé (2012); McQuinn and White (2011) in modelling the 3D correlations of HCD absorbers on large scales as the convolution of a linear model in real space that does not include the damping wings of HCD absorbers [akin to Eq. (5.1)] with the profiles of the wings. Therefore, the 3D flux power spectrum of a set of HCD

absorbers with column densities in the interval $[N(\mathrm{HI})_{\mathrm{min}}, N(\mathrm{HI})_{\mathrm{max}}]$ is given as

$$P^{\mathrm{3D}}_{\mathrm{HCD}}(|\mathbf{k}|, \mu, z) = b^2_{\mathrm{HCD}}(1 + \beta_{\mathrm{HCD}}\mu^2)^2 P^{\mathrm{3D}}_{\mathrm{Linear}}(|\mathbf{k}|, z) F^2_{\mathrm{HCD}}(k_{||}, z), \tag{5.2}$$

where b_{HCD} and β_{HCD} are the bias and redshift space distortion parameters of the absorption caused by HCD absorbers (these will in general depend on redshift z). F_{HCD} is a function of the line-of-sight wavenumber $k_{||} = |\mathbf{k}|\mu$, since it is caused by the absorption profiles of HCD absorbers which only manifest along the line-of-sight:

$$F^{\mathrm{Voigt}}_{\mathrm{HCD}}(k_{||}, z) = \int_{N(\mathrm{HI})_{\mathrm{min}}}^{N(\mathrm{HI})_{\mathrm{max}}} \mathrm{d}N(\mathrm{HI}) f(N(\mathrm{HI}), z) V(k_{||}, N(\mathrm{HI})). \tag{5.3}$$

Here, $V(k_{||}, N(\mathrm{HI}))$ is the Fourier transform of the HCD absorbers' wing profiles as they manifest in the flux fluctuation field and $f(N(\mathrm{HI}), z)$ is the column density distribution function (CDDF). The model we consider in this study uses the profile of HCD absorbers, which is a Voigt function in optical depth (see *e. g.*, Appendix A of Rogers et al. 2017 for the full expression), the convolution of a Gaussian profile (caused by Doppler broadening) and a Lorentzian profile (caused by natural or collisional broadening). We will also consider the approximation made by the BOSS Collaboration (Bautista et al., 2017; du Mas des Bourboux et al., 2017) that the absorption profiles of HCD absorbers are top-hat filters ["BOSS model"]:

$$F^{\mathrm{BOSS}}_{\mathrm{HCD}}(k_{||}, z) = \frac{\sin(L_{\mathrm{HCD}} k_{||})}{L_{\mathrm{HCD}} k_{||}}, \tag{5.4}$$

where L_{HCD} is a free parameter setting the effective width of these filters.

By combining Eqs. (5.1) and (5.2) (and additionally remembering the cross-correlation between the Lyman-alpha forest and HCD absorber fields), the 3D flux power spectrum for the Lyman-alpha forest contaminated by a set of HCD absorbers is given as:

$$\begin{aligned} P^{\mathrm{3D}}_{\mathrm{Contaminated}}(|\mathbf{k}|, \mu, z) = P^{\mathrm{3D}}_{\mathrm{Linear}}(|\mathbf{k}|, z) \\ [\tilde{b}^2_{\mathrm{Forest}} D_{\mathrm{NL}}(|\mathbf{k}|, \mu) + 2\tilde{b}_{\mathrm{Forest}}\tilde{b}_{\mathrm{HCD}} + \tilde{b}^2_{\mathrm{HCD}}], \end{aligned} \tag{5.5}$$

where $\tilde{b}_{\mathrm{Forest}} = b_{\mathrm{Forest}}(1 + \beta_{\mathrm{Forest}}\mu^2)$ and $\tilde{b}_{\mathrm{HCD}} = b_{\mathrm{HCD}}(1 + \beta_{\mathrm{HCD}}\mu^2) F_{\mathrm{HCD}}(k_{||}, z)$. If there was uncertainty in the CDDF of a given sample of spectra, it will be preferable to sub-divide the column density integrals evaluated in the calculation of F_{HCD} in Eq. (5.3) and allow for extra terms in

Eq. (5.5), with bias parameters ($\tilde{b}_{\mathrm{HCD},i}$) for the N categories of HCD absorbers:

$$P^{\mathrm{3D}}_{\mathrm{Contaminated}}(|\mathbf{k}|,\mu,z) = P^{\mathrm{3D}}_{\mathrm{Linear}}(|\mathbf{k}|,z)\left[\tilde{b}^2_{\mathrm{Forest}} D_{\mathrm{NL}}(|\mathbf{k}|,\mu) + \sum_{i=1}^{N}\left(2\tilde{b}_{\mathrm{Forest}}\tilde{b}_{\mathrm{HCD},i} + \sum_{j=1}^{N}\tilde{b}_{\mathrm{HCD},i}\tilde{b}_{\mathrm{HCD},j}\right)\right]. \tag{5.6}$$

We mention also two possible additions that could be made to this model. First, the model in Eq. (5.5) does not consider any non-linear evolution in the clustering of HCD absorbers[2]. Second, as noted in Font-Ribera and Miralda-Escudé (2012), in the two-point function of the total contaminated flux, there will arise three- and four-point functions of the Lyman-alpha forest and HCD absorber fluctuations. This is because the forest and HCD absorption terms are multiplied: $\langle\mathcal{F}_{\mathrm{Total}}\rangle(1+\delta_{\mathrm{Total}}) = \langle\mathcal{F}_{\mathrm{Forest}}\rangle(1+\delta_{\mathrm{Forest}})\langle\mathcal{F}_{\mathrm{HCD}}\rangle(1+\delta_{\mathrm{HCD}})$. It follows that in the total flux power spectrum (Eq. (5.5)), there will be three- and four-point correlations involving δ_{Forest} and δ_{HCD}. The model presented in Eq. (5.5) only accounts for the leading two-point correlations; Font-Ribera and Miralda-Escudé (2012), however, found that the three-point term $\langle\delta_{\mathrm{Forest}}(\mathbf{x_1})\delta_{\mathrm{HCD}}(\mathbf{x_1})\delta_{\mathrm{Forest}}(\mathbf{x_2})\rangle$ is an important term on smaller scales (separations $r < 40\,\mathrm{Mpc}\,h^{-1}$). We discuss the possible impact of this additional term in § 5.6.

5.4 Method

We first outline the method we have used and then explain the steps in more detail in the following subsections (§ 5.4.1 to 5.4.4).

(1). We use a cosmological hydrodynamical simulation from the Illustris project (Nelson et al., 2015; Vogelsberger et al., 2014a) and generate mock spectra on a grid (562,500 in total, each at a velocity resolution of $10\,\mathrm{km\,s^{-1}}$ and with a typical length of $\sim 8,000\,\mathrm{km\,s^{-1}}$). We calculate these at two redshift slices ($z = [2.44, 3.49]$). (See § 5.4.1)

(2). The mock spectra we generate contain absorption from the Lyman-alpha forest and HCD absorbers. For our analysis, it is useful to have a set of spectra containing only the Lyman-alpha forest (still forming a regular grid to allow the use of fast Fourier transforms; FFTs). To achieve this, we replace spectra contaminated by HCD absorbers by a nearby spectrum

[2]Indeed, the full "BOSS model" as used by Bautista et al. (2017); du Mas des Bourboux et al. (2017) multiplies the last two terms in Eq. (5.5) by D_{NL}, the non-linear function calibrated by simulations of the Lyman-alpha forest only. We do not in the first instance include this correction to the linear Voigt model.

containing only the forest. Furthermore, we are able to construct boxes of spectra containing only the Lyman-alpha forest and a particular category of HCD absorber (*i. e.*, restricted to a particular column density interval) by replacing back the original spectra containing only that category of contamination. The details of this HCD "dodging" procedure are explained in § 5.4.2.

(3). For each box of spectra that we generate, we measure the three-dimensional (3D) flux power spectrum using an FFT. (See § 5.4.3.)

(4). Using these measurements of 3D flux power spectra, we fit the proposed model [Eq. (5.5)] using a Markov chain Monte Carlo (MCMC) method. (See § 5.4.4 and Appendix B.2.)

5.4.1 Hydrodynamical simulations and mock spectra

We use snapshots from the highest-resolution cosmological hydrodynamical simulation of the original Illustris project (Nelson et al., 2015; Vogelsberger et al., 2014a, Illustris-1[3]). The simulation adopts the following cosmological parameters: $\Omega_{\rm m} = 0.2726$, $\Omega_\Lambda = 0.7274$, $\Omega_{\rm b} = 0.0456$, $\sigma_8 = 0.809$, $n_{\rm s} = 0.963$ and $H_0 = 100\,h\,{\rm km\,s^{-1}\,Mpc^{-1}}$, where $h = 0.704$ (Vogelsberger et al., 2014b). The box has a comoving volume of $(106.5\,{\rm Mpc})^3$ and we consider snapshots at redshifts $z = 2.44$ and 3.49. The simulations are in broad agreement with observations of the HI CDDF (Vogelsberger et al., 2014a), the clustering of DLA halos (Bird et al., 2014) and the kinematics of HCD absorbers (Bird et al., 2015). For a summary of the relevant physics in the simulation and comparisons to observations, see Rogers et al. (2017).

For each snapshot, we generate mock spectra (using the `fake_spectra` code; Bird, 2017) containing only the Lyman-alpha absorption line, on a square grid of 562,500 spectra. Each spectrum extends the full length of the simulation box with periodic boundary conditions, giving a size in velocity space of $7,501$ and $8,420\,{\rm km\,s^{-1}}$ respectively at $z = 2.44$ and 3.49. We measure the optical depth τ in velocity bins of size $10\,{\rm km\,s^{-1}}$ along the spectrum. We then calculate the transmitted flux $\mathcal{F} = e^{-\tau}$. We convolve our spectra with a Gaussian kernel of ${\rm FWHM} = 8\,{\rm km\,s^{-1}}$, setting the simulated spectrographic resolution.

5.4.2 Dodging high column density absorbers

We associate with each mock spectrum that is contaminated by HCD absorbers a nearby spectrum containing only Lyman-alpha forest absorption. Indeed, if the transverse distance necessary to

[3]The simulation we use is publically available at `http://www.illustris-project.org/data`.

"dodge" the contaminating HCD absorber is small (as is expected considering the physical sizes of HCD absorbers, Krogager et al., 2012), the large-scale cosmological modes (*i. e.*, the Lyman-alpha forest modes) in the replacement spectrum should be identical to the original spectrum and the difference will be the HCD absorber modes only. The "dodging" procedure iteratively proposes a nearby replacement spectrum until one is found with no HCD absorber contamination (*i. e.*, there are no column densities, integrated over $100\,\mathrm{km\,s^{-1}}$ [4], exceeding $1.6 \times 10^{17}\mathrm{atoms\,cm^{-2}}$, the threshold for HCD absorbers). It searches for replacement spectra by successively generating spectra further away in a transverse direction from the original spectrum in steps of $10\,\mathrm{kpc}\,h^{-1}$ until a suitable spectrum is found. In this way, we are able to generate a box of spectra containing only the Lyman-alpha forest.

We are also able to generate boxes of spectra containing Lyman-alpha forest and HCD absorbers of a certain category (*i. e.*, column densities in a certain interval) by replacing back original spectra containing this particular category. We categorise spectra according to the maximum column density (again integrated over $100\,\mathrm{km\,s^{-1}}$) in each spectrum; there may be less dense HCD absorbers in each category but their effect will be sub-dominant since their damping wings are narrower.

Having generated these new boxes of spectra, we compute FFTs (§ 5.4.3), ignoring the transverse dodging distances and assuming that the dodged spectra lie on the original grid. Since the dodging distances are in general small (we find only $\sim 1\%$ to be $> 500\,\mathrm{kpc}\,h^{-1}$; see Appendix B.1), the error associated with this approximation is restricted to small scales, *i. e.*, large $|\mathbf{k}|$. We conduct an analysis of the error that arises from the irregular grid resulting from the dodging distances in Appendix B.1. Following these tests, we study only scales $|\mathbf{k}| < |\mathbf{k}|_{\mathrm{max}}$, where $|\mathbf{k}|_{\mathrm{max}} = 1\,h\,\mathrm{Mpc}^{-1}$; at these small values of $|\mathbf{k}|$, the dodging error is negligible compared to the effect of HCD absorbers that we wish to measure (see Fig. 5.3).

5.4.3 Three-dimensional flux power spectrum

We measure the 3D flux power spectrum at each redshift slice for our Lyman-alpha forest box of spectra ($P^{\mathrm{3D}}_{\mathrm{Forest}}$) and for our contaminated boxes of spectra for a number of HCD absorber categories ($P^{\mathrm{3D}}_{\mathrm{Contaminated}}$), the column density ranges of which we give in Table 5.1. We estimate the 3D flux power spectrum in bins of $|\mathbf{k}|$ and μ, $P^{\mathrm{3D}}_{\mathrm{Flux},i} = \frac{1}{N_i}\sum_n |\hat{\delta}_{\mathrm{Flux}}(\mathbf{k}_n)|^2$, where $\mathbf{k}_n$ lie within a given $(|\mathbf{k}|, \mu)$ bin and N_i are the number of modes in each bin i. $\hat{\delta}_{\mathrm{Flux}}(\mathbf{k})$ is the Fourier transform of the

[4] This is the same integration length as we used in our 1D flux power spectrum analysis in Rogers et al. (2017) and it amounts to ten neighbouring bins or a comoving length much larger than the most extensive HCD absorbers (Krogager et al., 2012). In Rogers et al. (2017), we also tested our sensitivity to the size of this integration length and found it made negligible difference to power spectrum estimates.

Table 5.1: The neutral hydrogen (HI) column density limits $[N(\mathrm{HI})_{\mathrm{min}}, N(\mathrm{HI})_{\mathrm{max}}]$ that define the categories of absorbing systems used in this work. The columns on the right show the percentage of spectra (at each redshift z that is considered) in our $(106.5\,\mathrm{Mpc})^3$ simulation box (Nelson et al., 2015; Vogelsberger et al., 2014a, Illustris-1) where the highest-density system belongs to a given category.

Absorber category	$N(\mathrm{HI})_{\mathrm{min}}$ [atoms cm^{-2}]	$N(\mathrm{HI})_{\mathrm{max}}$ [atoms cm^{-2}]	% of spectra in box at $z = 2.44$	% of spectra in box at $z = 3.49$
Lyman-α forest	0	1.6×10^{17}	69.6	45.7
LLS	1.6×10^{17}	1×10^{19}	14.9	27.0
Sub-DLA	1×10^{19}	2×10^{20}	8.1	14.3
Small DLA	2×10^{20}	1×10^{21}	4.1	7.8
Large DLA	1×10^{21}	∞	3.3	5.2

flux fluctuation field $\delta_{\mathrm{Flux}}(\mathbf{x}) = \frac{\mathcal{F}(\mathbf{x})}{\langle \mathcal{F} \rangle} - 1$. Here, for the mean flux $\langle \mathcal{F} \rangle$ we always use the mean flux of the original box of spectra (with no dodging) so that our modelling assumption that the flux fluctuations can be sub-divided into different absorber categories $\delta_{\mathrm{Total}} = \sum_i \delta_i$ holds true, where i indexes the different absorber categories. We use the convention of absorbing the $(2\pi)^3$ into the conjugate variable, *i. e.*, we define the Fourier transform as $\delta(k) = \int \delta(x) e^{-ikx} \mathrm{d}x$.

5.4.4 Modelling and Markov chain Monte Carlo sampling

We optimise the parameters of our model (the Lyman-alpha forest and HCD contamination biases and redshift-space distortion parameters) using MCMC sampling. The details of the construction of our likelihood function and prior probability distributions are given in Appendix B.2. We use MCMC sampling in order to estimate parameter uncertainties and to understand parameter degeneracies. The results are shown in § 5.5.2.

5.5 Results

5.5.1 Measuring the effect of HCD absorbers

Figure 5.2 shows the measured 3D flux power spectra as a function of scale $|\mathbf{k}|$ and the cosine of the angle away from the line-of-sight μ [5]. The anisotropic behaviour arises due to linear redshift-space distortions on larger scales (enhancing power towards the line-of-sight) and non-linear effects on smaller scales (suppressing power along the line-of-sight due to non-linear peculiar velocities and thermal broadening of absorption lines). The non-linear effects are more manifest in the bottom

[5] *I. e.*, $\mu = 1$ is along the line-of-sight and $\mu = 0$ is transverse to the line-of-sight.

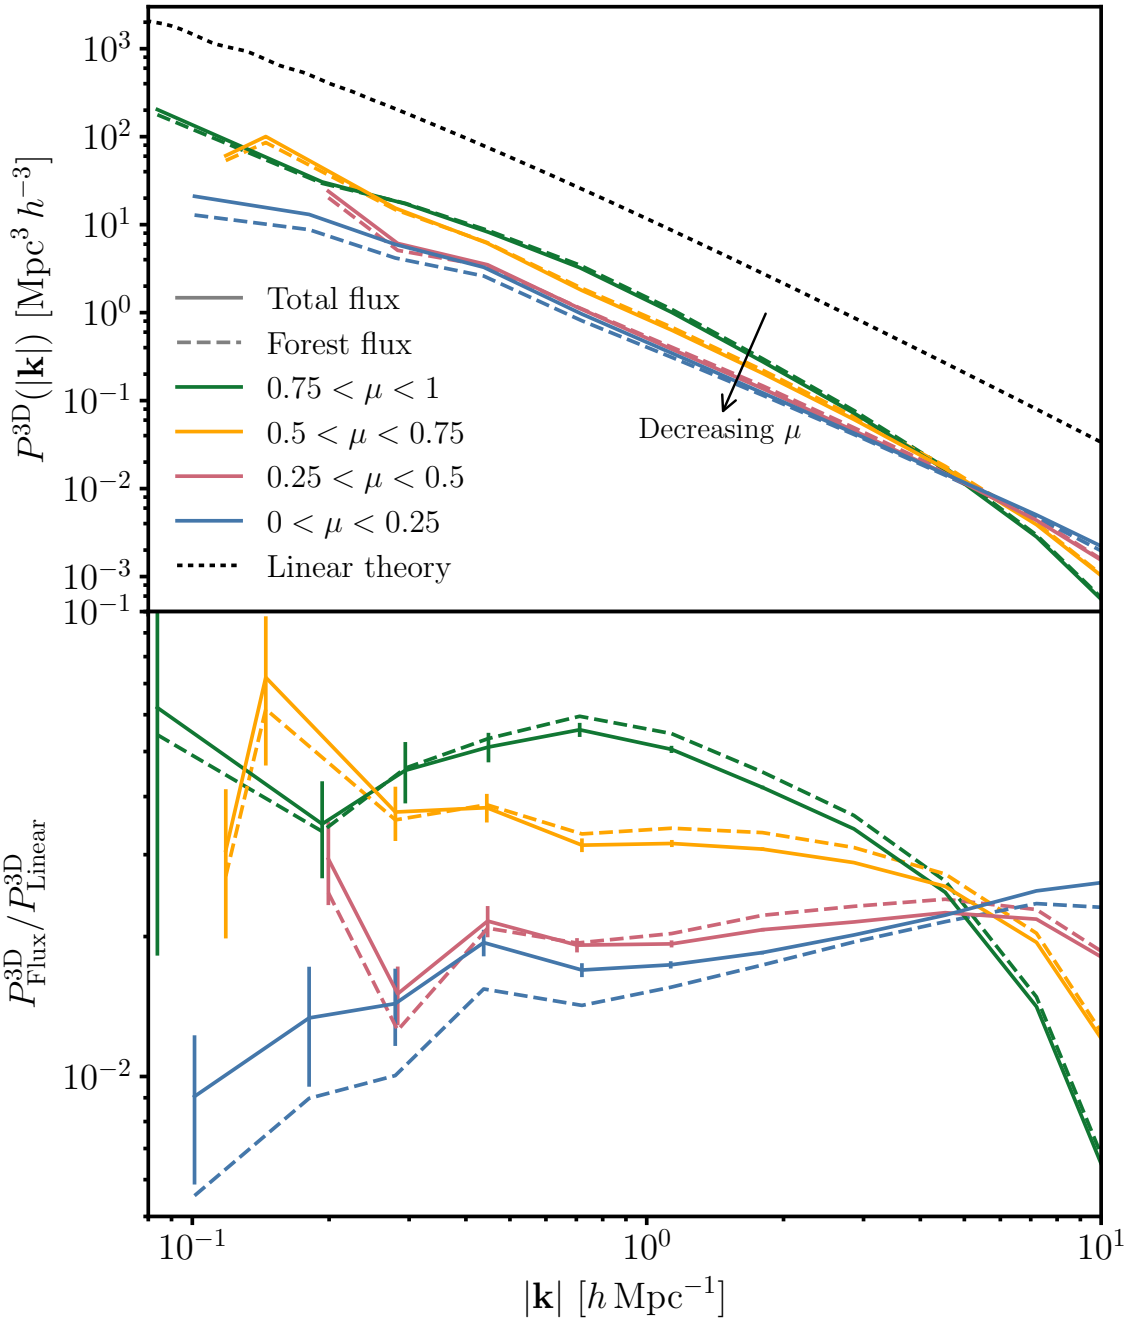

Figure 5.2: *Above*: the three-dimensional power spectra of the total flux from the Lyman-alpha forest and HCD absorbers (solid lines); and of the flux from the Lyman-alpha forest only (dashed lines); and the total linear theory matter power spectrum. For the flux power spectra, we show the anisotropic behaviour as a function of μ. *Below*: the flux power spectra in ratio to the linear power spectrum. Flux measurements are made from a simulation box at redshift $z = 2.44$.

panel, where the flux power spectra are shown in ratio to the linear theory matter power spectrum. On the largest scales, these ratios should tend towards constant values (*i. e.*, $b^2(1+\beta\mu^2)^2$ for the linear models presented in § 5.3). However, we are only able to probe a small number of these large scale modes in our $75\,\mathrm{Mpc}\,h^{-1}$ simulation box and so our measurement of large-scale bias has a large variance. The isotropic enhancement of power due to non-linear collapse of structure is broadly observable on larger scales; the anisotropic suppression of power towards the line-of-sight mentioned above is clearly observable for scales $|\mathbf{k}| > 1\,h\,\mathrm{Mpc}^{-1}$, leading to the characteristic cross-over in the curves on small scales (Arinyo-i-Prats et al., 2015; McDonald, 2003).

Figure 5.2 also compares the 3D flux power spectra of contaminated and uncontaminated Lyman-alpha forest absorption (solid and dashed lines respectively; see § 5.4.2 for more details about how a box of spectra without HCD absorber contamination is constructed). The effect of HCD absorber contamination adds power in some regimes (especially in the transverse direction) and suppresses power in others (especially on smaller scales towards the line-of-sight). Following the tests of the error caused by the dodging procedure in forming the Lyman-alpha forest box of spectra (see Appendix B.1 and § 5.4.2), we cut our data-vectors at $|\mathbf{k}|_{\mathrm{max}} = 1$ and we will throw away smaller scales in our following analysis. The forest flux power spectra shown in Fig. 5.2 (dashed lines) do not quantify the additional systematic and statistical error arising from the dodging (which is only significant for $|\mathbf{k}| > 1\,h\,\mathrm{Mpc}^{-1}$).

Figure 5.3 shows the fractional effect of HCD absorber contamination on the 3D Lyman-alpha forest flux power spectrum. The fractional effect of the full ensemble of HCD absorbers [panel *(a)*] can be as large as a 60% correction to $P^{\mathrm{3D}}(|\mathbf{k}|)$ at $|\mathbf{k}| = 0.1\,h\,\mathrm{Mpc}^{-1}$ in the transverse direction. The fractional effect is smaller at higher redshift because the Lyman-alpha forest power spectrum (in the denominator) has a larger amplitude (since neutral hydrogen is more abundant and so there is stronger Lyman-alpha absorption). There is a larger fractional effect in the transverse direction driven also by the Lyman-alpha forest power spectrum, which has less power in this direction due to redshift-space distortions. The scale-dependence in Fig. 5.3 is partly driven by the non-linear effects in the Lyman-alpha forest power spectrum discussed above; in particular, the Lyman-alpha forest power spectrum is boosted on small scales due to non-linear growth and so the fractional effect decreases.

The bottom panel Fig. 5.3 *(b)* shows the equivalent effects but for a mock residual contamination of HCD absorbers after the largest HCD absorbers have been "clipped" out (*i. e.*, only LLS and sub-DLAs remaining). The same trends are observed as above, but the overall amplitude is smaller since the largest damping wings have been removed. Nonetheless, the effect at $|\mathbf{k}| = 0.1\,h\,\mathrm{Mpc}^{-1}$

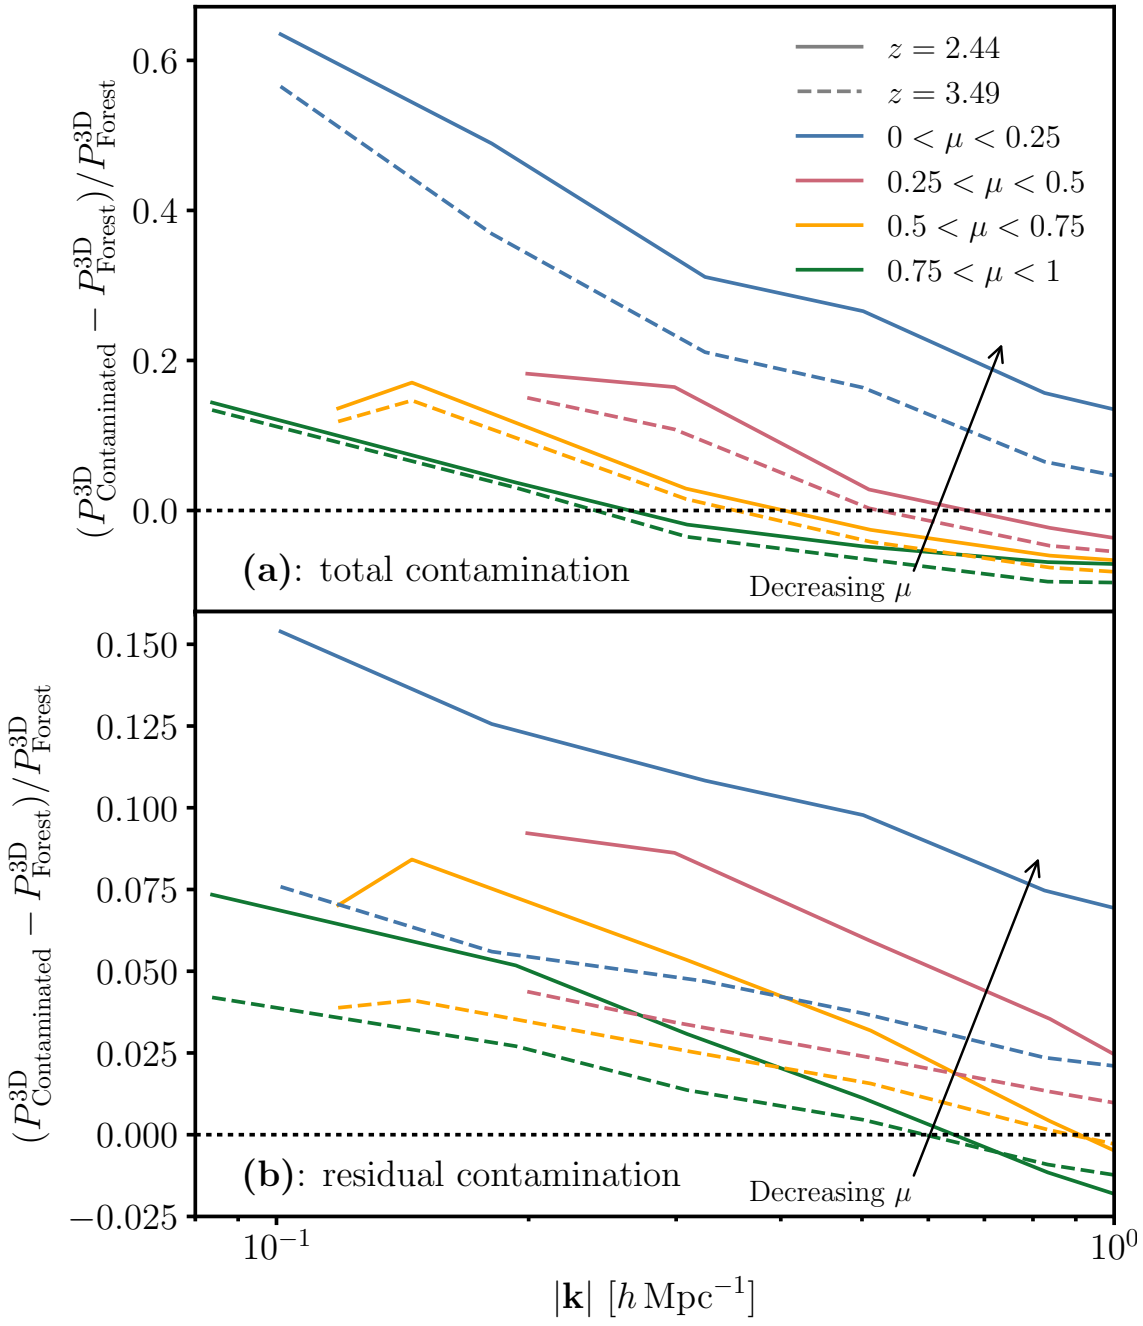

Figure 5.3: The change in the flux power spectrum from contamination of quasar spectra by HCD absorbers, as a fraction of the Lyman-alpha forest power spectrum. **(a)** *above*: the effect of the total contamination from all HCD absorbers in our simulation box; **(b)** *below*: the effect of a mock residual contamination after the largest HCD absorbers have been "clipped" from quasar spectra (*i. e.*, only LLS and sub-DLAs remaining).

in the transverse direction still constitutes a 15% correction; it is therefore necessary to model this effect for robust cosmological inference from the Lyman-alpha forest (see § 5.5.2).

5.5.2 Modelling the effect of HCD absorbers

We carry out the modelling as explained in § 5.4.4, using the simulated measurements of 3D flux power spectra we have presented in § 5.5.1. In particular, we investigate the linear Voigt model presented in Eq. (5.5) for the 3D flux power spectra contaminated by different categories of HCD absorber (LLS, sub-DLAs, small and large DLAs, as defined in Table 5.1); the Lyman-alpha forest bias parameters are additionally constrained by the forest-only 3D flux power spectrum (see § 5.4.4 for more details).

We show the maximum posterior values of the linear Voigt model and the measurements made in our simulations in Fig. 5.4. To emphasise the effect of HCD absorbers, we compare the part of the model for the auto-correlations of HCD absorbers and their cross-correlation to the Lyman-alpha forest [*i. e.*, the last two terms in Eq. (5.5)] to the difference between the 3D flux power spectra of the contaminated and uncontaminated boxes of spectra. We plot the results as a function of column density (by showing the effect for different HCD absorber categories from top to bottom) and as a function of redshift (from left to right). The widths of the error bars are the widths of the 1σ credible intervals of the model's posterior distribution. (This can be estimated from the distribution of the samples formed by evaluating the model at the parameter values of the samples from the parameter MCMC chains.) In this way, the uncertainty shown at each point in the data-space quantifies the full likelihood function that we have constructed and the prior information that we have assumed; thereby coherently propagating errors and correlations.

A measure of the goodness-of-fit is the values of the reduced chi-squared statistic: from top to bottom, left to right, $\chi^2_{\rm red} = (\rm a)\,1.55$; $(\rm b)\,1.56$; $(\rm c)\,1.57$; $(\rm d)\,1.73$; $(\rm e)\,1.04$; $(\rm f)\,1.03$; $(\rm g)\,1.21$; $(\rm h)\,1.24$, all indicating a good fit for the linear Voigt model for all the column densities and redshifts we have considered, excluding certain regimes as explained below. The number of degrees of freedom is (a), (b), (e), (f): 34; (c), (g): 26; (d), (h): 22. We wish to emphasise that this is an estimate of the goodness-of-fit of the full model (as given by Eq. (5.5)). This should not be compared to a "chi-by-eye" on Fig. 5.4, which is deliberately constructed to show an interesting sub-set of the model. The linear Voigt model is discrepant with the simulation results for a small part of the data space ($|\mathbf{k}| \gtrsim 0.4\,h\,\mathrm{Mpc}^{-1}$ for small DLAs and $|\mathbf{k}| \gtrsim 0.25\,h\,\mathrm{Mpc}^{-1}$ for large DLAs); these exceptions are indicated by the dotted lines in Fig. 5.4 and are restricted to small scales (particularly towards

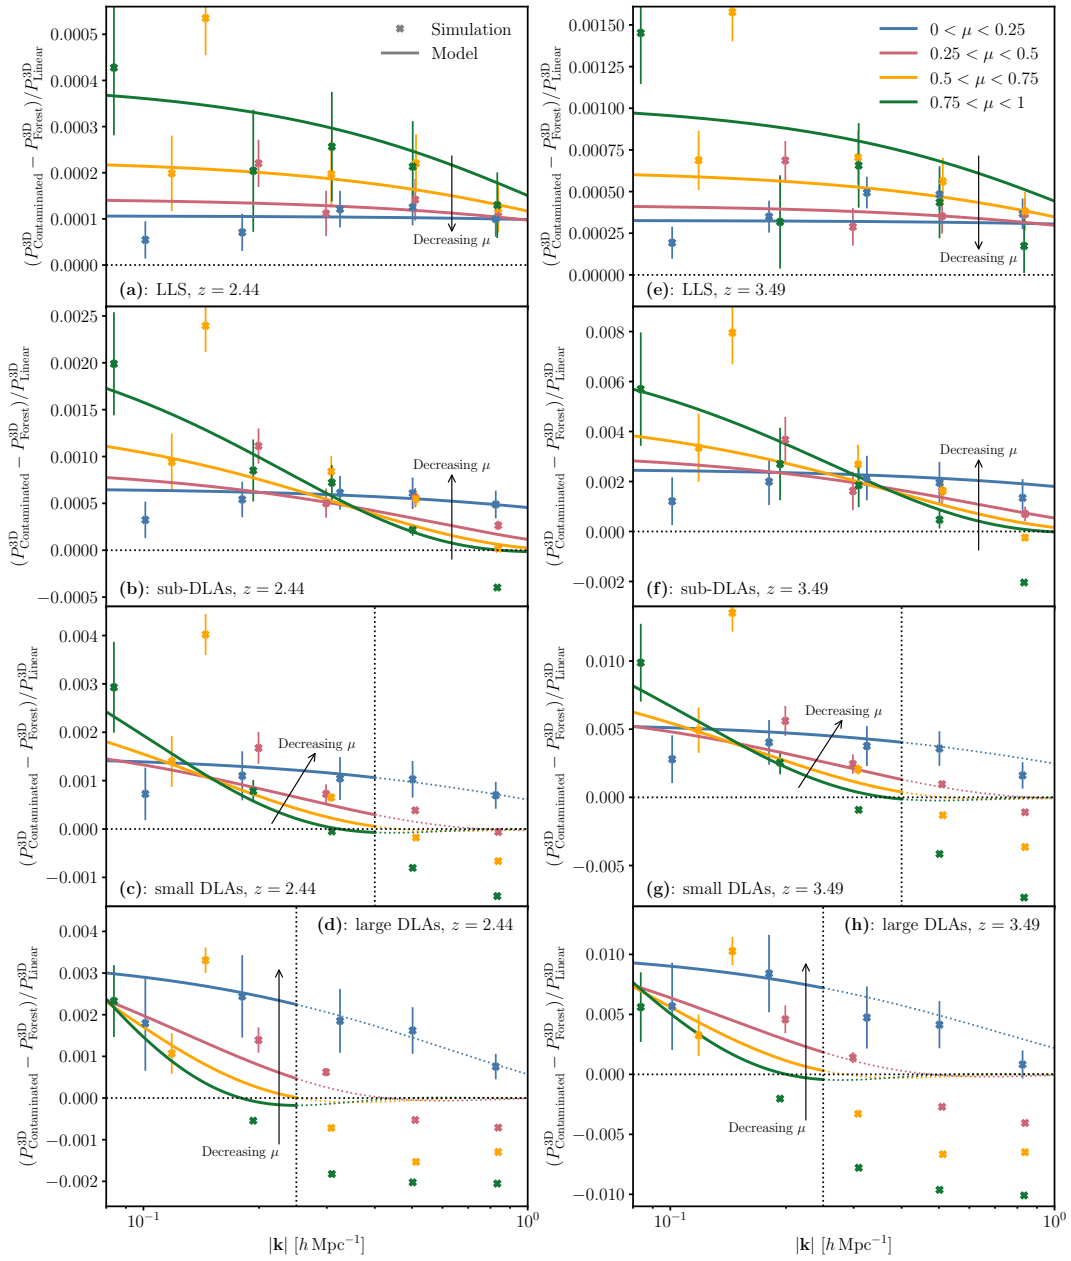

Figure 5.4: The change in the flux power spectrum from contamination of quasar spectra by different categories of HCD absorbers, in ratio to the linear power spectrum. The points are measurements from our simulation boxes. The lines are maximum posterior values of our preferred model. *From top to bottom*, we show the effect of different categories of HCD absorbers; *from left to right*, we show the effect at different redshifts z. The vertical dotted lines for the two largest HCD absorber categories indicate the smallest scale which we include in our data-vector from our HCD-contaminated simulation boxes for those categories. Our preferred model does not correctly characterise the simulation results for these categories on smaller scales towards the line-of-sight. The dotted lines show an extrapolation of this model, highlighting the discrepancy.

the line-of-sight) for the largest HCD absorbers (small and large DLAs). Indeed, we exclude these parts of the data-vector for the contaminated 3D flux power spectra in our parameter inference; we discuss the implications of this small-scale discrepancy for the highest column density absorbers in § 5.6.

There is further discrepancy between the best-fit model and (simulated) data points in regions of data-space which are not excluded. In particular, it is noteworthy that the yellow data-point second from the left in each panel is discrepant by several sigmas. This is possibly a reflection of an over-strong prior on the HCD absorber bias parameters in this part of the data-space. Indeed, without the prior information, the uncertainty on this data point is very large because of the very small number of large-scale modes and the model and data would (naively) be more consistent. Our prior distributions are constructed from the latest inferred values from observations, in the absence of previous such analyses with simulations. Considering that there are observed differences between the values of the Lyman-alpha forest bias parameters inferred from observations and simulations, it would not be surprising if this extended to the high-density biases. A future solution to this problem would be to measure the large-scale absorber bias of HCD systems on a larger simulation box (than is currently available). We however reiterate that although the amplitude of our best-fit model is determined by our prior information, the scale-dependence is physically determined. Furthermore, the effective absorber bias of HCD systems will ultimately be survey-specific after the attempted removal of the largest damping wings. There is also a discrepancy on small scales towards the line-of-sight in some of the un-excluded data-space; this is part of the same problem that persists in the excluded data-space. It reflects a fundamental limitation of the simple linear model we are testing. We discuss in some detail in § 5.6 the limitations of this simple model and possible sophistications that could be made. We highlight here however that the fundamental limitation of this analysis is an insufficient physical understanding of the small-scale (non-linear) effect of HCD absorbers (rather than the statistics of our analysis). We finally note that part of the discrepancy between simulation and model in Fig. 5.4 is also driven by the range of μ values within each μ bin.

For completeness, we quote the maximum (marginalised) posterior values and 1σ credible intervals of the (linear) bias parameters of the Lyman-alpha forest[6] at $z = 2.44$:

$$b_{\rm Forest}(1+\beta_{\rm Forest}) = -0.270 \pm 0.004; \quad \beta_{\rm Forest} = 1.722 \pm 0.072$$

[6]Rather than $b_{\rm Forest}$, following *e. g.*, Bautista et al. (2017); Slosar et al. (2011), we sample the combination $b_{\rm Forest}(1+\beta_{\rm Forest})$, which is less correlated with $\beta_{\rm Forest}$.

and at $z = 3.49$:

$$b_{\rm Forest}(1 + \beta_{\rm Forest}) = -0.511 \pm 0.006; \quad \beta_{\rm Forest} = 1.249 \pm 0.043.$$

The posterior on $\beta_{\rm Forest}$ at $z = 2.44$ is in 1σ agreement with the best-fit value from BOSS DR12 spectra (Bautista et al., 2017), $\beta_{\rm Forest}^{\rm BOSS} = 1.663 \pm 0.085$ at a central redshift of $z = 2.3$. However, the posterior on $b_{\rm Forest}(1 + \beta_{\rm Forest})$ is lower than the value measured from data $b_{\rm Forest}^{\rm BOSS}(1 + \beta_{\rm Forest}^{\rm BOSS}) = -0.325 \pm 0.004$ at $z = 2.3$; this difference has been observed in other studies with hydrodynamical simulations (*e. g.*, Arinyo-i-Prats et al., 2015). The redshift evolution in $b_{\rm Forest}$ observed in our simulations (modelled as $b_{\rm Forest} \propto (1 + z)^{\gamma}$) implies $\gamma = 3.1$, roughly matching the value currently assumed in data analyses $\gamma^{\rm BOSS} = 2.9$ (see § 5.3.1). We find that $\beta_{\rm Forest}$ decreases at higher redshift, also as observed in previous studies with simulations (Arinyo-i-Prats et al., 2015). In Appendix B.3, we test the sensitivity of our inference of the bias parameters of the Lyman-alpha forest to the smallest scale included in our analysis $|\mathbf{k}|_{\rm max}$; we find that our inferences are overall insensitive to this, suggesting that our results are robust to our modelling of non-linear effects. We also recover the same posterior distributions on the Lyman-alpha forest bias parameters for each of the HCD absorber categories of contaminated flux power spectra that we consider. These parameters also match those inferred from the 3D Lyman-alpha forest flux power spectrum only.

As discussed in § 5.4.4, we place Gaussian priors on the bias parameters of the different categories of HCD absorber (but not the forest biases), which are otherwise poorly constrained, since the amplitude of their effect is sub-dominant to the Lyman-alpha forest flux power spectrum. These prior distributions are returned almost exactly in the marginalised posteriors. The scale-dependence of the effect of HCD absorbers, meanwhile, is fully determined by the physics of their absorption profiles and the appropriate CDDF.

Figure 5.5 compares the maximum posterior values of the linear Voigt model to the effect of a mock residual contamination of HCD absorbers on the 3D flux power spectrum, at the two redshifts we consider. This mock residual contamination (of LLS and sub-DLAs) approximately matches the column densities assumed remaining in BOSS spectra (Bautista et al., 2017) after the largest damping wings have been removed. We find that the linear Voigt model is in statistical agreement with our simulation measurements ($\chi^2_{\rm red} =$ (a) 1.56; (b) 1.03). The number of degrees of freedom is (for both panels) 34. We highlight this configuration because, although Fig. 5.4 shows that there are some small scales towards the line-of-sight for the largest HCD absorbers where the simple linear Voigt model is not appropriate, these are the HCD absorber categories most efficiently removed in

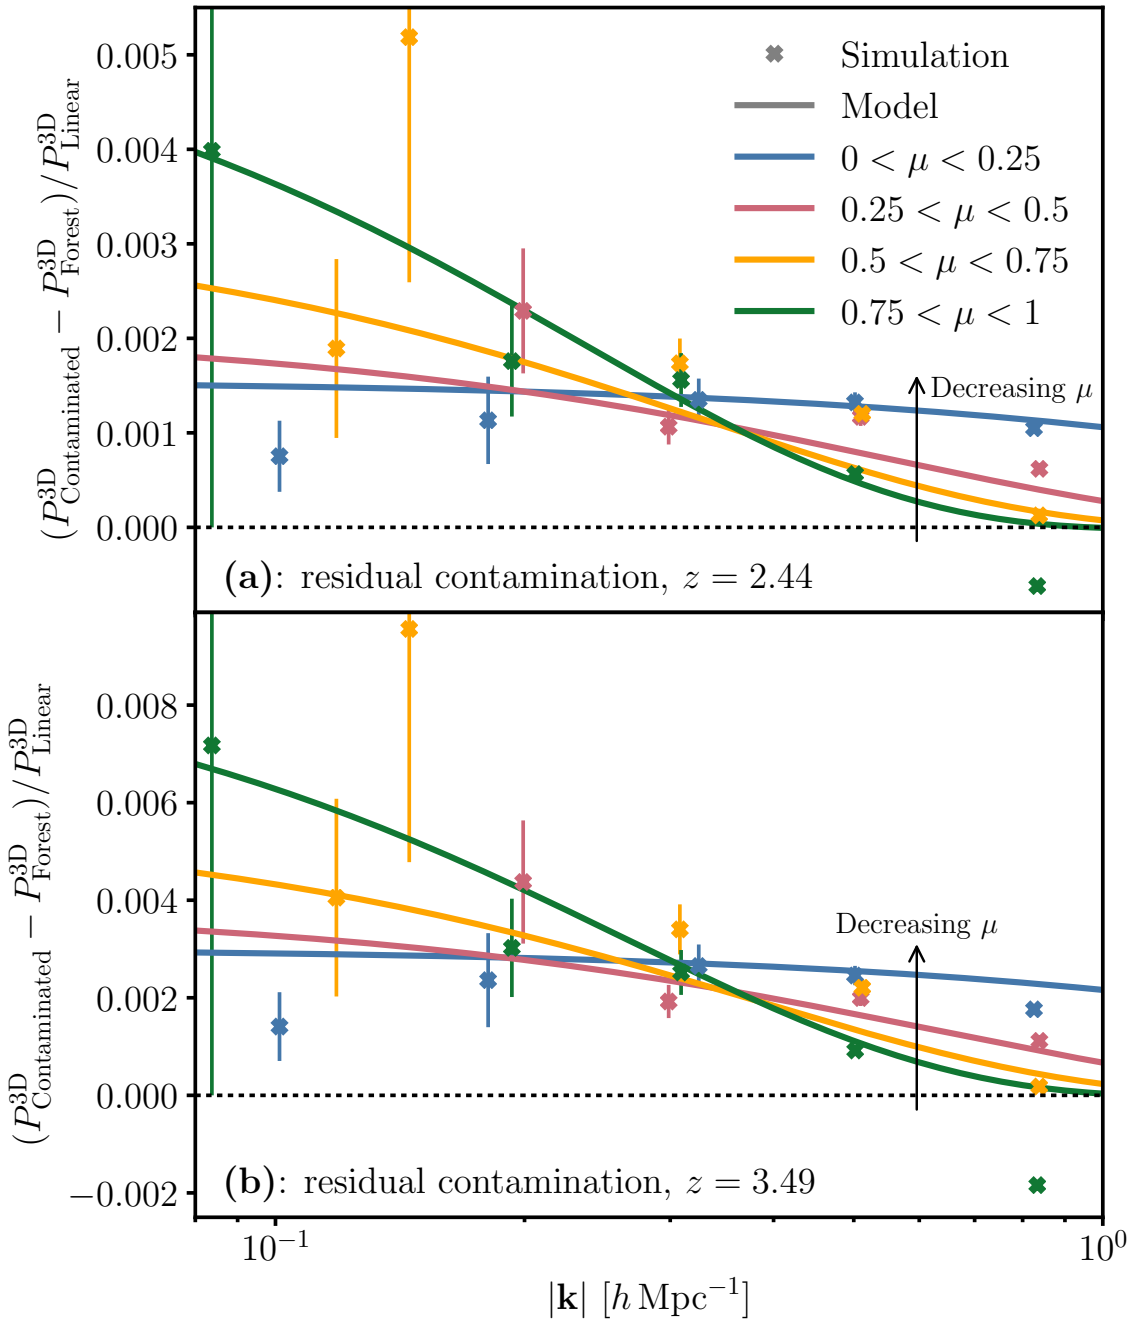

Figure 5.5: As Fig. 5.4, but for a mock residual contamination after the largest HCD absorbers have been "clipped" from quasar spectra (*i. e.*, only LLS and sub-DLAs remaining). *From top to bottom*, we show the effect at different redshifts z.

the “clipping” process in data analysis.

5.6 Discussion

In § 5.5.2, we showed the regimes in scale and column density where the simple linear Voigt model (see § 5.3.2) can characterise the effect of HCD absorber contamination on correlations in the 3D Lyman-alpha forest. The linear Voigt model is probably the simplest model that can be constructed to take account of the true absorption line profiles of HCD absorbers, convolving a linear model of HCD absorbers (as a biased tracer of the matter density distribution with redshift-space distortions) with the Voigt profiles of HCD absorbers' damping wings, integrated over the CDDF of the HCD absorbers. The linear Voigt model manifests as a suppression in power due to damping wings which remove structure in the spectra. This effect is stronger towards the line-of-sight since this is the direction in which the wings appear. The scales at which the suppression starts (and the overall amplitude) are larger for more dense absorber categories since their wings are wider. The suppression is such that it overcomes the boost in power towards the line-of-sight on large scales due to redshift-space distortions; consequently, there is a characteristic cross-over in the different curves for each absorber category. The effect of HCD absorbers transverse to the line-of-sight is scale-independent, since there is no projection of the damping wings in this direction (our lowest μ bin does include some modes slightly away from the transverse direction). The amplitude increases with redshift, mainly because the cross-correlation with the Lyman-alpha forest is stronger (there being overall more absorption at higher redshift).

However, on small scales towards the line-of-sight for the largest HCD absorber categories ($|\mathbf{k}| \gtrsim 0.4\,h\,\mathrm{Mpc}^{-1}$ for small DLAs and $|\mathbf{k}| \gtrsim 0.25\,h\,\mathrm{Mpc}^{-1}$ for large DLAs), the linear Voigt model cannot characterise our simulation results. Power is suppressed towards the line-of-sight more strongly than our model allows such that there is less power than without the HCD absorbers. We consider two possible causes of this discrepancy with the linear Voigt model (as mentioned in § 5.3.2). First, our model does not consider any non-linear clustering of the gas or halos associated with HCD absorbers. A comprehensive model for the clustering of HCD absorbers should certainly account for this effect, but we found no preference for a parametric form (akin to that used for the Lyman-alpha forest; Arinyo-i-Prats et al., 2015; McDonald, 2003) that would improve the fit to our simulation results, suggesting that such a closed form cannot alone account for the discrepancy. Second, we consider the non-linear effect of a three-point correlation between a Lyman-alpha forest fluctuation and an HCD absorber fluctuation at the same position and a Lyman-alpha forest

fluctuation at a second position (see § 5.3.2). This term was shown by Font-Ribera and Miralda-Escudé (2012) to be at least as significant as the HCD absorber auto-correlation on small scales (separations $r < 40\,\mathrm{Mpc}\,h^{-1}$) and to have the correct (negative) sign to account for the additional suppression of power observed on small scales towards the line-of-sight for small and large DLAs. (This is because it is the correlation between three negatively-biased tracers.) It is intuitively understood as the effect of the damping wings in masking regions of the Lyman-alpha forest and so suppressing auto-correlations in the Lyman-alpha forest that would otherwise occur on scales within the widths of individual wings. This effect will be stronger for more dense HCD absorbers since their damping wings are wider and so mask more of the Lyman-alpha forest; stronger on scales smaller than the widths of wings; and stronger towards the line-of-sight since this is the direction in which the masking occurs. This seems a qualitative match to the observed discrepancies with the linear Voigt model, but as yet there exists no simple model for this higher-order effect and we have not explicitly tested whether it can account for the observed discrepancies. As discussed above and in § 5.5.2, the effect is restricted to the highest column densities, which are in any case mostly removed in the clipping pre-processing of spectra.

Figure 5.6 compares the linear Voigt model as inferred from our simulations contaminated by a mock residual contamination of HCD absorbers (only LLS and sub-DLAs remaining, approximating the effect of clipping out the damping wings of more dense absorbers as is done with survey spectra), with the model used by the BOSS Collaboration (Bautista et al., 2017; du Mas des Bourboux et al., 2017) for the same effect. This "BOSS model" approximates the damping wings as top-hats and so the effect on the Fourier space correlations (*i. e.*, the flux power spectrum) is a sinc function (see § 5.3.2). We rescale the BOSS model to have the same bias and redshift-space distortions as inferred in our simulation box (for a fair comparison to the linear Voigt model), but use the shape parameter as found in BOSS mock spectra with a residual contamination and in data ($L_{\mathrm{HCD}} = 24.341\,\mathrm{Mpc}\,h^{-1}$)[7]. We extrapolate the BOSS model to smaller scales than considered in their analysis, where the minimum separations measured were $r = 10\,\mathrm{Mpc}\,h^{-1}$. Although our inference on the linear Voigt model is only constrained by the scales accessible in our simulation box, we extrapolate this model to larger scales of relevance to a BAO analysis. Although we are not able to explicitly test the model on these larger scales, it is expected to correctly characterise the effect as it constitutes the physical expectation on large scales. We conclude from Fig. 5.6 that

[7]We use this shape parameter value because the shape parameter of the BOSS model as inferred from our simulation box was considerably smaller than the BOSS best-fit value (in order to fit the small-scale correlations) and gave unphysical results on scales larger than the size of our box (*i. e.*, the curves were flat on large scales, indicating no effect of damping wings in contradiction to the physical linear Voigt model).

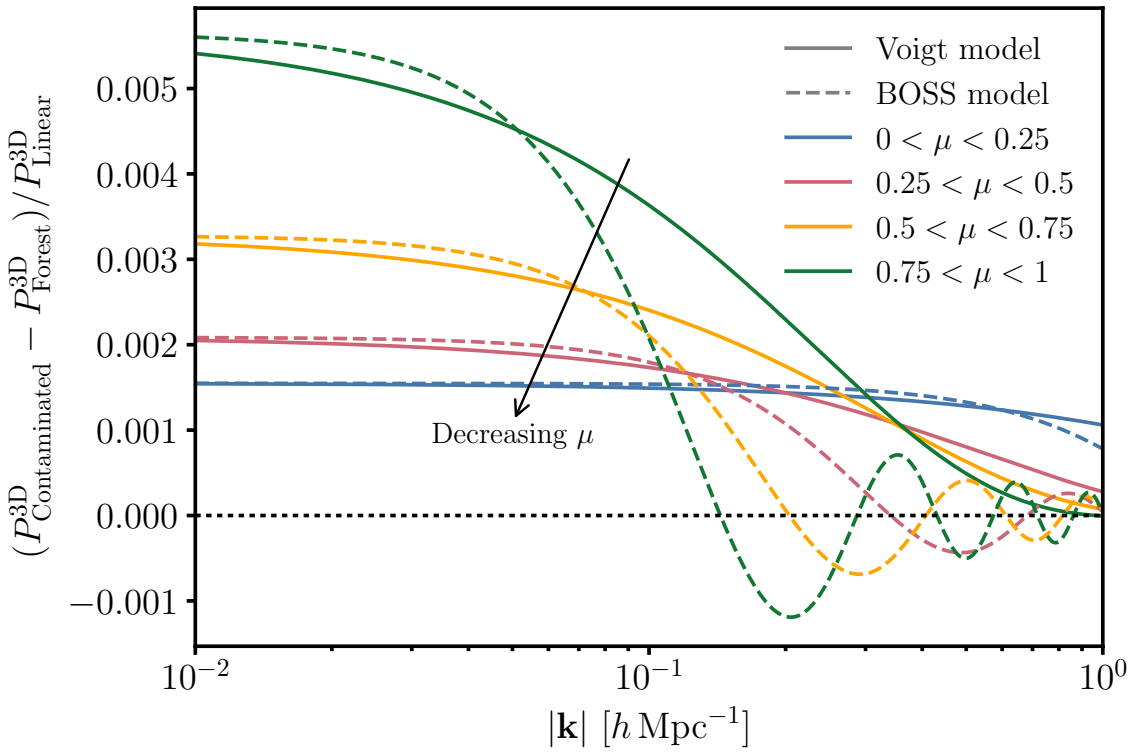

Figure 5.6: A comparison of the existing model as used by the BOSS Collaboration (Bautista et al., 2017; Pérez-Ràfols et al., 2017) and the linear Voigt model presented in this study for the additive effect of residual HCD absorbers (after the "clipping" of the largest absorbers from quasar spectra) on the three-dimensional flux power spectrum. For the linear Voigt model, we show the maximum posterior values as inferred from a mock residual contamination in our simulation box at $z = 2.44$. For the BOSS model, we rescale to match the bias and redshift-space distortions inferred in our box, but use the best-fit value of the shape parameter as found in BOSS mock spectra with a residual contamination and data. The maximum posterior value of the BOSS model as inferred from our simulation gives unphysical results on scales larger than the size of our box.

the BOSS model constitutes a good approximation for scales of relevance for a BAO analysis, but that on smaller scales, the linear Voigt model should be used in order to account for the effect of extended damping wings for a residual contamination of HCD absorbers.

5.7 Conclusions

We have measured the effect of contamination of quasar spectra by the damping wings of high column density (HCD) absorbing regions of neutral hydrogen on correlations in the 3D Lyman-alpha forest. We accomplished this by measuring 3D flux power spectra from a cosmological hydrodynamical simulation (Illustris; Nelson et al., 2015; Vogelsberger et al., 2014a) as a function of the column density of the HCD absorber contamination and redshift. We found that, even after the largest damping wings have been removed (as performed by survey pipelines), that the effect of the residual contamination can be as large as a 15% correction to the 3D Lyman-alpha forest flux power spectrum (at $|\mathbf{k}| = 0.1\,h\,\mathrm{Mpc}^{-1}$). We found that the effect of this residual contamination can be characterised by a simple linear model (with bias and redshift-space distortions) convolved with the Voigt profiles of the damping wings and integrated over the column density distribution function of the HCD absorbers. This model also successfully characterises the contamination effect on large scales for the highest column densities; however, on smaller scales (*e. g.*, $|\mathbf{k}| > 0.4\,h\,\mathrm{Mpc}^{-1}$ for small DLAs) towards the line-of-sight, the model fails possibly due to additional suppression in power by the most massive systems due to the effective masking of auto-correlations in the Lyman-alpha forest by their damping wings. Font-Ribera et al. (2017) found that there is much more constraining power in the 3D flux power spectrum than the 1D power spectrum for BOSS for $|\mathbf{k}| < 1\,h\,\mathrm{Mpc}^{-1}$, underlying the importance of accurately modelling systematics up to small scales. We therefore find that this linear Voigt model will help with precision measurements of BAO in future surveys (eBOSS/DESI) and will be essential for reconstructing the power spectrum shape beyond BAO.

Acknowledgements

KKR thanks the organisers of the COSMO21 symposium in 2016, where this project was devised.

A

Appendix to Chapter 4

A.1 One-dimensional flux power spectrum of a Voigt profile

As discussed in § 4.3, the broadened absorption lines of high column density absorbers are usually modelled by a Voigt profile. A Voigt profile is a convolution of a Lorentzian profile and a Gaussian profile. It therefore appropriately models the combination of the main physical processes that broaden atomic transition lines: the Lorentzian profile from *e. g.*, natural or collisional broadening and the Gaussian profile from *e. g.*, Doppler broadening. The optical depth as a function of wavelength $\tau(\lambda)$ is the product of the line-of-sight column density N and the atomic absorption coefficient $\alpha(\lambda)$ (*e. g.*, Humlicek, 1979)[1]:

$$\tau(\lambda) = N\alpha(\lambda) = N\frac{\sqrt{\pi}e^2}{4\pi\epsilon_0 m_\mathrm{e} c^2}\frac{f\lambda_\mathrm{t}^2}{\Delta\lambda_\mathrm{D}}u(x,y), \tag{A.1}$$

where the fundamental physical constants have their usual meaning, f is the oscillator strength of the atomic transition, λ_t is the transition wavelength and the Doppler wavelength "shift" associated with a gas of temperature T for an ion of mass m_ion,

$$\Delta\lambda_\mathrm{D} = \frac{\lambda_\mathrm{t}}{c}\left(\frac{2k_\mathrm{B}T}{m_\mathrm{ion}}\right)^{\frac{1}{2}}. \tag{A.2}$$

$u(x,y)$ is an unnormalised form of the Voigt function (the normalisation is already expressed in the pre-factors of Eq. (A.1)), specifically the real part of the Faddeeva function:

$$w(z) = e^{-z^2}\mathrm{erfc}(-iz) = u(x,y) + iv(x,y), \tag{A.3}$$

[1] Eq. (A.1) is valid in SI units.

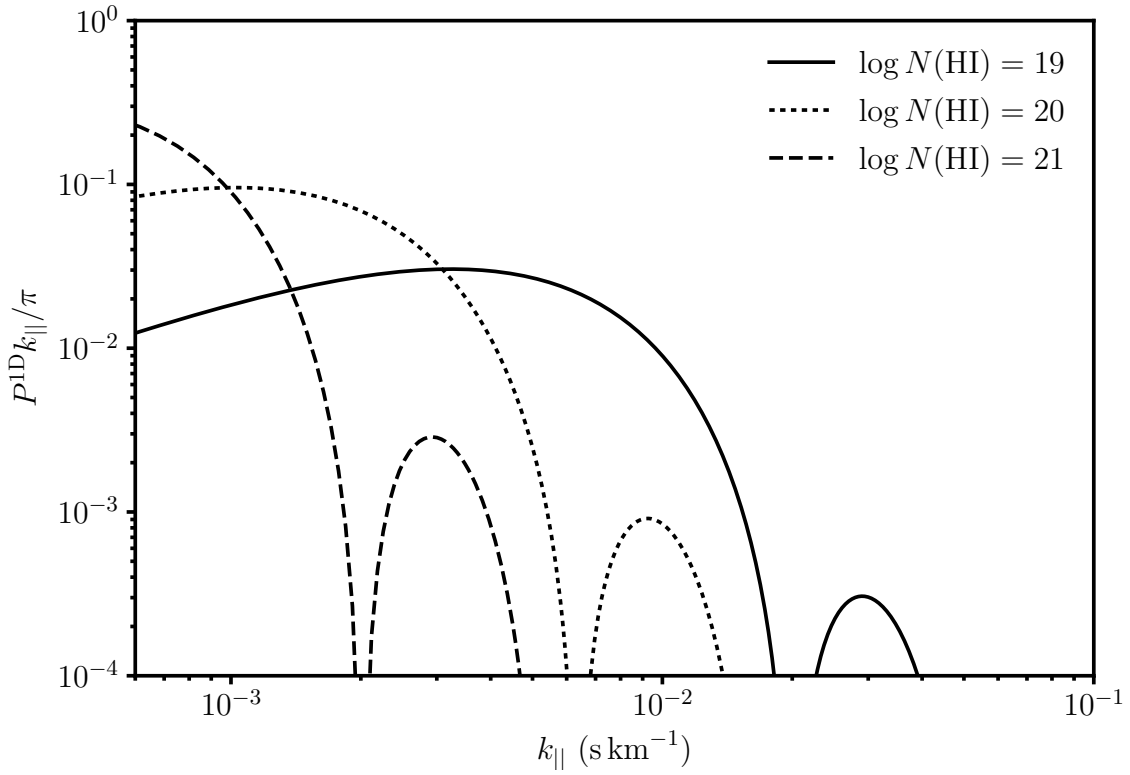

Figure A.1: The one-dimensional flux power spectra of Voigt profiles of broadened Lyman-alpha absorption lines as generated by different column densities of neutral hydrogen $N(\mathrm{HI})$, as a function of line-of-sight scale $k_{||}$ (the units of $N(\mathrm{HI})$ are $\mathrm{atoms\,cm^{-2}}$).

where $\mathrm{erfc}(x)$ is the complementary error function and $z = x + iy$. x and y are respectively the wavelength difference from the line centre λ_c and the natural width of the transition, in units of the Doppler shift:

$$x(\lambda) = \frac{\lambda - \lambda_\mathrm{c}}{\Delta\lambda_\mathrm{D}}; \quad y = \frac{\Gamma\lambda_\mathrm{t}^2}{4\pi c}\frac{1}{\Delta\lambda_\mathrm{D}}, \tag{A.4}$$

where Γ is the damping constant of the transition, *i. e.*, the inverse of the time scale for the electron to remain in the upper level of the transition in the vacuum. For the Lyman-alpha transition, $f = 0.4164$, $\lambda_\mathrm{t} = 1215.67$Å, $m_\mathrm{ion} = m_\mathrm{proton}$ and $\Gamma = 6.265 \times 10^8$ Hz (Morton, 2004). For the column densities that we consider, we assume a gas temperature $T \approx 10^4 K$. In order to calculate the 1D flux power spectrum arising from these Voigt profiles, the same procedure is followed as in § 4.4.2, *i. e.*, we form flux spectra and carry out a Fourier transform. We transform from wavelengths to velocities by $\Delta v/c = \Delta\lambda/\lambda$.

Figure A.1 shows the 1D flux power spectra of Voigt profiles as given by Eq. (A.1) for the Lyman-alpha absorption line for three different column densities of neutral hydrogen $N(\mathrm{HI}) = [10^{19}, 10^{20}, 10^{21}]\,\mathrm{atoms\,cm^{-2}}$, spanning the column densities for LLS and DLAs. This figure should be compared with Fig. 4.1 in § 4.5, which shows the 1D flux power spectra we have measured in the hydrodynamical simulations. The trends in Fig. A.1 broadly support the arguments made in § 4.6, relating the large-scale power spectrum of simulated spectra contaminated by high column density absorbers to the power spectrum of relevant Voigt profiles. The shape of the large-scale power spectrum of the Voigt profiles is similar in amplitude and scale-dependence as the excesses on large scales for the 1D flux power spectra of simulated spectra in high column density absorber categories. Moreover, these excesses get steeper, increase in amplitude and become prominent on larger scales for higher column densities, both in the simulated and analytic spectra. This reflects the fact that a higher column density means wider damping wings and so correlations on larger scales. In the analytic power spectra in Fig. A.1, we observe oscillations in the power spectrum on smaller scales that rapidly decrease in amplitude. These are not observed in the fully-simulated power spectra since the oscillations are orders of magnitude lower in amplitude than the flux power spectrum of residual Lyman-alpha forest (see Fig. 4.1). Furthermore, in our results, we are effectively averaging over a number of column densities in each column density bin (or absorber category) that we consider; this will have the additional effect of averaging out these smaller-scale oscillations in the power spectrum to form a smoother scaling.

B

Appendix to Chapter 5

B.1 Tests of HCD absorber dodging

In this Appendix, we test the effect of replacing simulated spectra contaminated by HCD absorption with nearby uncontaminated spectra on our measurements of 3D flux power spectra. The measurements of these power spectra are made computationally simple by the use of FFTs, which in turn require a regular grid of samples. However, the transverse HCD absorber "dodging" of some spectra makes this grid irregular. An error therefore arises from treating this irregular grid as the original regular grid (*i. e.*, to ignore the transverse dodging distances) in computing the necessary FFTs.

Figure B.1 shows the distribution of the transverse dodging distances required to find replacement mock spectra uncontaminated by HCD absorbers, for the simulation boxes at the two redshifts we consider. (See § 5.4.2 for more details about why and how we dodge HCD absorbers.) Replacement spectra are trialled increasingly further away from the original spectrum in steps of $10\,\mathrm{kpc}\,h^{-1}$ until an uncontaminated spectrum is found. Many of the final replacement spectra require many iterations to be found; this is exacerbated by dodging one HCD absorber but then finding another HCD absorber elsewhere along the spectrum which then requires further dodging. More dodging is required at higher redshift because neutral hydrogen is more abundant.

We then replicate the exact movement of lines-of-sight as we carry out in our simulations to dodge HCD absorbers on the estimation of a cosmological power spectrum (the same as input to the simulations) from a Gaussian random field (in order to approximate the effect on the flux power spectrum). An error is introduced in estimating the power spectrum and ignoring the transverse distances that some of the samples of the field have moved (this error could be avoided if we didn't rely in the calculation of power spectra on FFTs which require evenly sampled functions). Figure B.2 shows the fractional error from this effect for the dodges we carry out at the different redshifts

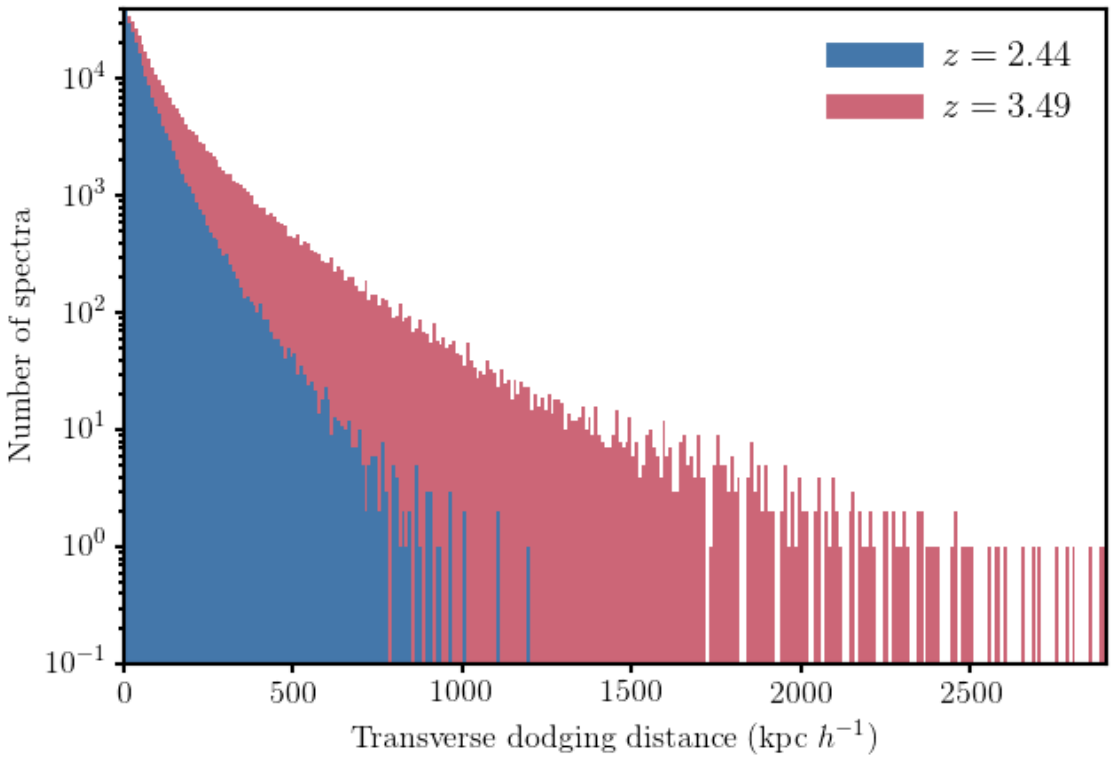

Figure B.1: Histogram of the transverse comoving distances "dodged" by each simulated spectrum in order to avoid HCD absorbers. The total number of spectra at each redshift z is 562,500. The number of spectra remaining un-dodged at $z = 2.44$ and 3.49 is respectively 391,500 (69.6%) and 257,063 (45.7%). There is a tail of large dodging distances, much larger than the physical size of the most massive HCD absorbers because sometimes, in dodging one absorber, the proposed replacement spectrum will coincide with another absorber, somewhere else along the line-of-sight, requiring further dodging.

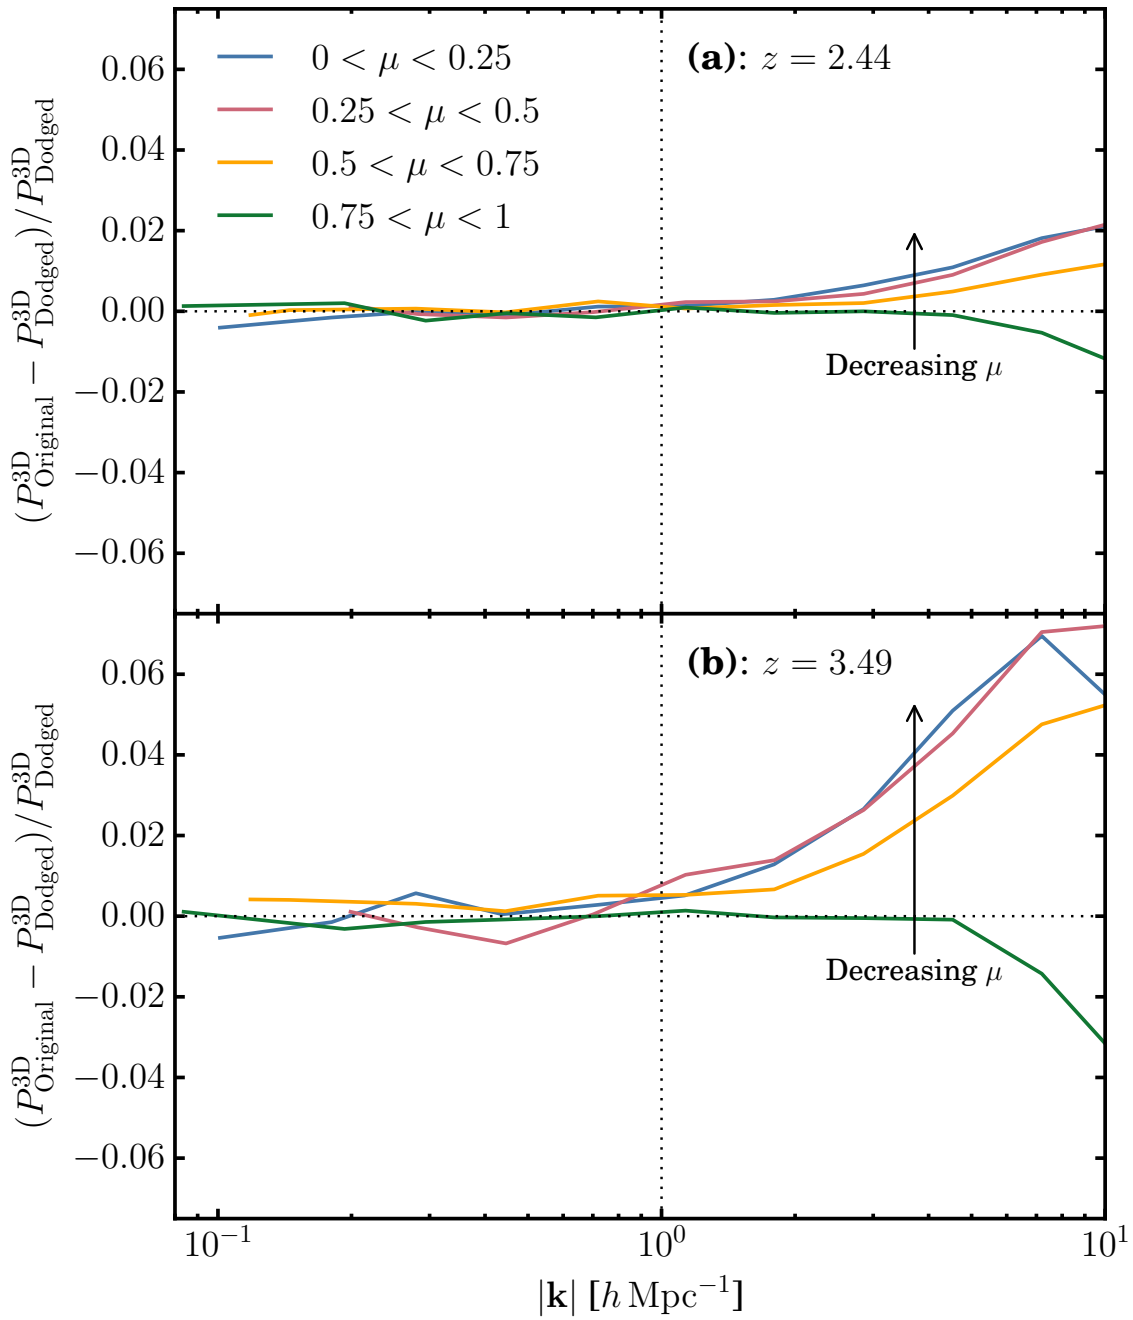

Figure B.2: The fractional error in the estimation of the power spectrum of a Gaussian random field (GRF) due to "dodging" lines-of-sight. We replicate in a GRF the exact movement of lines-of-sight that we carry out in our simulation box in order to dodge HCD absorbers. We then calculate the error in the estimation of the power spectrum due to ignoring the changes in position of lines-of-sight when calculating the necessary (fast) Fourier transforms. We note that the error remains small (sub-percent) for scales of interest in our study ($|\mathbf{k}| < 1\,h\,\text{Mpc}^{-1}$). **(a)** *above*: we replicate the dodging in our box at $z = 2.44$; **(b)** *below*: at $z = 3.49$.

we consider. There is more error at higher redshift because there is more dodging (see Fig. B.1), but the error remains sub-percent for scales of interest for our study ($|\mathbf{k}| < 1\,h\,\mathrm{Mpc}^{-1}$). This figure should be compared to Fig. 5.3 *(b)* (the fractional effect of a mock residual contamination of HCD absorbers on the 3D Lyman-alpha forest flux power spectrum); for the scales of interest, the effect of the HCD absorbers that we wish to measure remains much larger than the error arising from the dodging. We therefore ignore only scales smaller than this cut-off in our analysis.

B.2 Details of modelling and MCMC sampling

In § 5.4.4, we wish to sample the joint posterior probability distribution of the parameters of our proposed model in Eq. (5.5), given our measured 3D flux power spectra. We therefore require a likelihood function for our simulated data given the model. We use a Gaussian likelihood function and assume the covariance matrix to be diagonal (*i. e.*, we ignore correlations between power spectrum bins). Each $\hat{\delta}_{\mathrm{Flux}}(\mathbf{k}_n)$ is well approximated by a Gaussian random variable and so, as explained in § 5.4.3, our flux power spectrum estimates in each bin i are the sums of the squares of N_i such variables. It follows that each of the elements of our data-vector is chi-squared distributed (ignoring the slightly different amplitudes of and possible correlations between the Fourier modes within each bin):[1]

$$P^{\mathrm{3D}}_{\mathrm{Flux},i} \sim \frac{P^{\mathrm{3D}}_{\mathrm{True},i}}{N_i}\chi^2(N_i), \tag{B.1}$$

where $P^{\mathrm{3D}}_{\mathrm{True},i}$ is the true (ensemble) value of the 3D flux power spectrum in bin i. The variance of this distribution is $2(P^{\mathrm{3D}}_{\mathrm{True},i})^2/N_i$ and these form the (diagonal) elements of our covariance matrix (substituting $P^{\mathrm{3D}}_{\mathrm{Flux},i}$ for $P^{\mathrm{3D}}_{\mathrm{True},i}$, which is otherwise *a priori* unknown). To constrain the contamination parameters while marginalising over intrinsic Lyman-alpha forest bias parameters, we combine the likelihoods for the $P^{\mathrm{3D}}_{\mathrm{Contaminated}}$ and $P^{\mathrm{3D}}_{\mathrm{Forest}}$ data-vectors (using Eq. (5.1) to model $P^{\mathrm{3D}}_{\mathrm{Forest}}$); we ignore correlations between the two data-vectors and simply add the log-likelihoods. This is sufficient for the level of accuracy of our study.

Our model requires evaluation of $P^{\mathrm{3D}}_{\mathrm{Linear}}$ in each bin; to improve the comparison to $P^{\mathrm{3D}}_{\mathrm{Flux},i}$, we similarly evaluate $P^{\mathrm{3D}}_{\mathrm{Linear}}$ at each individual mode and bin in the same way. We associate with each bin the average values of $|\mathbf{k}|$ and μ from the contributing modes. We use uniform prior probability distributions for the Lyman-alpha forest bias parameters b_{Forest} and β_{Forest} and use Gaussian priors for the HCD absorber bias parameters b_{HCD} and β_{HCD}, which are otherwise poorly

[1] For practical purposes, for $N_i > 50$, the distribution is close to a Gaussian distribution.

constrained. The mean and 1σ values of the Gaussian priors on $b_{\rm HCD}$ are (following the best-fit values on the total $b_{\rm HCD}$ from Bautista et al. 2017) *from top to bottom, left to right* in Fig. 5.4, (a) -0.0005 ± 0.0002; (b) -0.003 ± 0.001; (c) -0.007 ± 0.003; (d) -0.016 ± 0.006; (e) -0.0007 ± 0.0002; (f) -0.005 ± 0.002; (g) -0.012 ± 0.004; (h) -0.022 ± 0.009. They are estimated by dividing the total $b_{\rm HCD}$ from Bautista et al. (2017) by the relative rest-frame equivalent widths of the damping wings of each absorber category; they are scaled up at higher redshift by the increased amount of HCD absorption (estimated from the fraction of contaminated spectra). Following the best-fit values found by Bautista et al. (2017), we place a Gaussian prior on $\beta_{\rm HCD} = 0.7 \pm 0.2$. Our prior distributions for these contamination parameters are almost exactly returned in their marginalised 1D posterior distributions (details are given in § 5.5.2), *i. e.*, we are very insensitive to the amplitude of the effect of HCD absorbers. The shape of the scale-dependent bias is fully determined by the physics of the Voigt absorption profiles and the CDDF (see § 5.3.2). We investigate the suitability of our model for the scale-dependent bias arising from the absorption profiles of HCD absorbers as a function of column density by repeating the posterior sampling for the $P^{\rm 3D}_{\rm Contaminated}$ constructed for each HCD absorber category.

We sample the posterior distributions using a Markov chain Monte Carlo method, specifically `emcee` (Foreman-Mackey et al., 2013), an implementation of the affine-invariant MCMC sampler. We initialise our chains uniformly within the (non-zero) bounds of our prior distributions and test for convergence using the Gelman-Rubin statistic (Brooks and Gelman, 1998; Gelman and Rubin, 1992).

B.3 Tests of robustness of inference of bias parameters of the Lyman-alpha forest

Figure B.3 shows the results of testing how changing the smallest scale that we include in our analysis $|\mathbf{k}|_{\rm max}$ affects the (marginalised 1D) posterior distributions inferred for the bias parameters of the Lyman-alpha forest. The 1σ credible intervals on the combination $b_{\rm Forest}(1+\beta_{\rm Forest})$ and $\beta_{\rm Forest}$ increase as $|\mathbf{k}|_{\rm max}$ decreases because the number of modes remaining on scales larger than $|\mathbf{k}|_{\rm max}$ falls off quite sharply as $|\mathbf{k}|_{\rm max}$ is reduced. Although the combination $b_{\rm Forest}(1+\beta_{\rm Forest})$ is sampled, rather than $b_{\rm Forest}$ alone, because it is less correlated with $\beta_{\rm Forest}$, there is evidently still correlation: as $\beta_{\rm Forest}$ decreases with $|\mathbf{k}|_{\rm max}$, so does also the amplitude of $b_{\rm Forest}(1+\beta_{\rm Forest})$. Nonetheless, the posteriors of both bias parameters are statistically consistent for all the values of

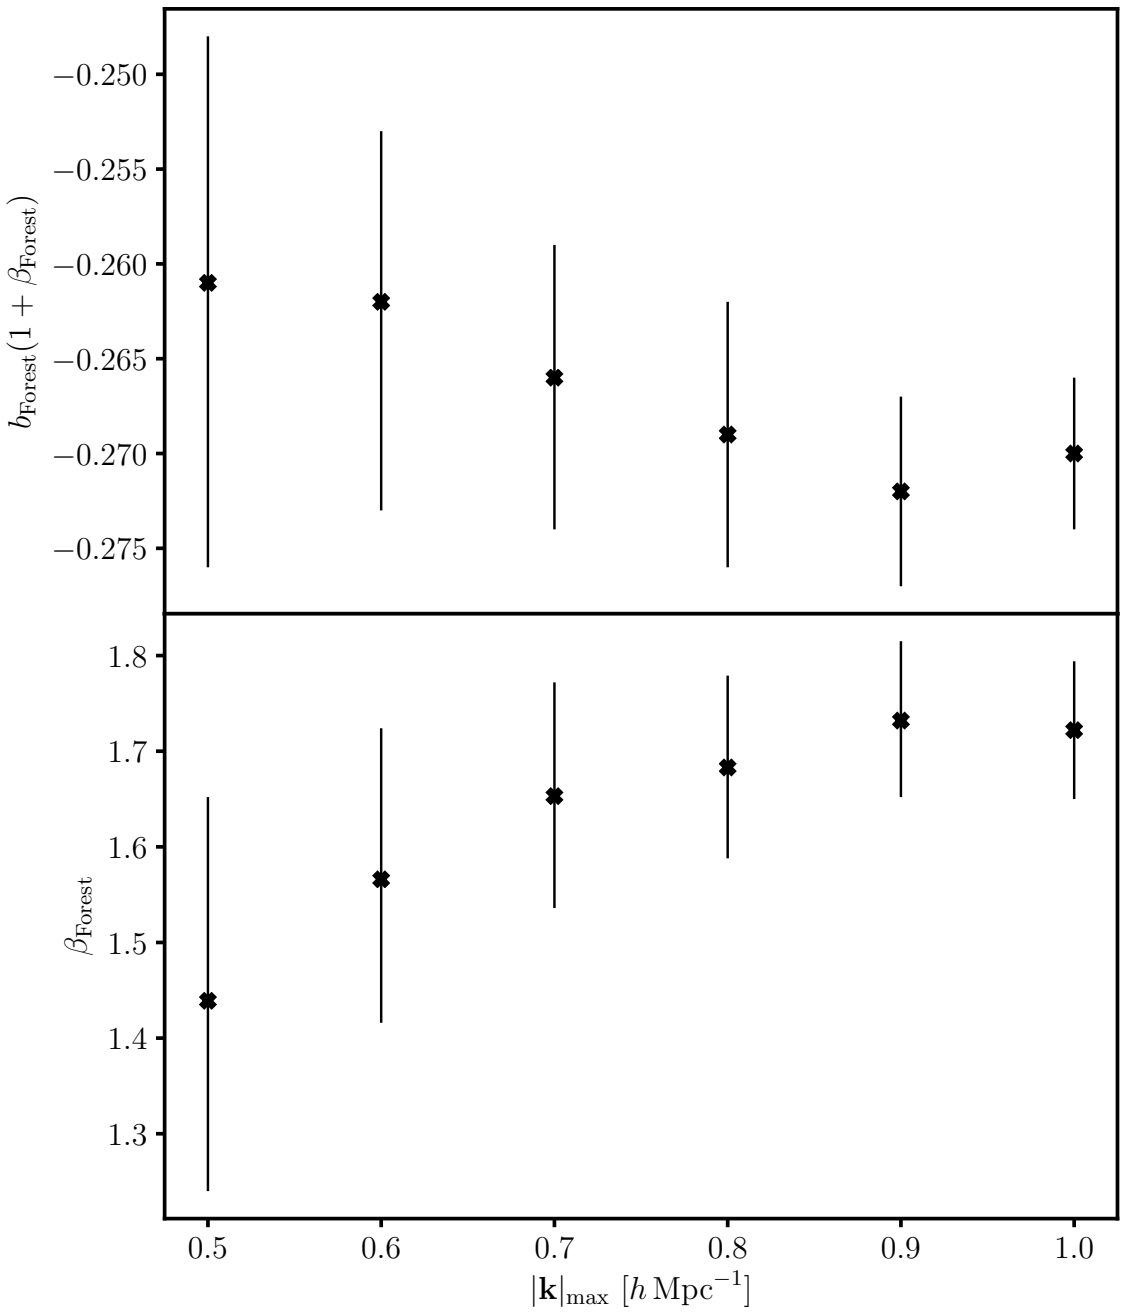

Figure B.3: The maximum posterior values with the 1σ credible intervals of the bias b_{Forest} (*above*) and redshift space distortion β_{Forest} (*below*) parameters of the Lyman-alpha forest, as inferred from our simulation box at $z = 2.44$, using different values of $|\mathbf{k}|_{\mathrm{max}}$, the smallest scale included in our analysis. Following *e. g.*, Bautista et al. (2017); Slosar et al. (2011), we sample the combination $b_{\mathrm{Forest}}(1 + \beta_{\mathrm{Forest}})$, which is less correlated with β_{Forest}. We find that our marginalised parameter posteriors are statistically consistent, irregardless of the smallest scale at which we cut our data vector.

$|\mathbf{k}|_{\max}$ that we consider, suggesting a degree of robustness in our inference on these parameters. Moreover, our conclusions on the scale-dependence of the effect of HCD absorbers on correlations in the 3D Lyman-alpha forest are insensitive to the overall amplitude, including the biases of the Lyman-alpha forest.

Bibliography

K. Abazajian, J. K. Adelman-McCarthy, M. A. Agüeros, S. S. Allam, S. F. Anderson, J. Annis, N. A. Bahcall, I. K. Baldry, S. Bastian, A. Berlind, M. Bernardi, M. R. Blanton, N. Blythe, J. J. Bochanski, Jr., W. N. Boroski, H. Brewington, J. W. Briggs, J. Brinkmann, R. J. Brunner, T. Budavári, L. N. Carey, M. A. Carr, F. J. Castander, K. Chiu, M. J. Collinge, A. J. Connolly, K. R. Covey, I. Csabai, J. J. Dalcanton, S. Dodelson, M. Doi, F. Dong, D. J. Eisenstein, M. L. Evans, X. Fan, P. D. Feldman, D. P. Finkbeiner, S. D. Friedman, J. A. Frieman, M. Fukugita, R. R. Gal, B. Gillespie, K. Glazebrook, C. F. Gonzalez, J. Gray, E. K. Grebel, L. Grodnicki, J. E. Gunn, V. K. Gurbani, P. B. Hall, L. Hao, D. Harbeck, F. H. Harris, H. C. Harris, M. Harvanek, S. L. Hawley, T. M. Heckman, J. F. Helmboldt, J. S. Hendry, G. S. Hennessy, R. B. Hindsley, D. W. Hogg, D. J. Holmgren, J. A. Holtzman, L. Homer, L. Hui, S.-i. Ichikawa, T. Ichikawa, J. P. Inkmann, Ž. Ivezić, S. Jester, D. E. Johnston, B. Jordan, W. P. Jordan, A. M. Jorgensen, M. Jurić, G. Kauffmann, S. M. Kent, S. J. Kleinman, G. R. Knapp, A. Y. Kniazev, R. G. Kron, J. Krzesiński, P. Z. Kunszt, N. Kuropatkin, D. Q. Lamb, H. Lampeitl, B. E. Laubscher, B. C. Lee, R. F. Leger, N. Li, A. Lidz, H. Lin, Y.-S. Loh, D. C. Long, J. Loveday, R. H. Lupton, T. Malik, B. Margon, P. M. McGehee, T. A. McKay, A. Meiksin, G. A. Miknaitis, B. K. Moorthy, J. A. Munn, T. Murphy, R. Nakajima, V. K. Narayanan, T. Nash, E. H. Neilsen, Jr., H. J. Newberg, P. R. Newman, R. C. Nichol, T. Nicinski, M. Nieto-Santisteban, A. Nitta, M. Odenkirchen, S. Okamura, J. P. Ostriker, R. Owen, N. Padmanabhan, J. Peoples, J. R. Pier, B. Pindor, A. C. Pope, T. R. Quinn, R. R. Rafikov, S. N. Raymond, G. T. Richards, M. W. Richmond, H.-W. Rix, C. M. Rockosi, J. Schaye, D. J. Schlegel, D. P. Schneider, J. Schroeder, R. Scranton, M. Sekiguchi, U. Seljak, G. Sergey, B. Sesar, E. Sheldon, K. Shimasaku, W. A. Siegmund, N. M. Silvestri, A. J. Sinisgalli, E. Sirko, J. A. Smith, V. Smolčić, S. A. Snedden, A. Stebbins, C. Steinhardt, G. Stinson, C. Stoughton, I. V. Strateva, M. A. Strauss, M. SubbaRao, A. S. Szalay, I. Szapudi, P. Szkody, L. Tasca, M. Tegmark, A. R. Thakar, C. Tremonti, D. L. Tucker, A. Uomoto, D. E. Vanden Berk, J. Vandenberg, M. S. Vogeley, W. Voges, N. P. Vogt, L. M. Walkowicz, D. H. Weinberg, A. A. West, S. D. M. White, B. C. Wilhite, B. Willman, Y. Xu, B. Yanny, J. Yarger, N. Yasuda, C.-W. Yip, D. R. Yocum, D. G. York, N. L. Zakamska, I. Zehavi, W. Zheng, S. Zibetti, and D. B. Zucker. The First Data Release of the Sloan Digital Sky Survey. AJ, 126:2081–2086, Oct. 2003. doi: 10.1086/378165.

B. Abolfathi, D. S. Aguado, G. Aguilar, C. Allende Prieto, A. Almeida, T. Tasnim Ananna, F. Anders, S. F. Anderson, B. H. Andrews, B. Anguiano, and et al. The Fourteenth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the extended Baryon Oscillation Sky Survey and from the second phase of the Apache Point Observatory Galactic Evolution Experiment. *ArXiv e-prints*, July 2017.

P. A. R. Ade, N. Aghanim, Z. Ahmed, R. W. Aikin, K. D. Alexander, M. Arnaud, J. Aumont, C. Baccigalupi, A. J. Banday, D. Barkats, R. B. Barreiro, J. G. Bartlett, N. Bartolo, E. Battaner, K. Benabed,

A. Benoît, A. Benoit-Lévy, S. J. Benton, J.-P. Bernard, M. Bersanelli, P. Bielewicz, C. A. Bischoff, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, J. A. Brevik, M. Bucher, I. Buder, E. Bullock, C. Burigana, R. C. Butler, V. Buza, E. Calabrese, J.-F. Cardoso, A. Catalano, A. Challinor, R.-R. Chary, H. C. Chiang, P. R. Christensen, L. P. L. Colombo, C. Combet, J. Connors, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J.-M. Delouis, F.-X. Désert, C. Dickinson, J. M. Diego, H. Dole, S. Donzelli, O. Doré, M. Douspis, C. D. Dowell, L. Duband, A. Ducout, J. Dunkley, X. Dupac, C. Dvorkin, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, E. Falgarone, J. P. Filippini, F. Finelli, S. Fliescher, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frejsel, S. Galeotta, S. Galli, K. Ganga, T. Ghosh, M. Giard, E. Gjerløw, S. R. Golwala, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, M. Halpern, F. K. Hansen, D. Hanson, D. L. Harrison, M. Hasselfield, G. Helou, S. Henrot-Versillé, D. Herranz, S. R. Hildebrandt, G. C. Hilton, E. Hivon, M. Hobson, W. A. Holmes, W. Hovest, V. V. Hristov, K. M. Huffenberger, H. Hui, G. Hurier, K. D. Irwin, A. H. Jaffe, T. R. Jaffe, J. Jewell, W. C. Jones, M. Juvela, A. Karakci, K. S. Karkare, J. P. Kaufman, B. G. Keating, S. Kefeli, E. Keihänen, S. A. Kernasovskiy, R. Keskitalo, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, J. M. Kovac, N. Krachmalnicoff, M. Kunz, C. L. Kuo, H. Kurki-Suonio, G. Lagache, A. Lähteenmäki, J.-M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, E. M. Leitch, R. Leonardi, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Linden-Vørnle, M. López-Caniego, P. M. Lubin, M. Lueker, J. F. Macías-Pérez, B. Maffei, D. Maino, N. Mandolesi, A. Mangilli, M. Maris, P. G. Martin, E. Martínez-González, S. Masi, P. Mason, S. Matarrese, K. G. Megerian, P. R. Meinhold, A. Melchiorri, L. Mendes, A. Mennella, M. Migliaccio, S. Mitra, M.-A. Miville-Deschênes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. T. Nguyen, H. U. Nørgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, R. O'Brient, R. W. Ogburn, A. Orlando, L. Pagano, F. Pajot, R. Paladini, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, T. J. Pearson, O. Perdereau, L. Perotto, V. Pettorino, F. Piacentini, M. Piat, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, N. Ponthieu, G. W. Pratt, S. Prunet, C. Pryke, J.-L. Puget, J. P. Rachen, W. T. Reach, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, A. Renzi, S. Richter, I. Ristorcelli, G. Rocha, M. Rossetti, G. Roudier, M. Rowan-Robinson, J. A. Rubiño Martín, B. Rusholme, M. Sandri, D. Santos, M. Savelainen, G. Savini, R. Schwarz, D. Scott, M. D. Seiffert, C. D. Sheehy, L. D. Spencer, Z. K. Staniszewski, V. Stolyarov, R. Sudiwala, R. Sunyaev, D. Sutton, A.-S. Suur-Uski, J.-F. Sygnet, J. A. Tauber, G. P. Teply, L. Terenzi, K. L. Thompson, L. Toffolatti, J. E. Tolan, M. Tomasi, M. Tristram, M. Tucci, A. D. Turner, L. Valenziano, J. Valiviita, B. Van Tent, L. Vibert, P. Vielva, A. G. Vieregg, F. Villa, L. A. Wade, B. D. Wandelt, R. Watson, A. C. Weber, I. K. Wehus, M. White, S. D. M. White, J. Willmert, C. L. Wong, K. W. Yoon, D. Yvon, A. Zacchei, and A. Zonca. Joint analysis of bicep2/keck array and planck data. *Phys. Rev. Lett.*, 114:101301, Mar 2015. doi: 10.1103/PhysRevLett.114.101301. URL `https://link.aps.org/doi/10.1103/PhysRevLett.114.101301`.

R. Adhikari, M. Agostini, N. A. Ky, T. Araki, M. Archidiacono, M. Bahr, J. Baur, J. Behrens, F. Bezrukov, P. S. Bhupal Dev, D. Borah, A. Boyarsky, A. de Gouvea, C. A. d. S. Pires, H. J. de Vega, A. G. Dias, P. Di Bari, Z. Djurcic, K. Dolde, H. Dorrer, M. Durero, O. Dragoun, M. Drewes, G. Drexlin, C. E. Düllmann, K. Eberhardt, S. Eliseev, C. Enss, N. W. Evans, A. Faessler, P. Filianin, V. Fischer, A. Fleischmann, J. A. Formaggio, J. Franse, F. M. Fraenkle, C. S. Frenk, G. Fuller, L. Gastaldo, A. Garzilli, C. Giunti, F. Glück, M. C. Goodman, M. C. Gonzalez-Garcia, D. Gorbunov, J. Hamann, V. Hannen, S. Hannestad, S. H. Hansen, C. Hassel, J. Heeck, F. Hofmann, T. Houdy, A. Huber, D. Iakubovskyi, A. Ianni, A. Ibarra, R. Jacobsson, T. Jeltema, J. Jochum, S. Kempf, T. Kieck, M. Korzeczek, V. Kornoukhov, T. Lachenmaier, M. Laine, P. Langacker, T. Lasserre, J. Lesgourgues, D. Lhuillier, Y. F. Li, W. Liao, A. W. Long, M. Maltoni, G. Mangano, N. E. Mavromatos, N. Menci, A. Merle, S. Mertens, A. Mirizzi, B. Monreal, A. Nozik, A. Neronov, V. Niro, Y. Novikov, L. Oberauer, E. Otten, N. Palanque-Delabrouille, M. Pallavicini, V. S. Pantuev, E. Papastergis, S. Parke, S. Pascoli, S. Pastor, A. Patwardhan, A. Pilaftsis, D. C. Radford, P. C.-O. Ranitzsch, O. Rest, D. J. Robinson, P. S. Rodrigues da Silva, O. Ruchayskiy, N. G. Sanchez, M. Sasaki, N. Saviano, A. Schneider, F. Schneider, T. Schwetz, S. Schönert, S. Scholl, F. Shankar,

R. Shrock, N. Steinbrink, L. Strigari, F. Suekane, B. Suerfu, R. Takahashi, N. T. H. Van, I. Tkachev, M. Totzauer, Y. Tsai, C. G. Tully, K. Valerius, J. W. F. Valle, D. Venos, M. Viel, M. Vivier, M. Y. Wang, C. Weinheimer, K. Wendt, L. Winslow, J. Wolf, M. Wurm, Z. Xing, S. Zhou, and K. Zuber. A White Paper on keV sterile neutrino Dark Matter. J. Cosmology Astropart. Phys., 1:025, Jan. 2017. doi: 10.1088/1475-7516/2017/01/025.

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 `https://link.aps.org/doi/10.1103/PhysRevLett.48.1220`.

C. Alcock and B. Paczynski. An evolution free test for non-zero cosmological constant. Nature, 281: 358, Oct. 1979. doi: 10.1038/281358a0.

A. Amblard, C. Vale, and M. White. Weak lensing of the CMB by large-scale structure. New A, 9: 687–704, Oct. 2004. doi: 10.1016/j.newast.2004.05.003.

P. Anninos, M. L. Norman, and D. A. Clarke. Hierarchical numerical cosmology with hydrodynamics: Methods and code tests. ApJ, 436:11–22, Nov. 1994. doi: 10.1086/174876.

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. *ACHA*, 7:1–30, 1999.

A. Arinyo-i-Prats, J. Miralda-Escudé, M. Viel, and R. Cen. The non-linear power spectrum of the Lyman alpha forest. J. Cosmology Astropart. Phys., 12:017, Dec. 2015. doi: 10.1088/1475-7516/2015/12/017.

E. Armengaud, N. Palanque-Delabrouille, C. Yèche, D. J. E. Marsh, and J. Baur. Constraining the mass of light bosonic dark matter using SDSS Lyman-α forest. *ArXiv e-prints*, Mar. 2017.

P. Baldi, G. Kerkyacharian, D. Marinucci, and D. Picard. Asymptotics for spherical needlets. *Ann. Stat.*, 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. *MNRAS*, 318:475–481, 2000.

S. Basak and J. Delabrouille. A needlet internal linear combination analysis of WMAP 7-year data: estimation of CMB temperature map and power spectrum. *MNRAS*, 419:1163–1175, Jan. 2012. doi: 10.1111/j.1365-2966.2011.19770.x.

S. Basak and J. Delabrouille. A needlet ILC analysis of WMAP 9-year polarization data: CMB polarization power spectra. MNRAS, 435:18–29, Oct. 2013. doi: 10.1093/mnras/stt1158.

D. Baumann. TASI Lectures on Inflation. *ArXiv e-prints*, July 2009.

J. E. Bautista, N. G. Busca, J. Guy, J. Rich, M. Blomqvist, H. du Mas des Bourboux, M. M. Pieri, A. Font-Ribera, S. Bailey, T. Delubac, D. Kirkby, J.-M. Le Goff, D. Margala, A. Slosar, J. A. Vazquez, J. R. Brownstein, K. S. Dawson, D. J. Eisenstein, J. Miralda-Escudé, P. Noterdaeme, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, N. P. Ross, D. P. Schneider, D. H. Weinberg, and C. Yèche. Measurement of BAO correlations at $z = 2.3$ with SDSS DR12 \lya-Forests. *ArXiv e-prints*, Feb. 2017.

G. D. Becker, J. S. Bolton, M. G. Haehnelt, and W. L. W. Sargent. Detection of extended He II reionization in the temperature evolution of the intergalactic medium. MNRAS, 410:1096–1112, Jan. 2011. doi: 10.1111/j.1365-2966.2010.17507.x.

R. H. Becker, X. Fan, R. L. White, M. A. Strauss, V. K. Narayanan, R. H. Lupton, J. E. Gunn, J. Annis, N. A. Bahcall, J. Brinkmann, A. J. Connolly, I. Csabai, P. C. Czarapata, M. Doi, T. M. Heckman, G. S. Hennessy, Ž. Ivezić, G. R. Knapp, D. Q. Lamb, T. A. McKay, J. A. Munn, T. Nash, R. Nichol, J. R. Pier, G. T. Richards, D. P. Schneider, C. Stoughton, A. S. Szalay, A. R. Thakar, and D. G. York. Evidence for Reionization at z˜6: Detection of a Gunn-Peterson Trough in a z=6.28 Quasar. AJ, 122:2850–2857, Dec. 2001. doi: 10.1086/324231.

L. Bedini, D. Herranz, E. Salerno, C. Baccigalupi, E. E. Kuruouglu, and A. Tonazzini. Separation of correlated astrophysical sources using multiple-lag data covariance matrices. *EURASIP Journal on Applied Signal Processing*, 2005:2400–2412, Aug. 2005. doi: 10.1155/ASP.2005.2400.

C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, D. N. Spergel, G. S. Tucker, E. Wollack, E. L. Wright, C. Barnes, M. R. Greason, R. S. Hill, E. Komatsu, M. R. Nolta, N. Odegard, H. V. Peiris, L. Verde, and J. L. Weiland. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Preliminary Maps and Basic Results. ApJS, 148:1–27, Sept. 2003a. doi: 10.1086/377253.

C. L. Bennett, R. S. Hill, G. Hinshaw, M. R. Nolta, N. Odegard, L. Page, D. N. Spergel, J. L. Weiland, E. L. Wright, M. Halpern, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. Wollack. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Emission. *ApJS*, 148:97–117, Sept. 2003b. doi: 10.1086/377252.

A. Benoit-Lévy, T. Déchelette, K. Benabed, J.-F. Cardoso, D. Hanson, and S. Prunet. Full-sky CMB lensing reconstruction in presence of sky-cuts. A&A, 555:A37, July 2013. doi: 10.1051/0004-6361/201321048.

BICEP2 and Keck Array Collaborations, P. A. R. Ade, Z. Ahmed, R. W. Aikin, K. D. Alexander, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, J. A. Brevik, I. Buder, E. Bullock, V. Buza, J. Connors, B. P. Crill, C. D. Dowell, C. Dvorkin, L. Duband, J. P. Filippini, S. Fliescher, S. R. Golwala, M. Halpern, S. Harrison, M. Hasselfield, S. R. Hildebrandt, G. C. Hilton, V. V. Hristov, H. Hui, K. D. Irwin, K. S. Karkare, J. P. Kaufman, B. G. Keating, S. Kefeli, S. A. Kernasovskiy, J. M. Kovac, C. L. Kuo, E. M. Leitch, M. Lueker, P. Mason, K. G. Megerian, 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, K. L. Thompson, J. E. Tolan, A. D. Turner, A. G. Vieregg, A. C. Weber, J. Willmert, C. L. Wong, and K. W. Yoon. BICEP2/Keck Array V: Measurements of B-mode Polarization at Degree Angular Scales and 150 GHz by the Keck Array. ApJ, 811:126, Oct. 2015. doi: 10.1088/0004-637X/811/2/126.

BICEP2 Collaboration, Keck Array Collaboration, P. A. R. Ade, Z. Ahmed, R. W. Aikin, K. D. Alexander, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, R. Bowens-Rubin, J. A. Brevik, I. Buder, E. Bullock, V. Buza, J. Connors, B. P. Crill, L. Duband, C. Dvorkin, J. P. Filippini, S. Fliescher, J. Grayson, M. Halpern, S. Harrison, G. C. Hilton, H. Hui, K. D. Irwin, K. S. Karkare, E. Karpel, J. P. Kaufman, B. G. Keating, S. Kefeli, S. A. Kernasovskiy, J. M. Kovac, C. L. Kuo, E. M. Leitch, M. Lueker, K. G. Megerian, C. B. Netterfield, H. T. Nguyen, R. O'Brient, R. W. Ogburn, A. Orlando, C. Pryke, S. Richter, R. Schwarz, C. D. Sheehy, Z. K. Staniszewski, B. Steinbach, R. V. Sudiwala, G. P. Teply, K. L. Thompson, J. E. Tolan, C. Tucker, A. D. Turner, A. G. Vieregg, A. C. Weber, D. V. Wiebe, J. Willmert, C. L. Wong, W. L. K. Wu, and K. W. Yoon. Improved Constraints on Cosmology and Foregrounds from BICEP2 and Keck Array Cosmic Microwave Background Data with Inclusion of 95 GHz Band. *Physical Review Letters*, 116(3):031302, Jan. 2016. doi: 10.1103/PhysRevLett.116.031302.

S. Bird. FSFE: Fake Spectra Flux Extractor. Astrophysics Source Code Library, Oct. 2017.

S. Bird, M. Vogelsberger, M. Haehnelt, D. Sijacki, S. Genel, P. Torrey, V. Springel, and L. Hernquist. Damped Lyman α absorbers as a probe of stellar feedback. MNRAS, 445:2313–2324, Dec. 2014. doi: 10.1093/mnras/stu1923.

S. Bird, M. Haehnelt, M. Neeleman, S. Genel, M. Vogelsberger, and L. Hernquist. Reproducing the kinematics of damped Lyman α systems. MNRAS, 447:1834–1846, Feb. 2015. doi: 10.1093/mnras/stu2542.

S. Bird, I. Cholis, J. B. Muñoz, Y. Ali-Haïmoud, M. Kamionkowski, E. D. Kovetz, A. Raccanelli, and A. G. Riess. Did LIGO Detect Dark Matter? *Physical Review Letters*, 116(20):201301, May 2016. doi: 10.1103/PhysRevLett.116.201301.

S. Bird, R. Garnett, and S. Ho. Statistical properties of damped Lyman-alpha systems from Sloan Digital Sky Survey DR12. MNRAS, 466:2111–2122, Apr. 2017. doi: 10.1093/mnras/stw3246.

J. Bobin, Y. Moudden, J.-L. Starck, J. Fadili, and N. Aghanim. SZ and CMB reconstruction using generalized morphological component analysis. *Statistical Methodology*, 5:307–317, July 2008. doi: 10.1016/j.stamet.2007.10.003.

J. Bobin, J.-L. Starck, F. Sureau, and S. Basak. Sparse component separation for accurate cosmic microwave background estimation. A&A, 550:A73, Feb. 2013. doi: 10.1051/0004-6361/201219781.

N. W. Boggess, J. C. Mather, R. Weiss, C. L. Bennett, E. S. Cheng, E. Dwek, S. Gulkis, M. G. Hauser, M. A. Janssen, T. Kelsall, S. S. Meyer, S. H. Moseley, T. L. Murdock, R. A. Shafer, R. F. Silverberg, G. F. Smoot, D. T. Wilkinson, and E. L. Wright. The COBE mission - Its design and performance two years after launch. ApJ, 397:420–429, Oct. 1992. doi: 10.1086/171797.

J. R. Bond and G. Efstathiou. The statistics of cosmic background radiation fluctuations. MNRAS, 226:655–687, June 1987.

H. Bondi and T. Gold. The Steady-State Theory of the Expanding Universe. MNRAS, 108:252, 1948. doi: 10.1093/mnras/108.3.252.

A. M. Brooks and A. Zolotov. Why Baryons Matter: The Kinematics of Dwarf Spheroidal Satellites. ApJ, 786:87, May 2014. doi: 10.1088/0004-637X/786/2/87.

S. P. Brooks and A. Gelman. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4):434–455, 1998. doi: 10.1080/10618600.1998.10474787.

E. F. Bunn, M. Zaldarriaga, M. Tegmark, and A. de Oliveira-Costa. E/B decomposition of finite pixelized CMB maps. Phys. Rev. D, 67(2):023501, Jan. 2003. doi: 10.1103/PhysRevD.67.023501.

N. G. Busca, T. Delubac, J. Rich, S. Bailey, A. Font-Ribera, D. Kirkby, J.-M. Le Goff, M. M. Pieri, A. Slosar, É. Aubourg, J. E. Bautista, D. Bizyaev, M. Blomqvist, A. S. Bolton, J. Bovy, H. Brewington, A. Borde, J. Brinkmann, B. Carithers, R. A. C. Croft, K. S. Dawson, G. Ebelke, D. J. Eisenstein, J.-C. Hamilton, S. Ho, D. W. Hogg, K. Honscheid, K.-G. Lee, B. Lundgren, E. Malanushenko, V. Malanushenko, D. Margala, C. Maraston, K. Mehta, J. Miralda-Escudé, A. D. Myers, R. C. Nichol, P. Noterdaeme, M. D. Olmstead, D. Oravetz, N. Palanque-Delabrouille, K. Pan, I. Pâris, W. J. Percival, P. Petitjean, N. A. Roe, E. Rollinde, N. P. Ross, G. Rossi, D. J. Schlegel, D. P. Schneider, A. Shelden, E. S. Sheldon, A. Simmons, S. Snedden, J. L. Tinker, M. Viel, B. A. Weaver, D. H. Weinberg, M. White, C. Yèche, and D. G. York. Baryon acoustic oscillations in the Lyα forest of BOSS quasars. A&A, 552:A96, Apr. 2013. doi: 10.1051/0004-6361/201220724.

F. Capozzi, E. Lisi, A. Marrone, D. Montanino, and A. Palazzo. Neutrino masses and mixings: Status of known and unknown 3ν parameters. *Nuclear Physics B*, 908:218–234, July 2016. doi: 10.1016/j.nuclphysb.2016.02.016.

J.-F. Cardoso, M. Martin, J. Delabrouille, M. Betoule, and G. Patanchon. Component separation with flexible models. Application to the separation of astrophysical emissions. *ArXiv e-prints*, Mar. 2008.

B. J. Carr and S. W. Hawking. Black holes in the early Universe. MNRAS, 168:399–416, Aug. 1974. doi: 10.1093/mnras/168.2.399.

S. M. Carroll. The Cosmological Constant. *Living Reviews in Relativity*, 4:1, Feb. 2001. doi: 10.12942/lrr-2001-1.

S. Caucci, S. Colombi, C. Pichon, E. Rollinde, P. Petitjean, and T. Sousbie. Recovering the topology of the intergalactic medium at z ˜ 2. MNRAS, 386:211–229, May 2008. doi: 10.1111/j.1365-2966.2008.13016.x.

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. MNRAS, 326:1243–1248, Oct. 2001. doi: 10.1111/j.1365-2966.2001.04641.x.

R. Cen. The Nature of Damped Lyα Systems and Their Hosts in the Standard Cold Dark Matter Universe. ApJ, 748:121, Apr. 2012. doi: 10.1088/0004-637X/748/2/121.

J. Y. H. Chan, B. Leistedt, T. D. Kitching, and J. D. McEwen. Second-generation curvelets on the sphere. *IEEE TSP, submitted*, 2015.

T. Clifton, P. G. Ferreira, A. Padilla, and C. Skordis. Modified gravity and cosmology. Phys. Rep., 513:1–189, Mar. 2012. doi: 10.1016/j.physrep.2012.01.001.

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. A. C. Croft, D. H. Weinberg, N. Katz, and L. Hernquist. Recovery of the Power Spectrum of Mass Fluctuations from Observations of the Lyα Forest. ApJ, 495:44–62, Mar. 1998. doi: 10.1086/305289.

R. A. C. Croft, D. H. Weinberg, M. Pettini, L. Hernquist, and N. Katz. The Power Spectrum of Mass Fluctuations Measured from the Lyα Forest at Redshift z = 2.5. ApJ, 520:1–23, July 1999. doi: 10.1086/307438.

S. Das, T. A. Marriage, 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, M. J. Devlin, S. R. Dicker, W. Bertrand Doriese, J. Dunkley, 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, R. Hlozek, K. M. Huffenberger, D. H. Hughes, J. P. Hughes, L. Infante, K. D. Irwin, J. Baptiste Juin, M. Kaul, J. Klein, A. Kosowsky, J. M. Lau, M. Limon, Y.-T. Lin, R. H. Lupton, D. Marsden, K. Martocci, 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. D. Sherwin, J. Sievers, 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: A Measurement of the Cosmic Microwave Background Power Spectrum at 148 and 218 GHz from the 2008 Southern Survey. ApJ, 729:62, Mar. 2011a. doi: 10.1088/0004-637X/729/1/62.

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 2011b. doi: 10.1103/PhysRevLett.107.021301.

M. Davis, J. Huchra, D. W. Latham, and J. Tonry. A survey of galaxy redshifts. II - The large scale space distribution. ApJ, 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. AJ, 145:10, Jan. 2013. doi: 10.1088/0004-6256/145/1/10.

P. de Bernardis, P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, K. Coble, B. P. Crill, G. De Gasperis, P. C. Farese, P. G. Ferreira, K. Ganga, M. Giacometti, E. Hivon, V. V. Hristov, A. Iacoangeli, A. H. Jaffe, A. E. Lange, L. Martinis, S. Masi, P. V. Mason, P. D. Mauskopf, A. Melchiorri, L. Miglio, T. Montroy, C. B. Netterfield, E. Pascale, F. Piacentini, D. Pogosyan, S. Prunet, S. Rao, G. Romeo, J. E. Ruhl, F. Scaramuzzi, D. Sforna, and N. Vittorio. A flat Universe from high-resolution maps of the cosmic microwave background radiation. Nature, 404:955–959, Apr. 2000. doi: 10.1038/35010035.

V. de Lapparent, M. J. Geller, and J. P. Huchra. A slice of the universe. ApJ, 302:L1–L5, Mar. 1986. doi: 10.1086/184625.

J. Delabrouille, J.-F. Cardoso, M. Le Jeune, M. Betoule, G. Fay, and F. Guilloux. A full sky, low foreground, high resolution CMB map from WMAP. *A&A*, 493:835–857, Jan. 2009. doi: 10.1051/0004-6361:200810514.

T. Delubac, J. E. Bautista, N. G. Busca, J. Rich, D. Kirkby, S. Bailey, A. Font-Ribera, A. Slosar, K.-G. Lee, M. M. Pieri, J.-C. Hamilton, É. Aubourg, M. Blomqvist, J. Bovy, J. Brinkmann, W. Carithers, K. S. Dawson, D. J. Eisenstein, S. G. A. Gontcho, J.-P. Kneib, J.-M. Le Goff, D. Margala, J. Miralda-Escudé, A. D. Myers, R. C. Nichol, P. Noterdaeme, R. O'Connell, M. D. Olmstead, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, N. P. Ross, G. Rossi, D. J. Schlegel, D. P. Schneider, D. H. Weinberg, C. Yèche, and D. G. York. Baryon acoustic oscillations in the Lyα forest of BOSS DR11 quasars. A&A, 574:A59, Feb. 2015. doi: 10.1051/0004-6361/201423969.

DESI Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland, and et al. The DESI Experiment Part I: Science,Targeting, and Survey Design. *ArXiv e-prints*, Oct. 2016a.

DESI Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland, and et al. The DESI Experiment Part II: Instrument Design. *ArXiv e-prints*, Oct. 2016b.

J. Dick, M. Remazeilles, and J. Delabrouille. Impact of calibration errors on CMB component separation using FastICA and ILC. *MNRAS*, 401:1602–1612, Jan. 2010. doi: 10.1111/j.1365-2966.2009.15798.x.

R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson. Cosmic Black-Body Radiation. ApJ, 142:414–419, July 1965. doi: 10.1086/148306.

S. Dodelson. *Modern cosmology*. Academic Press, San Diego, CA, 2003. URL `https://cds.cern.ch/record/1282338`.

A. G. Doroshkevich, E. V. Kotok, A. N. Poliudov, S. F. Shandarin, I. S. Sigov, and I. D. Novikov. Two-dimensional simulation of the gravitational system dynamics and formation of the large-scale structure of the universe. MNRAS, 192:321–337, Aug. 1980. doi: 10.1093/mnras/192.2.321.

B. T. Draine and A. Lazarian. Electric Dipole Radiation from Spinning Dust Grains. ApJ, 508:157–179, Nov. 1998. doi: 10.1086/306387.

J. Driscoll and D. Healy. Computing fourier transforms and convolutions on the 2-sphere. *Advances in Applied Mathematics*, 15(2):202 – 250, 1994. ISSN 0196-8858. doi: http://dx.doi.org/10.1006/aama.1994.1008. URL `http://www.sciencedirect.com/science/article/pii/S0196885884710086`.

H. du Mas des Bourboux, J.-M. Le Goff, M. Blomqvist, N. G. Busca, J. Guy, J. Rich, C. Yèche, J. E. Bautista, É. Burtin, K. S. Dawson, D. J. Eisenstein, A. Font-Ribera, D. Kirkby, J. Miralda-Escudé, P. Noterdaeme, I. Pâris, P. Petitjean, I. Pérez-Ràfols, M. M. Pieri, N. P. Ross, D. J. Schlegel, D. P. Schneider, A. Slosar, D. H. Weinberg, and P. Zarrouk. Baryon acoustic oscillations from the complete SDSS-III Lyα-quasar cross-correlation function at $z = 2.4$. *ArXiv e-prints*, Aug. 2017.

J. Dunkley, A. Amblard, C. Baccigalupi, M. Betoule, D. Chuss, A. Cooray, J. Delabrouille, C. Dickinson, G. Dobler, J. Dotson, H. K. Eriksen, D. Finkbeiner, D. Fixsen, P. Fosalba, A. Fraisse, C. Hirata, A. Kogut, J. Kristiansen, C. Lawrence, A. M. Magalha˜Es, M. A. Miville-Deschenes, S. Meyer, A. Miller, S. K. Naess, L. Page, H. V. Peiris, N. Phillips, E. Pierpaoli, G. Rocha, J. E. Vaillancourt, and L. Verde. Prospects for polarized foreground removal. In S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, and K. Smith, editors, *American Institute of Physics Conference Series*, volume 1141 of *American Institute of Physics Conference Series*, pages 222–264, June 2009. doi: 10.1063/1.3160888.

G. Efstathiou and J. W. Eastwood. On the clustering of particles in an expanding universe. MNRAS, 194:503–525, Feb. 1981. doi: 10.1093/mnras/194.3.503.

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. AJ, 142:72, Sept. 2011. doi: 10.1088/0004-6256/142/3/72.

H. K. Eriksen, A. J. Banday, K. M. Górski, and P. B. Lilje. On Foreground Removal from the Wilkinson Microwave Anisotropy Probe Data by an Internal Linear Combination Method: Limitations and Implications. *ApJ*, 612:633–646, Sept. 2004. doi: 10.1086/422807.

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. ApJ, 641:665–682, Apr. 2006. doi: 10.1086/500499.

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. ApJ, 676:10–32, Mar. 2008. doi: 10.1086/525277.

J. Errard, S. M. Feeney, H. V. Peiris, and A. H. Jaffe. Robust forecasts on fundamental physics from the foreground-obscured, gravitationally-lensed CMB polarization. *ArXiv e-prints*, Sept. 2015.

I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz. Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity. *Journal of High Energy Physics*, 1:87, Jan. 2017. doi: 10.1007/JHEP01(2017)087.

A. E. Evrard. Beyond N-body - 3D cosmological gas dynamics. MNRAS, 235:911–934, Dec. 1988. doi: 10.1093/mnras/235.3.911.

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. Phys. Rev. D, 84(4):043507, Aug. 2011. doi: 10.1103/PhysRevD.84.043507.

R. Fernández-Cobos, P. Vielva, R. B. Barreiro, and E. Martínez-González. Multiresolution internal template cleaning: an application to the Wilkinson Microwave Anisotropy Probe 7-yr polarization data. MNRAS, 420:2162–2169, Mar. 2012. doi: 10.1111/j.1365-2966.2011.20182.x.

R. Fernández-Cobos, A. Marcos-Caballero, P. Vielva, E. Martínez-González, and R. B. Barreiro. Exploring 2-spin internal linear combinations for the recovery of the CMB polarization. *ArXiv e-prints*, Jan. 2016.

A. Ferté, J. Grain, M. Tristram, and R. Stompor. Efficiency of pseudospectrum methods for estimation of the cosmic microwave background B-mode power spectrum. Phys. Rev. D, 88(2):023524, July 2013. doi: 10.1103/PhysRevD.88.023524.

A. Font-Ribera and J. Miralda-Escudé. The effect of high column density systems on the measurement of the Lyman-α forest correlation function. J. Cosmology Astropart. Phys., 7:028, July 2012. doi: 10.1088/1475-7516/2012/07/028.

A. Font-Ribera, P. McDonald, and J. Miralda-Escudé. Generating mock data sets for large-scale Lyman-α forest correlation measurements. J. Cosmology Astropart. Phys., 1:001, Jan. 2012a. doi: 10.1088/1475-7516/2012/01/001.

A. Font-Ribera, J. Miralda-Escudé, E. Arnau, B. Carithers, K.-G. Lee, P. Noterdaeme, I. Pâris, P. Petitjean, J. Rich, E. Rollinde, N. P. Ross, D. P. Schneider, M. White, and D. G. York. The large-scale cross-correlation of Damped Lyman alpha systems with the Lyman alpha forest: first measurements from BOSS. J. Cosmology Astropart. Phys., 11:059, Nov. 2012b. doi: 10.1088/1475-7516/2012/11/059.

A. Font-Ribera, E. Arnau, J. Miralda-Escudé, E. Rollinde, J. Brinkmann, J. R. Brownstein, K.-G. Lee, A. D. Myers, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, J. Rich, N. P. Ross, D. P. Schneider, and M. White. The large-scale quasar-Lyman α forest cross-correlation from BOSS. J. Cosmology Astropart. Phys., 5:018, May 2013. doi: 10.1088/1475-7516/2013/05/018.

A. Font-Ribera, D. Kirkby, N. Busca, J. Miralda-Escudé, N. P. Ross, A. Slosar, J. Rich, É. Aubourg, S. Bailey, V. Bhardwaj, J. Bautista, F. Beutler, D. Bizyaev, M. Blomqvist, H. Brewington, J. Brinkmann, J. R. Brownstein, B. Carithers, K. S. Dawson, T. Delubac, G. Ebelke, D. J. Eisenstein, J. Ge, K. Kinemuchi, K.-G. Lee, V. Malanushenko, E. Malanushenko, M. Marchante, D. Margala, D. Muna, A. D. Myers, P. Noterdaeme, D. Oravetz, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, M. M. Pieri, G. Rossi, D. P. Schneider, A. Simmons, M. Viel, C. Yeche, and D. G. York. Quasar-Lyman α forest cross-correlation from BOSS DR11: Baryon Acoustic Oscillations. J. Cosmology Astropart. Phys., 5:027, May 2014a. doi: 10.1088/1475-7516/2014/05/027.

A. Font-Ribera, P. McDonald, N. Mostek, B. A. Reid, H.-J. Seo, and A. Slosar. DESI and other Dark Energy experiments in the era of neutrino mass measurements. J. Cosmology Astropart. Phys., 5: 023, May 2014b. doi: 10.1088/1475-7516/2014/05/023.

A. Font-Ribera, P. McDonald, and A. Slosar. How to estimate the 3D power spectrum of the Lyman-α forest. *ArXiv e-prints*, Oct. 2017.

D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. PASP, 125:306, Mar. 2013. doi: 10.1086/670067.

D. V. Forero, M. Tórtola, and J. W. F. Valle. Global status of neutrino oscillation parameters after Neutrino-2012. Phys. Rev. D, 86(7):073012, Oct. 2012. doi: 10.1103/PhysRevD.86.073012.

D. V. Forero, M. Tórtola, and J. W. F. Valle. Neutrino oscillations refitted. Phys. Rev. D, 90(9): 093006, Nov. 2014. doi: 10.1103/PhysRevD.90.093006.

M. Fumagalli, J. X. Prochaska, D. Kasen, A. Dekel, D. Ceverino, and J. R. Primack. Absorption-line systems in simulated galaxies fed by cold streams. MNRAS, 418:1796–1821, Dec. 2011. doi: 10.1111/j.1365-2966.2011.19599.x.

M. Fumagalli, J. M. O'Meara, J. X. Prochaska, M. Rafelski, and N. Kanekar. Directly imaging damped Ly α galaxies at z ¿ 2 - III. The star formation rates of neutral gas reservoirs at z $\sim$ 2.7. MNRAS, 446:3178–3198, Jan. 2015. doi: 10.1093/mnras/stu2325.

S. Galli, K. Benabed, F. Bouchet, J.-F. Cardoso, F. Elsner, E. Hivon, A. Mangilli, S. Prunet, and B. Wandelt. CMB polarization can constrain cosmology better than CMB temperature. Phys. Rev. D, 90(6):063504, Sept. 2014. doi: 10.1103/PhysRevD.90.063504.

R. Garnett, S. Ho, S. Bird, and J. Schneider. Detecting damped Ly α absorbers with Gaussian processes. MNRAS, 472:1850–1865, Dec. 2017. doi: 10.1093/mnras/stx1958.

D. Geller and D. Marinucci. Spin Wavelets on the Sphere. *JFAA*, 16(6):840–884, Nov. 2010.

D. Geller and D. Marinucci. Mixed Needlets. *J. Math. Anal. Appl.*, 375(2):610–630, June 2011.

D. Geller, F. K. Hansen, D. Marinucci, G. Kerkyacharian, and D. Picard. Spin needlets for cosmic microwave background polarization data analysis. *PRD*, 78(12):123533–+, 2008. doi: 10.1103/PhysRevD.78.123533.

M. J. Geller and J. P. Huchra. Mapping the universe. *Science*, 246:897–903, Nov. 1989. doi: 10.1126/science.246.4932.897.

A. Gelman and D. B. Rubin. Inference from iterative simulation using multiple sequences. *Statist. Sci.*, 7(4):457–472, 11 1992. doi: 10.1214/ss/1177011136. URL `http://dx.doi.org/10.1214/ss/1177011136`.

R. A. Gingold and J. J. Monaghan. Smoothed particle hydrodynamics - Theory and application to non-spherical stars. MNRAS, 181:375–389, Nov. 1977. doi: 10.1093/mnras/181.3.375.

M. C. Gonzalez-Garcia, M. Maltoni, and T. Schwetz. Updated fit to three neutrino mixing: status of leptonic CP violation. *Journal of High Energy Physics*, 11:52, Nov. 2014. doi: 10.1007/JHEP11(2014)052.

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. ApJ, 622:759–771, Apr. 2005. doi: 10.1086/427976.

J. Grain, M. Tristram, and R. Stompor. CMB EB and TB cross-spectrum estimation via pseudospectrum techniques. Phys. Rev. D, 86(7):076005, Oct. 2012. doi: 10.1103/PhysRevD.86.076005.

J. E. Gunn and B. A. Peterson. On the Density of Neutral Hydrogen in Intergalactic Space. ApJ, 142: 1633–1641, Nov. 1965. doi: 10.1086/148444.

A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. Phys. Rev. D, 23:347–356, Jan. 1981. doi: 10.1103/PhysRevD.23.347.

A. H. Guth and E. J. Weinberg. Could the universe have recovered from a slow first-order phase transition? *Nuclear Physics B*, 212:321–364, Feb. 1983. doi: 10.1016/0550-3213(83)90307-3.

M. G. Haehnelt. Are the Lyman alpha forest "clouds" expanding pancakes? *ArXiv Astrophysics e-prints*, Dec. 1995.

S. Hanany, P. Ade, A. Balbi, J. Bock, J. Borrill, A. Boscaleri, P. de Bernardis, P. G. Ferreira, V. V. Hristov, A. H. Jaffe, A. E. Lange, A. T. Lee, P. D. Mauskopf, C. B. Netterfield, S. Oh, E. Pascale, B. Rabii, P. L. Richards, G. F. Smoot, R. Stompor, C. D. Winant, and J. H. P. Wu. MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of 10'-5$\deg$. ApJ, 545:L5–L9, Dec. 2000. doi: 10.1086/317322.

E. R. Harrison. Fluctuations at the Threshold of Classical Cosmology. Phys. Rev. D, 1:2726–2730, May 1970. doi: 10.1103/PhysRevD.1.2726.

C. G. T. Haslam, C. J. Salter, H. Stoffel, and W. E. Wilson. A 408 MHz all-sky continuum survey. II - The atlas of contour maps. A&AS, 47:1, Jan. 1982.

J. W. Henning, J. T. Sayre, C. L. Reichardt, P. A. R. Ade, A. J. Anderson, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. C. Chiang, H. Cho, R. Citron, C. Corbett Moran, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, W. Everett, J. Gallicchio, E. M. George, A. Gilbert, N. W. Halverson, N. Harrington, G. C. Hilton, G. P. Holder, W. L. Holzapfel, S. Hoover, Z. Hou, J. D. Hrubes, N. Huang, J. Hubmayr, K. D. Irwin, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, D. Li, A. Lowitz, A. Manzotti, J. J. McMahon, S. S. Meyer, L. Mocanu, J. Montgomery, A. Nadolski, T. Natoli, J. P. Nibarger, V. Novosad, S. Padin, C. Pryke, J. E. Ruhl, B. R. Saliwanchik, K. K. Schaffer, C. Sievers, G. Smecher, A. A. Stark, K. T. Story, C. Tucker, K. Vanderlinde, T. Veach, J. D. Vieira, G. Wang, N. Whitehorn, W. L. K. Wu, and V. Yefremenko. Measurements of the Temperature and E-Mode Polarization of the CMB from 500 Square Degrees of SPTpol Data. *ArXiv e-prints*, July 2017.

L. Hernquist and N. Katz. TREESPH - A unification of SPH with the hierarchical tree method. ApJS, 70:419–446, June 1989. doi: 10.1086/191344.

L. Hernquist, N. Katz, D. H. Weinberg, and J. Miralda-Escudé. The Lyman-Alpha Forest in the Cold Dark Matter Model. ApJ, 457:L51, Feb. 1996. doi: 10.1086/309899.

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. ApJ, 749:90, Apr. 2012. doi: 10.1088/0004-637X/749/1/90.

M. P. Hobson, A. W. Jones, A. N. Lasenby, and F. R. Bouchet. Foreground separation methods for satellite observations of the cosmic microwave background. MNRAS, 300:1–29, Oct. 1998. doi: 10.1046/j.1365-8711.1998.01777.x.

Y. Hoffman and E. Ribak. Constrained realizations of Gaussian fields - A simple algorithm. ApJ, 380:L5–L8, Oct. 1991. doi: 10.1086/186160.

F. Hoyle. A New Model for the Expanding Universe. MNRAS, 108:372, 1948. doi: 10.1093/mnras/108.5.372.

W. Hu. Weak lensing of the CMB: A harmonic approach. Phys. Rev. D, 62(4):043007, Aug. 2000. doi: 10.1103/PhysRevD.62.043007.

W. Hu. Mapping the Dark Matter through the Cosmic Microwave Background Damping Tail. ApJ, 557:L79–L83, Aug. 2001a. doi: 10.1086/323253.

W. Hu. Angular trispectrum of the cosmic microwave background. Phys. Rev. D, 64(8):083005, Oct. 2001b. doi: 10.1103/PhysRevD.64.083005.

W. Hu. Dark synergy: Gravitational lensing and the CMB. Phys. Rev. D, 65(2):023003, Jan. 2002. doi: 10.1103/PhysRevD.65.023003.

W. Hu and T. Okamoto. Mass Reconstruction with Cosmic Microwave Background Polarization. ApJ, 574:566–574, Aug. 2002. doi: 10.1086/341110.

L. Hui and N. Y. Gnedin. Equation of state of the photoionized intergalactic medium. MNRAS, 292: 27, Nov. 1997. doi: 10.1093/mnras/292.1.27.

L. Hui, A. Stebbins, and S. Burles. A Geometrical Test of the Cosmological Energy Contents Using the Lyα Forest. ApJ, 511:L5–L8, Jan. 1999. doi: 10.1086/311826.

J. Humlicek. An efficient method for evaluation of the complex probability function: The voigt function and its derivatives. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 21(4): 309 – 313, 1979. ISSN 0022-4073. doi: http://dx.doi.org/10.1016/0022-4073(79)90062-1. URL `http://www.sciencedirect.com/science/article/pii/0022407379900621`.

V. Iršič, M. Viel, T. A. M. Berg, V. D'Odorico, M. G. Haehnelt, S. Cristiani, G. Cupani, T.-S. Kim, S. López, S. Ellison, G. D. Becker, L. Christensen, K. D. Denney, G. Worseck, and J. S. Bolton. The Lyman-alpha forest power spectrum from the XQ-100 Legacy Survey. MNRAS, Dec. 2016. doi: 10.1093/mnras/stw3372.

V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, and G. D. Becker. First constraints on fuzzy dark matter from Lyman-α forest data and hydrodynamical simulations. *ArXiv e-prints*, Mar. 2017a.

V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, S. Cristiani, G. D. Becker, V. D'Odorico, G. Cupani, T.-S. Kim, T. A. M. Berg, S. López, S. Ellison, L. Christensen, K. D. Denny, and G. Worseck. New Constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman-α forest data. *ArXiv e-prints*, Feb. 2017b.

B. Jain, U. Seljak, and S. White. Ray-tracing Simulations of Weak Lensing by Large-Scale Structure. ApJ, 530:547–577, Feb. 2000. doi: 10.1086/308384.

G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric dark matter. Phys. Rep., 267: 195–373, Mar. 1996. doi: 10.1016/0370-1573(95)00058-5.

N. Kaiser. On the spatial correlations of Abell clusters. ApJ, 284:L9–L12, Sept. 1984. doi: 10.1086/184341.

N. Kaiser. Clustering in real space and in redshift space. MNRAS, 227:1–21, July 1987. doi: 10.1093/mnras/227.1.1.

M. Kamionkowski, A. Kosowsky, and A. Stebbins. A Probe of Primordial Gravity Waves and Vorticity. *Physical Review Letters*, 78:2058–2061, Mar. 1997. doi: 10.1103/PhysRevLett.78.2058.

M. Kaplinghat, L. Knox, and Y.-S. Song. Determining Neutrino Mass from the Cosmic Microwave Background Alone. *Physical Review Letters*, 91(24):241301, Dec. 2003. doi: 10.1103/PhysRevLett.91.241301.

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. ApJ, 743:28, Dec. 2011. doi: 10.1088/0004-637X/743/1/28.

R. Keisler, S. Hoover, N. Harrington, J. W. Henning, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. C. Chiang, H.-M. Cho, R. Citron, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, W. Everett, J. Gallicchio, J. Gao, E. M. George, A. Gilbert, N. W. Halverson, D. Hanson, G. C. Hilton, G. P. Holder, W. L. Holzapfel, Z. Hou, J. D. Hrubes, N. Huang, J. Hubmayr, K. D. Irwin, L. Knox, A. T. Lee, E. M. Leitch, D. Li, D. Luong-Van, D. P. Marrone, J. J. McMahon, J. Mehl, S. S. Meyer, L. Mocanu, T. Natoli, J. P. Nibarger, V. Novosad, S. Padin, C. Pryke, C. L. Reichardt, J. E. Ruhl, B. R. Saliwanchik, J. T. Sayre, K. K. Schaffer, E. Shirokoff, G. Smecher, A. A. Stark, K. T. Story, C. Tucker, K. Vanderlinde, J. D. Vieira, G. Wang, N. Whitehorn, V. Yefremenko, and O. Zahn. Measurements of Sub-degree B-mode Polarization in the Cosmic Microwave Background from 100 Square Degrees of SPTpol Data. ApJ, 807:151, July 2015. doi: 10.1088/0004-637X/807/2/151.

T.-S. Kim, M. Viel, M. G. Haehnelt, R. F. Carswell, and S. Cristiani. The power spectrum of the flux distribution in the Lyman α forest of a large sample of UVES QSO absorption spectra (LUQAS). MNRAS, 347:355–366, Jan. 2004. doi: 10.1111/j.1365-2966.2004.07221.x.

D. Kirkby, D. Margala, A. Slosar, S. Bailey, N. G. Busca, T. Delubac, J. Rich, J. E. Bautista, M. Blomqvist, J. R. Brownstein, B. Carithers, R. A. C. Croft, K. S. Dawson, A. Font-Ribera, J. Miralda-Escudé, A. D. Myers, R. C. Nichol, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, G. Rossi, D. J. Schlegel, D. P. Schneider, M. Viel, D. H. Weinberg, and C. Yèche. Fitting methods for baryon acoustic oscillations in the Lyman-α forest fluctuations in BOSS data release 9. J. Cosmology Astropart. Phys., 3:024, Mar. 2013. doi: 10.1088/1475-7516/2013/03/024.

A. Kogut, A. J. Banday, C. L. Bennett, K. M. Gorski, G. Hinshaw, and W. T. Reach. High-Latitude Galactic Emission in the COBE Differential Microwave Radiometer 2 Year Sky Maps. ApJ, 460:1, Mar. 1996. doi: 10.1086/176947.

A. Kogut, D. N. Spergel, C. Barnes, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, E. Wollack, and E. L. Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Temperature-Polarization Correlation. ApJS, 148: 161–173, Sept. 2003. doi: 10.1086/377219.

J. M. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzapfel. Detection of polarization in the cosmic microwave background using DASI. Nature, 420:772–787, Dec. 2002. doi: 10.1038/nature01269.

J.-K. Krogager, J. P. U. Fynbo, P. Møller, C. Ledoux, P. Noterdaeme, L. Christensen, B. Milvang-Jensen, and M. Sparre. On the sizes of z¿2 damped Lyα absorbing galaxies. MNRAS, 424:L1–L5, July 2012. doi: 10.1111/j.1745-3933.2012.01272.x.

A. Krolewski, K.-G. Lee, M. White, J. Hennawi, D. J. Schlegel, P. E. Nugent, Z. Lukić, C. W. Stark, O. Le Fèvre, B. C. Lemaux, C. Maier, M. Salvato, and L. Tasca. A Detection of z˜2.3 Cosmic Voids from 3D Lyman-α Forest Tomography in the COSMOS Field. *ArXiv e-prints*, Oct. 2017.

S. M. Leach, J.-F. Cardoso, C. Baccigalupi, R. B. Barreiro, M. Betoule, J. Bobin, A. Bonaldi, J. Delabrouille, G. de Zotti, C. Dickinson, H. K. Eriksen, J. González-Nuevo, F. K. Hansen, D. Herranz, M. Le Jeune, M. López-Caniego, E. Martínez-González, M. Massardi, J.-B. Melin, M.-A. Miville-Deschênes, G. Patanchon, S. Prunet, S. Ricciardi, E. Salerno, J. L. Sanz, J.-L. Starck, F. Stivoli,

V. Stolyarov, R. Stompor, and P. Vielva. Component separation methods for the PLANCK mission. A&A, 491:597–615, Nov. 2008. doi: 10.1051/0004-6361:200810116.

K.-G. Lee, S. Bailey, L. E. Bartsch, W. Carithers, K. S. Dawson, D. Kirkby, B. Lundgren, D. Margala, N. Palanque-Delabrouille, M. M. Pieri, D. J. Schlegel, D. H. Weinberg, C. Yèche, É. Aubourg, J. Bautista, D. Bizyaev, M. Blomqvist, A. S. Bolton, A. Borde, H. Brewington, N. G. Busca, R. A. C. Croft, T. Delubac, G. Ebelke, D. J. Eisenstein, A. Font-Ribera, J. Ge, J.-C. Hamilton, J. F. Hennawi, S. Ho, K. Honscheid, J.-M. Le Goff, E. Malanushenko, V. Malanushenko, J. Miralda-Escudé, A. D. Myers, P. Noterdaeme, D. Oravetz, K. Pan, I. Pâris, P. Petitjean, J. Rich, E. Rollinde, N. P. Ross, G. Rossi, D. P. Schneider, A. Simmons, S. Snedden, A. Slosar, D. N. Spergel, N. Suzuki, M. Viel, and B. A. Weaver. The BOSS Lyα Forest Sample from SDSS Data Release 9. AJ, 145:69, Mar. 2013. doi: 10.1088/0004-6256/145/3/69.

K.-G. Lee, J. F. Hennawi, M. White, R. A. C. Croft, and M. Ozbek. Observational Requirements for Lyα Forest Tomographic Mapping of Large-scale Structure at z ˜ 2. ApJ, 788:49, June 2014. doi: 10.1088/0004-637X/788/1/49.

B. Leistedt, J. D. McEwen, P. Vandergheynst, and Y. Wiaux. S2LET: A code to perform fast wavelet analysis on the sphere. *A&A*, 558:A128, Oct. 2013. doi: 10.1051/0004-6361/201220729.

B. Leistedt, J. D. McEwen, M. Büttner, H. V. Peiris, P. Vandergheynst, and Y. Wiaux. Analysing the polarisation of the CMB with spin scale-discretised wavelets. *ArXiv e-prints*, Jan. 2015.

B. Leistedt, J. D. McEwen, M. Büttner, and H. V. Peiris. Wavelet reconstruction of E and B modes for CMB polarisation and cosmic shear analyses. *ArXiv e-prints*, May 2016.

K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly Journal of Applied Mathmatics*, II(2):164–168, 1944.

A. Lewis and A. Challinor. Weak gravitational lensing of the cmb. *Phys. Rept.*, 429:1–65, 2006. doi: 10.1016/j.physrep.2006.03.002.

A. Lewis, A. Challinor, and N. Turok. Analysis of CMB polarization on an incomplete sky. Phys. Rev. D, 65(2):023505, Jan. 2002. doi: 10.1103/PhysRevD.65.023505.

A. D. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108:389–393, Feb. 1982. doi: 10.1016/0370-2693(82)91219-9.

T. Louis, E. Grace, M. Hasselfield, M. Lungu, L. Maurin, G. E. Addison, P. A. R. Ade, S. Aiola, R. Allison, M. Amiri, E. Angile, N. Battaglia, J. A. Beall, F. de Bernardis, J. R. Bond, J. Britton, E. Calabrese, H.-m. Cho, S. K. Choi, K. Coughlin, D. Crichton, K. Crowley, R. Datta, M. J. Devlin, S. R. Dicker, J. Dunkley, R. Dünner, S. Ferraro, A. E. Fox, P. Gallardo, M. Gralla, M. Halpern, S. Henderson, J. C. Hill, G. C. Hilton, M. Hilton, A. D. Hincks, R. Hlozek, S. P. P. Ho, Z. Huang, J. Hubmayr, K. M. Huffenberger, J. P. Hughes, L. Infante, K. Irwin, S. Muya Kasanda, J. Klein, B. Koopman, A. Kosowsky, D. Li, M. Madhavacheril, T. A. Marriage, J. McMahon, F. Menanteau, K. Moodley, C. Munson, S. Naess, F. Nati, L. Newburgh, J. Nibarger, M. D. Niemack, M. R. Nolta, C. Nuñez, L. A. Page, C. Pappas, B. Partridge, F. Rojas, E. Schaan, B. L. Schmitt, N. Sehgal, B. D. Sherwin, J. Sievers, S. Simon, D. N. Spergel, S. T. Staggs, E. R. Switzer, R. Thornton, H. Trac, J. Treu, C. Tucker, A. Van Engelen, J. T. Ward, and E. J. Wollack. The Atacama Cosmology Telescope: two-season ACTPol spectra and parameters. J. Cosmology Astropart. Phys., 6:031, June 2017. doi: 10.1088/1475-7516/2017/06/031.

L. B. Lucy. A numerical approach to the testing of the fission hypothesis. AJ, 82:1013–1024, Dec. 1977. doi: 10.1086/112164.

S. J. Maddox, G. Efstathiou, W. J. Sutherland, and J. Loveday. The APM galaxy survey. I - APM measurements and star-galaxy separation. MNRAS, 243:692–712, Apr. 1990.

J. Magueijo and R. Brandenberger. Cosmic defects and cosmology. *ArXiv Astrophysics e-prints*, Feb. 2000.

E. Majorana. Teoria simmetrica dell'elettrone e del positrone. *Il Nuovo Cimento*, 14:171–184, Apr. 1937. doi: 10.1007/BF02961314.

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. *MNRAS*, 383: 539–545, 2008. doi: 10.1111/j.1365-2966.2007.12550.x.

D. W. Marquardt. An algorithm for Least-Squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963. ISSN 03684245. doi: 10.2307/2098941. URL `http://dx.doi.org/10.2307/2098941`.

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.

E. Martínez-González, J. M. Diego, P. Vielva, and J. Silk. Cosmic microwave background power spectrum estimation and map reconstruction with the expectation-maximization algorithm. MNRAS, 345:1101–1109, Nov. 2003. doi: 10.1046/j.1365-2966.2003.06885.x.

J. C. Mather, E. S. Cheng, R. E. Eplee, Jr., R. B. Isaacman, S. S. Meyer, R. A. Shafer, R. Weiss, E. L. Wright, C. L. Bennett, N. W. Boggess, E. Dwek, S. Gulkis, M. G. Hauser, M. Janssen, T. Kelsall, P. M. Lubin, S. H. Moseley, Jr., T. L. Murdock, R. F. Silverberg, G. F. Smoot, and D. T. Wilkinson. A preliminary measurement of the cosmic microwave background spectrum by the Cosmic Background Explorer (COBE) satellite. ApJ, 354:L37–L40, May 1990. doi: 10.1086/185717.

J. C. Mather, E. S. Cheng, D. A. Cottingham, R. E. Eplee, Jr., D. J. Fixsen, T. Hewagama, R. B. Isaacman, K. A. Jensen, S. S. Meyer, P. D. Noerdlinger, S. M. Read, L. P. Rosen, R. A. Shafer, E. L. Wright, C. L. Bennett, N. W. Boggess, M. G. Hauser, T. Kelsall, S. H. Moseley, Jr., R. F. Silverberg, G. F. Smoot, R. Weiss, and D. T. Wilkinson. Measurement of the cosmic microwave background spectrum by the COBE FIRAS instrument. ApJ, 420:439–444, Jan. 1994. doi: 10.1086/173574.

P. McDonald. Toward a Measurement of the Cosmological Geometry at z ˜ 2: Predicting Lyα Forest Correlation in Three Dimensions and the Potential of Future Data Sets. ApJ, 585:34–51, Mar. 2003. doi: 10.1086/345945.

P. McDonald and J. Miralda-Escudé. Measuring the Cosmological Geometry from the Lyα Forest along Parallel Lines of Sight. ApJ, 518:24–31, June 1999. doi: 10.1086/307264.

P. McDonald, J. Miralda-Escudé, M. Rauch, W. L. W. Sargent, T. A. Barlow, R. Cen, and J. P. Ostriker. The Observed Probability Distribution Function, Power Spectrum, and Correlation Function of the Transmitted Flux in the Lyα Forest. ApJ, 543:1–23, Nov. 2000. doi: 10.1086/317079.

P. McDonald, U. Seljak, R. Cen, P. Bode, and J. P. Ostriker. Physical effects on the Lyα forest flux power spectrum: damping wings, ionizing radiation fluctuations and galactic winds. MNRAS, 360:1471–1482, July 2005a. doi: 10.1111/j.1365-2966.2005.09141.x.

P. McDonald, U. Seljak, R. Cen, D. Shih, D. H. Weinberg, S. Burles, D. P. Schneider, D. J. Schlegel, N. A. Bahcall, J. W. Briggs, J. Brinkmann, M. Fukugita, Ž. Ivezić, S. Kent, and D. E. Vanden Berk. The Linear Theory Power Spectrum from the Lyα Forest in the Sloan Digital Sky Survey. ApJ, 635:761–783, Dec. 2005b. doi: 10.1086/497563.

P. McDonald, U. Seljak, S. Burles, D. J. Schlegel, D. H. Weinberg, R. Cen, D. Shih, J. Schaye, D. P. Schneider, N. A. Bahcall, J. W. Briggs, J. Brinkmann, R. J. Brunner, M. Fukugita, J. E. Gunn, Ž. Ivezić, S. Kent, R. H. Lupton, and D. E. Vanden Berk. The Lyα Forest Power Spectrum from the Sloan Digital Sky Survey. ApJS, 163:80–109, Mar. 2006. doi: 10.1086/444361.

J. D. McEwen and A. M. M. Scaife. Simulating full-sky interferometric observations. *MNRAS*, 389 (3):1163–1178, 2008. doi: 10.1111/j.1365-2966.2008.13690.x.

J. D. McEwen and Y. Wiaux. A Novel Sampling Theorem on the Sphere. *IEEE Transactions on Signal Processing*, 59:5876–5887, Dec. 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*, 2006.

J. D. McEwen, P. Vielva, M. P. Hobson, E. Martínez-González, and A. N. Lasenby. Detection of the integrated Sachs-Wolfe effect and corresponding dark energy constraints made with directional spherical wavelets. MNRAS, 376:1211–1226, Apr. 2007. doi: 10.1111/j.1365-2966.2007.11505.x.

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. MNRAS, 384:1289–1300, Mar. 2008. doi: 10.1111/j.1365-2966.2007.12776.x.

J. D. McEwen, Y. Wiaux, and D. M. Eyers. Data compression on the sphere. *A&A*, 531:A98, 2011. doi: 10.1051/0004-6361/201015728.

J. D. McEwen, P. Vandergheynst, and Y. t. c. Wiaux. ation of directional scale-discretized wavelet transforms on the sphere. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8858 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 0, Sept. 2013. doi: 10.1117/12.2022889.

J. D. McEwen, M. Büttner, B. Leistedt, H. V. Peiris, P. Vandergheynst, and Y. Wiaux. On spin scale-discretised wavelets on the sphere for the analysis of CMB polarisation. In *IAU Symposium*, volume 306 of *IAU Symposium*, pages 64–67, May 2014. doi: 10.1017/S1743921314011107.

J. D. McEwen, M. Buttner, B. Leistedt, H. V. Peiris, and Y. Wiaux. A Novel Sampling Theorem on the Rotation Group. *IEEE Signal Processing Letters*, 22:2425–2429, Dec. 2015. doi: 10.1109/LSP.2015.2490676.

J. D. McEwen, C. Durastanti, and Y. Wiaux. Localisation of directional scale-discretised wavelets on the sphere. *ACHA, submitted*, 2015a.

J. D. McEwen, B. Leistedt, M. Büttner, H. V. Peiris, and Y. Wiaux. Directional spin wavelets on the sphere. *IEEE TSP, submitted*, 2015b.

M. McQuinn and M. White. On estimating Lyα forest correlations between multiple sightlines. MNRAS, 415:2257–2269, Aug. 2011. doi: 10.1111/j.1365-2966.2011.18855.x.

R. H. Miller and K. H. Prendergast. Stellar Dynamics in a Discrete Phase Space. ApJ, 151:699, Feb. 1968. doi: 10.1086/149469.

R. N. Mohapatra and G. Senjanović. Neutrino mass and spontaneous parity nonconservation. *Phys. Rev. Lett.*, 44:912–915, Apr 1980. doi: 10.1103/PhysRevLett.44.912. URL `https://link.aps.org/doi/10.1103/PhysRevLett.44.912`.

D. C. Morton. Erratum: "atomic data for resonance absorption lines. iii. wavelengths longward of the lyman limit for the elements hydrogen to gallium" (apjs, 149, 205 [2003]). *The Astrophysical Journal Supplement Series*, 151(2):403, 2004. URL `http://stacks.iop.org/0067-0049/151/i=2/a=403`.

Y. Moudden, J.-F. Cardoso, J.-L. Starck, and J. Delabrouille. Blind Component Separation in Wavelet Space: Application to CMB Analysis. *EURASIP Journal on Applied Signal Processing*, 2005:484606, Dec. 2005. doi: 10.1155/ASP.2005.2437.

V. F. Mukhanov and G. V. Chibisov. Quantum fluctuations and a nonsingular universe. *ZhETF Pisma Redaktsiiu*, 33:549–553, May 1981.

F. Narcowich, P. Petrushev, and J. Ward. Localized Tight Frames on Spheres. *SIAM J. Math. Anal.*, 38:574–594, Jan. 2006. doi: 10.1086/422807.

M. Neeleman, A. M. Wolfe, J. X. Prochaska, and M. Rafelski. The Fundamental Plane of Damped Lyα Systems. ApJ, 769:54, May 2013. doi: 10.1088/0004-637X/769/1/54.

M. Neeleman, N. Kanekar, J. X. Prochaska, M. Rafelski, C. L. Carilli, and A. M. Wolfe. [C II] 158-μm emission from the host galaxies of damped Lyman-alpha systems. *Science*, 355:1285–1288, Mar. 2017. doi: 10.1126/science.aal1737.

D. Nelson, A. Pillepich, S. Genel, M. Vogelsberger, V. Springel, P. Torrey, V. Rodriguez-Gomez, D. Sijacki, G. F. Snyder, B. Griffen, F. Marinacci, L. Blecha, L. Sales, D. Xu, and L. Hernquist. The illustris simulation: Public data release. *Astronomy and Computing*, 13:12–37, Nov. 2015. doi: 10.1016/j.ascom.2015.09.003.

C. B. Netterfield, M. J. Devlin, N. Jarosik, L. Page, and E. J. Wollack. A Measurement of the Angular Power Spectrum of the Anisotropy in the Cosmic Microwave Background. ApJ, 474:47–66, Jan. 1997. doi: 10.1086/303438.

P. Noterdaeme, P. Petitjean, C. Ledoux, and R. Srianand. Evolution of the cosmological mass density of neutral gas from Sloan Digital Sky Survey II - Data Release 7. A&A, 505:1087–1098, Oct. 2009. doi: 10.1051/0004-6361/200912768.

P. Noterdaeme, P. Petitjean, W. C. Carithers, I. Pâris, A. Font-Ribera, S. Bailey, E. Aubourg, D. Bizyaev, G. Ebelke, H. Finley, J. Ge, E. Malanushenko, V. Malanushenko, J. Miralda-Escudé, A. D. Myers, D. Oravetz, K. Pan, M. M. Pieri, N. P. Ross, D. P. Schneider, A. Simmons, and D. G. York. Column density distribution and cosmological mass density of neutral gas: Sloan Digital Sky Survey-III Data Release 9. A&A, 547:L1, Nov. 2012. doi: 10.1051/0004-6361/201220259.

J. H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. Bull. Astron. Inst. Netherlands, 6:249, Aug. 1932.

B. W. O'Shea, G. Bryan, J. Bordner, M. L. Norman, T. Abel, R. Harkness, and A. Kritsuk. Introducing Enzo, an AMR Cosmology Application. *ArXiv Astrophysics e-prints*, Mar. 2004.

N. Palanque-Delabrouille, C. Yèche, A. Borde, J.-M. Le Goff, G. Rossi, M. Viel, É. Aubourg, S. Bailey, J. Bautista, M. Blomqvist, A. Bolton, J. S. Bolton, N. G. Busca, B. Carithers, R. A. C. Croft, K. S. Dawson, T. Delubac, A. Font-Ribera, S. Ho, D. Kirkby, K.-G. Lee, D. Margala, J. Miralda-Escudé, D. Muna, A. D. Myers, P. Noterdaeme, I. Pâris, P. Petitjean, M. M. Pieri, J. Rich, E. Rollinde, N. P. Ross, D. J. Schlegel, D. P. Schneider, A. Slosar, and D. H. Weinberg. The one-dimensional Lyα forest power spectrum from BOSS. A&A, 559:A85, Nov. 2013. doi: 10.1051/0004-6361/201322130.

N. Palanque-Delabrouille, C. Yèche, J. Baur, C. Magneville, G. Rossi, J. Lesgourgues, A. Borde, E. Burtin, J.-M. LeGoff, J. Rich, M. Viel, and D. Weinberg. Neutrino masses and cosmology with Lyman-alpha forest power spectrum. J. Cosmology Astropart. Phys., 11:011, Nov. 2015. doi: 10.1088/1475-7516/2015/11/011.

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.

J. A. Peacock, S. Cole, P. Norberg, C. M. Baugh, J. Bland-Hawthorn, T. Bridges, R. D. Cannon, M. Colless, C. Collins, W. Couch, G. Dalton, K. Deeley, R. De Propris, S. P. Driver, G. Efstathiou, R. S. Ellis, C. S. Frenk, K. Glazebrook, C. Jackson, O. Lahav, I. Lewis, S. Lumsden, S. Maddox, W. J. Percival, B. A. Peterson, I. Price, W. Sutherland, and K. Taylor. A measurement of the cosmological mass density from clustering in the 2dF Galaxy Redshift Survey. Nature, 410:169–173, Mar. 2001.

R. D. Peccei and H. R. Quinn. Constraints imposed by CP conservation in the presence of pseudoparticles. Phys. Rev. D, 16:1791–1797, Sept. 1977a. doi: 10.1103/PhysRevD.16.1791.

R. D. Peccei and H. R. Quinn. CP conservation in the presence of pseudoparticles. *Physical Review Letters*, 38:1440–1443, June 1977b. doi: 10.1103/PhysRevLett.38.1440.

P. J. E. Peebles and J. T. Yu. Primeval Adiabatic Perturbation in an Expanding Universe. ApJ, 162: 815, Dec. 1970. doi: 10.1086/150713.

A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. ApJ, 142:419–421, July 1965. doi: 10.1086/148307.

I. Pérez-Ràfols, A. Font-Ribera, J. Miralda-Escudé, M. Blomqvist, S. Bird, N. Busca, H. du-Mas-des-Bourboux, L. Mas-Ribas, P. Noterdaeme, P. Petitjean, J. Rich, and D. P. Schneider. The SDSS-DR12 large-scale cross-correlation of Damped Lyman Alpha Systems with the Lyman Alpha Forest. *ArXiv e-prints*, Sept. 2017.

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. ApJ, 517:565–586, June 1999. doi: 10.1086/307221.

C. Pichon, J. L. Vergely, E. Rollinde, S. Colombi, and P. Petitjean. Inversion of the Lyman α forest: three-dimensional investigation of the intergalactic medium. MNRAS, 326:597–620, Sept. 2001. doi: 10.1046/j.1365-8711.2001.04595.x.

O. Pisanti, A. Cirillo, S. Esposito, F. Iocco, G. Mangano, G. Miele, and P. D. Serpico. PArthENoPE: Public algorithm evaluating the nucleosynthesis of primordial elements. *Computer Physics Communications*, 178:956–971, June 2008. doi: 10.1016/j.cpc.2008.02.015.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, M. Baker, A. Balbi, A. J. Banday, and et al. Planck early results. I. The Planck mission. A&A, 536: A1, Dec. 2011. doi: 10.1051/0004-6361/201116464.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. I. R. Alves, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, H. Aussel, and et al. Planck 2013 results. I. Overview of products and scientific results. A&A, 571:A1, Nov. 2014a. doi: 10.1051/0004-6361/201321529.

Planck Collaboration, P. A. R. Ade, N. Aghanim, F. Argüeso, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, and et al. Planck 2013 results. XXVIII. The Planck Catalogue of Compact Sources. A&A, 571:A28, Nov. 2014b. doi: 10.1051/0004-6361/201321524.

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. XII. Diffuse component separation. A&A, 571:A12, Nov. 2014c. doi: 10.1051/0004-6361/201321580.

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. XVI. Cosmological parameters. A&A, 571:A16, Nov. 2014d. doi: 10.1051/0004-6361/201321591.

Planck Collaboration, R. Adam, P. A. R. Ade, N. Aghanim, Y. Akrami, M. I. R. Alves, M. Arnaud, F. Arroja, J. Aumont, C. Baccigalupi, and et al. Planck 2015 results. I. Overview of products and scientific results. *ArXiv e-prints*, Feb. 2015a.

Planck Collaboration, R. Adam, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, and et al. Planck 2015 results. VIII. High Frequency Instrument data processing: Calibration and maps. *submitted to* A&A, Feb. 2015b.

Planck Collaboration, R. Adam, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, and et al. Planck 2015 results. IX. Diffuse component separation: CMB maps. *ArXiv e-prints*, Feb. 2015c.

Planck Collaboration, P. A. R. Ade, N. Aghanim, Y. Akrami, P. K. Aluri, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, and et al. Planck 2015 results. XVI. Isotropy and statistics of the CMB. *ArXiv e-prints*, June 2015d.

Planck Collaboration, P. A. R. Ade, N. Aghanim, D. Alina, M. I. R. Alves, G. Aniano, C. Armitage-Caplan, M. Arnaud, D. Arzoumanian, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, E. Battaner, K. Benabed, A. Benoit-Lévy, J.-P. Bernard, M. Bersanelli, P. Bielewicz, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, A. Bracco, C. Burigana, J.-F. Cardoso, A. Catalano, A. Chamballu, H. C. Chiang, P. R. Christensen, S. Colombi, L. P. L. Colombo, C. Combet, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, C. Dickinson, J. M. Diego, S. Donzelli, O. Doré, M. Douspis, X. Dupac, G. Efstathiou, T. A. Enßlin, H. K. Eriksen, E. Falgarone, L. Fanciullo, K. Ferrière, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, S. Galeotta, K. Ganga, T. Ghosh, M. Giard, Y. Giraud-Héraud, J. González-Nuevo, K. M. Górski, A. Gregorio, A. Gruppuso, V. Guillet, F. K. Hansen, D. L. Harrison, G. Helou, C. Hernández-Monteagudo, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, K. M. Huffenberger, A. H. Jaffe, T. R. Jaffe, W. C. Jones, M. Juvela, E. Keihänen, R. Keskitalo, T. S. Kisner, R. Kneissl, J. Knoche, M. Kunz, H. Kurki-Suonio, G. Lagache, J.-M. Lamarre, A. Lasenby, C. R. Lawrence, R. Leonardi, F. Levrier, M. Liguori, P. B. Lilje, M. Linden-Vørnle, M. López-Caniego, P. M. Lubin, J. F. Macías-Pérez, D. Maino, N. Mandolesi, M. Maris, D. J. Marshall, P. G. Martin, E. Martínez-González, S. Masi, S. Matarrese, P. Mazzotta, A. Melchiorri, L. Mendes, A. Mennella, M. Migliaccio, M.-A. Miville-Deschênes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, F. Noviello, D. Novikov, I. Novikov, C. A. Oxborrow, L. Pagano, F. Pajot, D. Paoletti, F. Pasian, V.-M. Pelkonen, O. Perdereau, L. Perotto, F. Perrotta, F. Piacentini, M. Piat, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, L. Popa, G. W. Pratt, S. Prunet, J.-L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, C. Renault, S. Ricciardi, T. Riller, I. Ristorcelli, G. Rocha, C. Rosset, G. Roudier, B. Rusholme, M. Sandri, D. Scott, J. D. Soler, L. D. Spencer, V. Stolyarov, R. Stompor, R. Sudiwala, D. Sutton, A.-S. Suur-Uski, J.-F. Sygnet, J. A. Tauber, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, M. Tucci, G. Umana, L. Valenziano, J. Valiviita, B. Van Tent, P. Vielva, F. Villa, L. A. Wade, B. D. Wandelt, and A. Zonca. Planck intermediate results. XX. Comparison of polarized thermal emission from Galactic dust with simulations of MHD turbulence. A&A, 576:A105, Apr. 2015e. doi: 10.1051/0004-6361/201424086.

Planck Collaboration, P. A. R. Ade, N. Aghanim, F. Argüeso, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, and et al. Planck 2015 results. XXVI. The Second Planck Catalogue of Compact Sources. *ArXiv e-prints*, July 2015f.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, and et al. Planck 2015 results. XVII. Constraints on primordial non-Gaussianity. *ArXiv e-prints*, Feb. 2015g.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and et al. Planck 2015 results. XII. Full Focal Plane simulations. *ArXiv e-prints*, Sept. 2015h.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, N. Bartolo, E. Battaner, K. Benabed, A. Benoît, A. Benoit-Lévy, J.-P. Bernard, M. Bersanelli, P. Bielewicz, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet,

M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J.-F. Cardoso, A. Catalano, A. Chamballu, R.-R. Chary, P. R. Christensen, S. Colombi, L. P. L. Colombo, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, C. Dickinson, J. M. Diego, H. Dole, S. Donzelli, O. Doré, M. Douspis, A. Ducout, X. Dupac, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, J. Fergusson, F. Finelli, O. Forni, M. Frailis, E. Franceschi, A. Frejsel, S. Galeotta, S. Galli, K. Ganga, M. Giard, Y. Giraud-Héraud, E. Gjerløw, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, F. K. Hansen, D. Hanson, D. L. Harrison, S. Henrot-Versillé, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, K. M. Huffenberger, G. Hurier, A. H. Jaffe, T. R. Jaffe, M. Juvela, E. Keihänen, R. Keskitalo, K. Kiiveri, T. S. Kisner, J. Knoche, M. Kunz, H. Kurki-Suonio, A. Lähteenmäki, J.-M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, J. P. Leahy, R. Leonardi, J. Lesgourgues, F. Levrier, M. Liguori, P. B. Lilje, M. Linden-Vørnle, V. Lindholm, M. López-Caniego, P. M. Lubin, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, P. G. Martin, E. Martínez-González, S. Masi, S. Matarrese, P. Mazzotta, P. McGehee, P. R. Meinhold, A. Melchiorri, L. Mendes, A. Mennella, M. Migliaccio, S. Mitra, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Nørgaard-Nielsen, D. Novikov, I. Novikov, F. Paci, L. Pagano, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, T. J. Pearson, O. Perdereau, L. Perotto, F. Perrotta, V. Pettorino, E. Pierpaoli, D. Pietrobon, E. Pointecouteau, G. Polenta, G. W. Pratt, G. Prézeau, S. Prunet, J.-L. Puget, J. P. Rachen, R. Rebolo, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, C. Rosset, M. Rossetti, G. Roudier, J. A. Rubiño-Martín, B. Rusholme, M. Sandri, D. Santos, M. Savelainen, D. Scott, M. D. Seiffert, E. P. S. Shellard, L. D. Spencer, V. Stolyarov, R. Stompor, D. Sutton, A.-S. Suur-Uski, J.-F. Sygnet, J. A. Tauber, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, M. Tucci, J. Tuovinen, L. Valenziano, J. Valiviita, B. Van Tent, T. Vassallo, P. Vielva, F. Villa, L. A. Wade, B. D. Wandelt, R. Watson, I. K. Wehus, D. Yvon, A. Zacchei, and A. Zonca. Planck 2015 results. VI. LFI mapmaking. *submitted to* A&A, Feb. 2015i.

Planck Collaboration, R. Adam, P. A. R. Ade, N. Aghanim, Y. Akrami, M. I. R. Alves, F. Argüeso, M. Arnaud, F. Arroja, M. Ashdown, and et al. Planck 2015 results. I. Overview of products and scientific results. A&A, 594:A1, Sept. 2016a. doi: 10.1051/0004-6361/201527101.

Planck Collaboration, R. Adam, P. A. R. Ade, N. Aghanim, M. I. R. Alves, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, and et al. Planck 2015 results. X. Diffuse component separation: Foreground maps. A&A, 594:A10, Sept. 2016b. doi: 10.1051/0004-6361/201525967.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, and et al. Planck 2015 results. XVII. Constraints on primordial non-Gaussianity. A&A, 594:A17, Sept. 2016c. doi: 10.1051/0004-6361/201525836.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, and et al. Planck 2015 results. XX. Constraints on inflation. A&A, 594:A20, Sept. 2016d. doi: 10.1051/0004-6361/201525898.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and et al. Planck 2015 results. XIII. Cosmological parameters. A&A, 594:A13, Sept. 2016e. doi: 10.1051/0004-6361/201525830.

Planck Collaboration, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo, and et al. Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters. A&A, 594:A11, Sept. 2016f. doi: 10.1051/0004-6361/201526926.

Planck Collaboration, N. Aghanim, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, R. Battye, K. Benabed, J.-P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J.-F. Cardoso, J. Carron, A. Challinor, H. C. Chiang, L. P. L. Colombo, C. Combet, B. Comis, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, R. J.

Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J.-M. Delouis, E. Di Valentino, C. Dickinson, J. M. Diego, O. Doré, M. Douspis, A. Ducout, X. Dupac, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, E. Falgarone, Y. Fantaye, F. Finelli, F. Forastieri, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gruppuso, J. E. Gudmundsson, F. K. Hansen, G. Helou, S. Henrot-Versillé, D. Herranz, E. Hivon, Z. Huang, S. Ilić, A. H. Jaffe, W. C. Jones, E. Keihänen, R. Keskitalo, T. S. Kisner, L. Knox, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, J.-M. Lamarre, M. Langer, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, J. P. Leahy, F. Levrier, M. Liguori, P. B. Lilje, M. López-Caniego, Y.-Z. Ma, J. F. Macías-Pérez, G. Maggio, A. Mangilli, M. Maris, P. G. Martin, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio, M.-A. Miville-Deschênes, D. Molinari, A. Moneti, L. Montier, G. Morgante, A. Moss, S. Mottet, P. Naselsky, P. Natoli, C. A. Oxborrow, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, L. Patrizii, O. Perdereau, L. Perotto, V. Pettorino, F. Piacentini, S. Plaszczynski, L. Polastri, G. Polenta, J.-L. Puget, J. P. Rachen, B. Racine, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, M. Rossetti, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, G. Sirri, R. Sunyaev, A.-S. Suur-Uski, J. A. Tauber, M. Tenti, L. Toffolatti, M. Tomasi, M. Tristram, T. Trombetti, J. Valiviita, F. Van Tent, L. Vibert, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, R. Watson, I. K. Wehus, M. White, A. Zacchei, and A. Zonca. Planck intermediate results. XLVI. Reduction of large-scale systematic effects in HFI polarization maps and estimation of the reionization optical depth. A&A, 596:A107, Dec. 2016g. doi: 10.1051/0004-6361/201628890.

A. Pontzen. Scale-dependent bias in the baryonic-acoustic-oscillation-scale intergalactic neutral hydrogen. Phys. Rev. D, 89(8):083010, Apr. 2014. doi: 10.1103/PhysRevD.89.083010.

A. Pontzen, F. Governato, M. Pettini, C. M. Booth, G. Stinson, J. Wadsley, A. Brooks, T. Quinn, and M. Haehnelt. Damped Lyman α systems in galaxy formation simulations. MNRAS, 390: 1349–1371, Nov. 2008. doi: 10.1111/j.1365-2966.2008.13782.x.

A. Pontzen, A. Deason, F. Governato, M. Pettini, J. Wadsley, T. Quinn, A. Brooks, J. Bellovary, and J. P. U. Fynbo. The nature of HIabsorbers in gamma-ray burst afterglows: clues from hydrodynamic simulations. MNRAS, 402:1523–1535, Mar. 2010. doi: 10.1111/j.1365-2966.2009.16017.x.

A. Pontzen, S. Bird, H. Peiris, and L. Verde. Constraints on Ionizing Photon Production from the Large-scale Lyα Forest. ApJ, 792:L34, Sept. 2014. doi: 10.1088/2041-8205/792/2/L34.

W. H. Press and P. Schechter. Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. ApJ, 187:425–438, Feb. 1974. doi: 10.1086/152650.

J. X. Prochaska, J. M. O'Meara, and G. Worseck. A Definitive Survey for Lyman Limit Systems at z ˜ 3.5 with the Sloan Digital Sky Survey. ApJ, 718:392–416, July 2010. doi: 10.1088/0004-637X/718/1/392.

M. J. Rees and D. W. Sciama. Large-scale Density Inhomogeneities in the Universe. Nature, 217: 511–516, Feb. 1968. doi: 10.1038/217511a0.

R. D. Richtmyer and K. W. Morton. *Difference methods for initial-value problems*. 1967.

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. AJ, 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. ApJ, 607:665–687, June 2004. doi: 10.1086/383612.

K. K. Rogers, H. V. Peiris, B. Leistedt, J. D. McEwen, and A. Pontzen. SILC: a new Planck Internal Linear Combination CMB temperature map using directional wavelets. MNRAS, 460:3014–3028, Jan. 2016. doi: 10.1093/mnras/stw1121.

K. K. Rogers, S. Bird, H. V. Peiris, A. Pontzen, A. Font-Ribera, and B. Leistedt. Simulating the effect of high column density absorbers on the one-dimensional Lyman-alpha forest flux power spectrum. *ArXiv e-prints*, June 2017.

V. C. Rubin and W. K. Ford, Jr. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. ApJ, 159:379, Feb. 1970. doi: 10.1086/150317.

J. L. Sanz, D. Herranz, M. López-Caniego, and F. Argüeso. Wavelets on the sphere – application to the detection problem. In *EUSIPCO*, Sept. 2006.

E. Sefusatti and R. Scoccimarro. Galaxy bias and halo-occupation numbers from large-scale clustering. Phys. Rev. D, 71(6):063001, Mar. 2005. doi: 10.1103/PhysRevD.71.063001.

M. Seldner, B. Siebers, E. J. Groth, and P. J. E. Peebles. New reduction of the Lick catalog of galaxies. AJ, 82:249–256, Apr. 1977. doi: 10.1086/112039.

U. Seljak and M. Zaldarriaga. Signature of Gravity Waves in the Polarization of the Microwave Background. *Physical Review Letters*, 78:2054–2057, Mar. 1997. doi: 10.1103/PhysRevLett.78.2054.

U. Seljak and M. Zaldarriaga. Direct signature of an evolving gravitational potential from the cosmic microwave background. Phys. Rev. D, 60(4):043504, Aug. 1999. doi: 10.1103/PhysRevD.60.043504.

U. Seljak, A. Makarov, P. McDonald, S. F. Anderson, N. A. Bahcall, J. Brinkmann, S. Burles, R. Cen, M. Doi, J. E. Gunn, Ž. Ivezić, S. Kent, J. Loveday, R. H. Lupton, J. A. Munn, R. C. Nichol, J. P. Ostriker, D. J. Schlegel, D. P. Schneider, M. Tegmark, D. E. Berk, D. H. Weinberg, and D. G. York. Cosmological parameter analysis including SDSS Lyα forest and galaxy bias: Constraints on the primordial spectrum of fluctuations, neutrino mass, and dark energy. Phys. Rev. D, 71(10):103515, May 2005. doi: 10.1103/PhysRevD.71.103515.

D. Sijacki, V. Springel, T. Di Matteo, and L. Hernquist. A unified model for AGN feedback in cosmological simulations of structure formation. MNRAS, 380:877–900, Sept. 2007. doi: 10.1111/j.1365-2966.2007.12153.x.

J. Silk. Cosmic Black-Body Radiation and Galaxy Formation. ApJ, 151:459, Feb. 1968. doi: 10.1086/149449.

A. Slosar, A. Font-Ribera, M. M. Pieri, J. Rich, J.-M. Le Goff, É. Aubourg, J. Brinkmann, N. Busca, B. Carithers, R. Charlassier, M. Cortês, R. Croft, K. S. Dawson, D. Eisenstein, J.-C. Hamilton, S. Ho, K.-G. Lee, R. Lupton, P. McDonald, B. Medolin, D. Muna, J. Miralda-Escudé, A. D. Myers, R. C. Nichol, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, Y. Piškur, E. Rollinde, N. P. Ross, D. J. Schlegel, D. P. Schneider, E. Sheldon, B. A. Weaver, D. H. Weinberg, C. Yeche, and D. G. York. The Lyman-α forest in three dimensions: measurements of large scale flux correlations from BOSS 1st-year data. J. Cosmology Astropart. Phys., 9:001, Sept. 2011. doi: 10.1088/1475-7516/2011/09/001.

A. Slosar, V. Iršič, D. Kirkby, S. Bailey, N. G. Busca, T. Delubac, J. Rich, É. Aubourg, J. E. Bautista, V. Bhardwaj, M. Blomqvist, A. S. Bolton, J. Bovy, J. Brownstein, B. Carithers, R. A. C. Croft, K. S. Dawson, A. Font-Ribera, J.-M. Le Goff, S. Ho, K. Honscheid, K.-G. Lee, D. Margala, P. McDonald, B. Medolin, J. Miralda-Escudé, A. D. Myers, R. C. Nichol, P. Noterdaeme, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, M. M. Pieri, Y. Piškur, N. A. Roe, N. P. Ross, G. Rossi, D. J. Schlegel, D. P. Schneider, N. Suzuki, E. S. Sheldon, U. Seljak, M. Viel, D. H. Weinberg, and C. Yèche. Measurement of baryon acoustic oscillations in the Lyman-α forest fluctuations in BOSS data release 9. J. Cosmology Astropart. Phys., 4:026, Apr. 2013. doi: 10.1088/1475-7516/2013/04/026.

K. M. Smith and M. Zaldarriaga. General solution to the E-B mixing problem. Phys. Rev. D, 76(4): 043001, Aug. 2007. doi: 10.1103/PhysRevD.76.043001.

K. M. Smith, O. Zahn, and O. Doré. Detection of gravitational lensing in the cosmic microwave background. Phys. Rev. D, 76(4):043510, Aug. 2007. doi: 10.1103/PhysRevD.76.043510.

R. E. Smith, J. A. Peacock, A. Jenkins, S. D. M. White, C. S. Frenk, F. R. Pearce, P. A. Thomas, G. Efstathiou, and H. M. P. Couchman. Stable clustering, the halo model and non-linear cosmological power spectra. MNRAS, 341:1311–1332, June 2003. doi: 10.1046/j.1365-8711.2003.06503.x.

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. ApJ, 396:L1–L5, Sept. 1992. doi: 10.1086/186504.

V. Springel. The cosmological simulation code GADGET-2. MNRAS, 364:1105–1134, Dec. 2005. doi: 10.1111/j.1365-2966.2005.09655.x.

V. Springel. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. MNRAS, 401:791–851, Jan. 2010. doi: 10.1111/j.1365-2966.2009.15715.x.

V. Springel and L. Hernquist. Cosmological smoothed particle hydrodynamics simulations: the entropy equation. MNRAS, 333:649–664, July 2002. doi: 10.1046/j.1365-8711.2002.05445.x.

V. Springel and L. Hernquist. Cosmological smoothed particle hydrodynamics simulations: a hybrid multiphase model for star formation. MNRAS, 339:289–311, Feb. 2003. doi: 10.1046/j.1365-8711.2003.06206.x.

V. Springel, T. Di Matteo, and L. Hernquist. Modelling feedback from stars and black holes in galaxy mergers. MNRAS, 361:776–794, Aug. 2005. doi: 10.1111/j.1365-2966.2005.09238.x.

J. Starck, Y. Moudden, and J. Bobin. Polarized wavelets and curvelets on the sphere. *A&A*, 497: 931–943, Apr. 2009.

J.-L. Starck, Y. Moudden, P. Abrial, and M. Nguyen. Wavelets, ridgelets and curvelets on the sphere. *A&A*, 446:1191–1204, Feb. 2006. doi: 10.1051/0004-6361:20053246.

A. A. Starobinsky. Dynamics of phase transition in the new inflationary universe scenario and generation of perturbations. *Physics Letters B*, 117:175–178, Nov. 1982. doi: 10.1016/0370-2693(82)90541-X.

R. A. Sunyaev and Y. B. Zeldovich. Small-Scale Fluctuations of Relic Radiation. Ap&SS, 7:3–19, Apr. 1970. doi: 10.1007/BF00653471.

A. N. Taylor, S. Dye, T. J. Broadhurst, N. Benítez, and E. van Kampen. Gravitational Lens Magnification and the Mass of Abell 1689. ApJ, 501:539–553, July 1998. doi: 10.1086/305827.

M. Tegmark, A. de Oliveira-Costa, and A. J. Hamilton. High resolution foreground cleaned CMB map from WMAP. *Phys. Rev. D*, 68(12):123523, Dec. 2003. doi: 10.1103/PhysRevD.68.123523.

R. Teyssier. Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. A&A, 385:337–364, Apr. 2002. doi: 10.1051/0004-6361:20011817.

The Polarbear Collaboration: P. A. R. Ade, Y. Akiba, A. E. Anthony, K. Arnold, M. Atlas, D. Barron, D. Boettger, J. Borrill, S. Chapman, Y. Chinone, M. Dobbs, T. Elleflot, J. Errard, G. Fabbian, C. Feng, D. Flanigan, A. Gilbert, W. Grainger, N. W. Halverson, M. Hasegawa, K. Hattori, M. Hazumi, W. L. Holzapfel, Y. Hori, J. Howard, P. Hyland, Y. Inoue, G. C. Jaehnig, A. H. Jaffe, B. Keating, Z. Kermish, R. Keskitalo, T. Kisner, M. Le Jeune, A. T. Lee, E. M. Leitch, E. Linder, M. Lungu,

F. Matsuda, T. Matsumura, X. Meng, N. J. Miller, H. Morii, S. Moyerman, M. J. Myers, M. Navaroli, H. Nishino, A. Orlando, H. Paar, J. Peloton, D. Poletti, E. Quealy, G. Rebeiz, C. L. Reichardt, P. L. Richards, C. Ross, I. Schanning, D. E. Schenck, B. D. Sherwin, A. Shimizu, C. Shimmin, M. Shimon, P. Siritanasak, G. Smecher, H. Spieler, N. Stebor, B. Steinbach, R. Stompor, A. Suzuki, S. Takakura, T. Tomaru, B. Wilson, A. Yadav, and O. Zahn. A Measurement of the Cosmic Microwave Background B-mode Polarization Power Spectrum at Sub-degree Scales with POLARBEAR. ApJ, 794:171, Oct. 2014. doi: 10.1088/0004-637X/794/2/171.

M. Viel, M. G. Haehnelt, R. F. Carswell, and T.-S. Kim. The effect of (strong) discrete absorption systems on the Lyman α forest flux power spectrum. MNRAS, 349:L33–L37, Apr. 2004a. doi: 10.1111/j.1365-2966.2004.07753.x.

M. Viel, M. G. Haehnelt, and V. Springel. Inferring the dark matter power spectrum from the Lyman α forest in high-resolution QSO absorption spectra. MNRAS, 354:684–694, Nov. 2004b. doi: 10.1111/j.1365-2966.2004.08224.x.

M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt. Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman-α forest data. Phys. Rev. D, 88(4):043502, Aug. 2013. doi: 10.1103/PhysRevD.88.043502.

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. ApJ, 609: 22–34, July 2004. doi: 10.1086/421007.

M. Vogelsberger, S. Genel, D. Sijacki, P. Torrey, V. Springel, and L. Hernquist. A model for cosmological simulations of galaxy formation physics. MNRAS, 436:3031–3067, Dec. 2013. doi: 10.1093/mnras/stt1789.

M. Vogelsberger, S. Genel, V. Springel, P. Torrey, D. Sijacki, D. Xu, G. Snyder, S. Bird, D. Nelson, and L. Hernquist. Properties of galaxies reproduced by a hydrodynamic simulation. Nature, 509: 177–182, May 2014a. doi: 10.1038/nature13316.

M. Vogelsberger, S. Genel, V. Springel, P. Torrey, D. Sijacki, D. Xu, G. Snyder, D. Nelson, and L. Hernquist. Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe. MNRAS, 444:1518–1547, Oct. 2014b. doi: 10.1093/mnras/stu1536.

S. von Hoerner. Die numerische Integration des n-Körper-Problemes für Sternhaufen. I. ZAp, 50, 1960.

D. H. Weinberg, J. S. Bullock, F. Governato, R. Kuzio de Naray, and A. H. G. Peter. Cold dark matter: Controversies on small scales. *Proceedings of the National Academy of Science*, 112:12249–12255, Oct. 2015. doi: 10.1073/pnas.1308716112.

S. D. M. White and M. J. Rees. Core condensation in heavy halos - A two-stage theory for galaxy formation and clustering. MNRAS, 183:341–358, May 1978. doi: 10.1093/mnras/183.3.341.

Y. Wiaux, L. Jacques, and P. Vandergheynst. Correspondence principle between spherical and Euclidean wavelets. *ApJ*, 632:15–28, 2005.

Y. Wiaux, J. D. McEwen, P. Vandergheynst, and O. Blanc. Exact reconstruction with directional wavelets on the sphere. *MNRAS*, 388:770–788, Aug. 2008. doi: 10.1111/j.1365-2966.2008.13448.x.

A. M. Wolfe, D. A. Turnshek, H. E. Smith, and R. D. Cohen. Damped Lyman-alpha absorption by disk galaxies with large redshifts. I - The Lick survey. ApJS, 61:249–304, June 1986. doi: 10.1086/191114.

C. Yeche, N. Palanque-Delabrouille, J. . Baur, and H. du Mas des BourBoux. Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100. *ArXiv e-prints*, Feb. 2017.

T. Zafar, C. Péroux, A. Popping, B. Milliard, J.-M. Deharveng, and S. Frank. The ESO UVES advanced data products quasar sample. II. Cosmological evolution of the neutral gas mass density. A&A, 556:A141, Aug. 2013. doi: 10.1051/0004-6361/201321154.

M. Zaldarriaga and U. Seljak. Gravitational lensing effect on cosmic microwave background polarization. Phys. Rev. D, 58(2):023003, July 1998. doi: 10.1103/PhysRevD.58.023003.

M. Zaldarriaga and U. Seljak. Reconstructing projected matter density power spectrum from cosmic microwave background. Phys. Rev. D, 59(12):123507, June 1999. doi: 10.1103/PhysRevD.59.123507.

Y. B. Zeldovich. A hypothesis, unifying the structure and the entropy of the Universe. MNRAS, 160: 1P, 1972. doi: 10.1093/mnras/160.1.1P.

F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.