

2008

Combined analysis of the integrated Sachs-Wolfe effect and cosmological implications

Tommaso Giannantonio

Ryan Scranton

Robert G. Crittenden

Stephen P. Boughn

Haverford College, sboughn@haverford.edu

Follow this and additional works at: http://scholarship.haverford.edu/astronomy_facpubs

Repository Citation

"Combined Analysis of the Integrated Sachs–Wolfe Effect and Cosmological Implications" (with T. Giannantonio, R. Scranton, R. Crittenden, R. Nichol, A. Myers, G. Richards), *Phys. Rev. D*, 77, 123520 (2008).

This Journal Article is brought to you for free and open access by the Astronomy at Haverford Scholarship. It has been accepted for inclusion in Faculty Publications by an authorized administrator of Haverford Scholarship. For more information, please contact nmedeiro@haverford.edu.

Combined analysis of the integrated Sachs-Wolfe effect and cosmological implications

Tommaso Giannantonio,^{1,*} Ryan Scranton,² Robert G. Crittenden,¹ Robert C. Nichol,¹ Stephen P. Boughn,³
Adam D. Myers,⁴ and Gordon T. Richards⁵

¹*Institute of Cosmology and Gravitation, University of Portsmouth, Mercantile House, Hampshire Terrace, Portsmouth PO1 2EG, UK*

²*Department of Physics and Astronomy, University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, Pennsylvania 15260, USA*

³*Haverford College, Haverford Pennsylvania 19041, USA*

⁴*Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, Illinois 61820, USA*

⁵*Department of Physics, Drexel University, 3141 Chestnut Street, Philadelphia, Pennsylvania 19104, USA*

(Received 5 February 2008; published 13 June 2008)

We present a global measurement of the integrated Sachs-Wolfe (ISW) effect obtained by cross correlating all relevant large-scale galaxy data sets with the cosmic microwave background radiation map provided by the Wilkinson Microwave Anisotropy Probe. With these measurements, the overall ISW signal is detected at the $\sim 4.5\sigma$ level. We also examine the cosmological implications of these measurements, particularly the dark energy equation of state w , its sound speed c_s , and the overall curvature of the Universe. The flat Λ CDM model is a good fit to the data and, assuming this model, we find that the ISW data constrain $\Omega_m = 0.20_{-0.11}^{+0.19}$ at the 95% confidence level. When we combine our ISW results with the latest baryon oscillation and supernovae measurements, we find that the result is still consistent with a flat Λ CDM model with $w = -1$ out to redshifts $z > 1$.

DOI: [10.1103/PhysRevD.77.123520](https://doi.org/10.1103/PhysRevD.77.123520)

PACS numbers: 98.80.Es, 98.54.Aj, 98.70.Vc

I. INTRODUCTION

There is growing evidence that the expansion of the Universe is dominated by an unknown *dark energy*, which accounts for more than 70% of the total matter that we observe. Measurements of the Hubble diagram of Type Ia supernovae [1,2] indicate this dark energy is causing the Universe's expansion to accelerate, while cosmic microwave background (CMB) anisotropies power spectrum experiments, such as WMAP [3], limit the amount of spatial curvature, which might otherwise explain the low dark matter density. Understanding the nature of this dark energy is essential, and many possible explanations have been proposed, from a cosmological constant, to quintessence, to modifying the laws of gravity.

To discriminate between possible explanations of dark energy, many alternative probes have been developed, including measuring the Universe's geometry through observations of the baryon oscillation (BAO) scale [4–6] and probing the dark matter density through observations of clusters of galaxies [7,8]. Here, we consider another impact of dark energy, the creation of CMB anisotropies at late times via the integrated Sachs-Wolfe (ISW) effect [9]. While most CMB anisotropies are generated at very early times, further fluctuations can be induced gravitationally at late times as photons pass through evolving gravitational potentials. If dark matter dominates, the gravitational potentials do not vary with time, but the presence of dark energy or spatial curvature will cause the potentials to evolve at late times, producing new temperature fluctuations at low redshifts (primarily at $z < 2$).

Directly observing these new CMB temperature anisotropies is challenging, primarily because their amplitudes are a fraction of the anisotropies arising from higher redshifts. They are also predominantly seen on large scales, where the uncertainty in the observations from cosmic variance is biggest. The search for the ISW effect has instead focused on finding correlations between the CMB temperature maps and maps of the density, which trace the local gravitational potentials [10]. Primordial anisotropies should be uncorrelated with the local density, making it possible to pull out the weaker ISW anisotropies.

Many groups have detected this correlation using the accurate WMAP CMB data and various density probes distributed at a range of redshifts and in different regions of the electromagnetic spectrum: from shallow infrared observations [11,12], to optical surveys such as the Automatic Plate Measuring Machine Survey (APM) and the Sloan Digital Sky Survey (SDSS) [13–17], radio galaxy catalogues [18,19], x-ray surveys [18], and the deepest quasars from the SDSS [20]. These measures span a range of redshift going from $z = 0.1$ to $z = 1.5$, where the ISW effect has been measured at significance levels, typically around $2\text{--}3\sigma$, and appear generally compatible with the expectation from the Λ CDM model.

Although indicative of the presence of dark energy, none of these measures alone has significant power to constrain models due to their low significance. Thus, it is important to combine the various observations; but some care must be taken in doing so. The surveys are often overlapping both in sky coverage and in redshift range, meaning there are likely covariances between them that may be important when considering a large-scale effect like the ISW. In addition, these measurements have been made with a vari-

*tommaso.giannantonio@port.ac.uk

ety of techniques, using angular correlations, Fourier modes, or a range of wavelet techniques [21–24]. The error bars themselves have also been estimated using different techniques, using both jackknife approaches and Monte Carlo simulations of the CMB sky.

A combined analysis has been attempted in the past adding several measures in order to extend the constraining power in redshift and learn more about the behavior of dark energy and other cosmological parameters [25–28]. However, this analysis largely ignored the differences in the observations and accounted for the covariances between experiments in a fairly arbitrary way. Here, we perform a combined analysis by reanalyzing all the observations in a consistent way, measuring directly the covariances between the different observations using a number of different methods and looking at the cross correlations between all the various data sets. In this way, we hope to give a definitive result for the ISW evidence for dark energy and the resulting cosmological constraints.

This paper is structured as follows: we begin in Sec. II by giving a brief theoretical description of the ISW effect, how the cross correlation is measured, and the important issue of estimating the covariance between observations. In Sec. III, we describe the catalogues used for the cross correlation, and in Sec. IV, we show the measurements of the various cross-correlation functions between the different catalogues and their cross correlation with the CMB. We discuss the significance of the measurements in Sec. V and show the resulting cosmological constraints in Sec. VI, before some concluding remarks in Sec. VII.

II. METHOD

A. The ISW effect

We begin with some theoretical background on the ISW effect and its detection. Most of the CMB anisotropies we observe were created at the redshift of last scattering ($z \sim 1100$) when the Universe was 400 ky old, as a result of fluctuations in the photon density, velocity, and potential energy. Since that time, the CMB photons have travelled largely untouched, but anisotropies can still be produced gravitationally if the photons pass through time varying potential wells

$$\frac{\Delta T}{T}(\hat{\mathbf{n}}) = -2 \int \dot{\Phi}[\tau, \hat{\mathbf{n}}(\tau_0 - \tau)] d\tau, \quad (1)$$

where τ is the conformal time, the dot represents a conformal-time derivative, and the integral is intended along the line of sight of the photon. (Throughout, we work in units where the speed of light is unity and for simplicity assume that the effects of anisotropic stress can be ignored.) When the gravitational potential decays due to its linear evolution, this is usually referred to as the ISW effect, and if the potential decay is a result of nonlinear evolution, as in clusters, it is referred to as the Rees-Sciama effect [29].

The physical picture is very straightforward; as a CMB photon falls into a gravitational potential well, it gains energy; as the photon climbs out of a potential well, it loses energy. These effects exactly cancel if the potential is time independent, but can result in a net kick if the potential evolves as the photon passes through it. In particular, we know that during the matter dominated era the gravitational potential stays constant, and so $\dot{\Phi} = 0$, which means that in that era there will not be any ISW produced. This changes if dark energy or curvature become important at later times: in this case, $\dot{\Phi} \neq 0$ and additional CMB anisotropies will be produced. In the usual case, the potential amplitudes decrease at late times, so that a temperature increase results from passing through potential wells, while a temperature deficits results from traversing potential hills.

Since we know from the CMB experiments, such as WMAP, that the Universe is very close to flat, we can attribute most of the late ISW to dark energy; observing it can therefore provide a probe of dark energy, its properties, and its evolution in time. In particular, the ISW effect could be one of the few ways of measuring the sound speed of dark energy [30]. However, it can also provide useful information about the curvature of the Universe [31].

Unfortunately, measurement of the ISW effect is made difficult by the embedding of the small ISW signal in the much larger (10 times) primary CMB anisotropies. Furthermore, the total ISW signal is due to all the density fluctuations, both positive and negative, along the line of sight. On small scales, the individual temperature differences are small, and they tend to cancel out. The most significant ISW effect results from the coherent large-scale potentials, but unfortunately these scales are precisely where cosmic variance is most troublesome.

This problem can be overcome by examining how the ISW temperature correlates with the density of galaxies, which should trace the potential wells and hills that bring about the anisotropies. The observed galaxy density contrast in a given direction $\hat{\mathbf{n}}_1$ will be

$$\delta_g(\hat{\mathbf{n}}_1) = \int b_g(z) \frac{dN}{dz}(z) \delta_m(\hat{\mathbf{n}}_1, z) dz, \quad (2)$$

where dN/dz is the selection function of the survey, $b_g(z)$ its galaxy bias relating the visible matter distribution to the underlying dark matter, and δ_m the matter density perturbations. Since the density δ_m is related to the potential Φ by the Poisson equation, the observed galaxy density will be correlated with the ISW temperature fluctuation in the nearby direction $\hat{\mathbf{n}}_2$, which is

$$\frac{\Delta T}{T}(\hat{\mathbf{n}}_2) = -2 \int e^{-\tau(z)} \frac{d\Phi}{dz}(\hat{\mathbf{n}}_2, z) dz, \quad (3)$$

where $e^{-\tau(z)}$ is the visibility function of the photons, which accounts for the effect of photons rescattering following reionization.

The galaxy bias $b_g(z)$ can evolve in time or as a function of scale; however, we will generally assume that it is time and scale independent for simplicity. For our purposes, a time-dependent bias is equivalent to changing the selection function of the survey. Scale dependence of the bias is more problematic, but on the very large scales (> 10 Mpc) we are considering, the scale dependence is expected to be weak [32,33].

Given a map of the CMB and a survey of galaxies, the angular autocorrelation and cross-correlation functions are defined as

$$C^{Tg}(\vartheta) \equiv \left\langle \frac{\Delta T}{T}(\hat{\mathbf{n}}_1) \delta_g(\hat{\mathbf{n}}_2) \right\rangle, \quad (4)$$

$$C^{gg}(\vartheta) \equiv \langle \delta_g(\hat{\mathbf{n}}_1) \delta_g(\hat{\mathbf{n}}_2) \rangle, \quad (5)$$

with the average carried over all the pairs at the same angular distance $\vartheta = |\hat{\mathbf{n}}_1 - \hat{\mathbf{n}}_2|$.

It is possible to express these quantities in the harmonic space with the use of the Legendre polynomials P_l

$$C^{Tg}(\vartheta) = \sum_{l=2}^{\infty} \frac{2l+1}{4\pi} C_l^{Tg} P_l[\cos(\vartheta)], \quad (6)$$

and the auto- and cross-correlation power spectra are given by

$$C_l^{Tg} = 4\pi \int \frac{dk}{k} \Delta^2(k) I_l^{ISW}(k) I_l^g(k), \quad (7)$$

$$C_l^{gg} = 4\pi \int \frac{dk}{k} \Delta^2(k) I_l^g(k) I_l^g(k), \quad (8)$$

where $\Delta(k)$ is the scale invariant matter power spectrum $\Delta^2(k) \equiv 4\pi k^3 P(k)/(2\pi)^3$, and the two integrands are respectively

$$I_l^{ISW}(k) = -2 \int e^{-\tau(z)} \frac{d\Phi_k}{dz} j_l[k\chi(z)] dz, \quad (9)$$

$$I_l^g(k) = \int b_g(z) \frac{dN}{dz}(z) \delta_m(k, z) j_l[k\chi(z)] dz, \quad (10)$$

where Φ_k , $\delta_m(k, z)$ are the Fourier components of the gravitational potential and matter perturbations, $j_l(x)$ are the spherical Bessel functions and χ is the comoving distance. These are calculated using a modified version of `cmbfast` [26,34].

B. Theoretical signal-to-noise ratio

Unfortunately, the ability to detect the cross correlation is limited because the signal falls off on small scales. Not only is cosmic variance an important factor, but there is also the problem of accidental correlations between the galaxy surveys and the CMB anisotropies produced at last scattering. Many independent measurements are needed to reduce the impact of such accidental correlations. The

signal-to-noise ratio (SNR) of the cross-correlation function (CCF) with a particular survey is given by

$$\left(\frac{S}{N}\right)^2 = \sum_l (2l+1) \frac{[C_l^{Tg}]^2}{C_l^{gg} C_l^{TT} + [C_l^{Tg}]^2}. \quad (11)$$

For the ISW, we are usually in the weak correlation regime, so that $C_l^{Tg}/\sqrt{C_l^{gg} C_l^{TT}} \ll 1$.

The SNR can be separated to obtain the contribution as a function of redshift; for a standard Λ CDM cosmology, most of the signal is expected at $z < 3$, with the peak around a redshift of $z \simeq 0.5$ [35]. While the signal is highest at low redshifts, more independent volumes are available for higher z . The SNR scales roughly as the square root of the fraction of the sky observed.

The most optimistic case is when the distribution of galaxies follows precisely the evolution of the ISW effect. In this case, $C_l^{Tg} = C_l^{gg} = C_l^{ISW}$ where C_l^{ISW} is the spectrum of the ISW temperature anisotropies alone, which is assumed to be much smaller than the total CMB anisotropy, $C_l^{ISW} \ll C_l^{TT}$. Thus, the SNR reduces to [10]

$$\left(\frac{S}{N}\right)^2 \simeq \sum_l (2l+1) \frac{C_l^{ISW}}{C_l^{TT}}, \quad (12)$$

This gives an optimistic total $S/N \simeq 7 - 10$ for a standard Λ CDM cosmology. The ISW constraints, which might arise from realistic future surveys, can be found in [36].

C. Correlation estimators

Our aim is to estimate the correlation between several galaxy surveys and the CMB: as described above, this measure can be performed in the real space using the CCF or in the harmonic space with the cross-correlation power spectrum. The two methods are theoretically equivalent for a full-sky analysis, and both have been used to detect the ISW cross-correlations. However, when one moves away from the ideal full-sky scenario, it is more straightforward to account for the sky mask using the real space correlations, and therefore, we will follow this approach here.

The matter density and CMB temperature as well as their projections onto the celestial sphere are in principle continuous fields; however, we only have access to the sampling of these fields experimentally obtained by measuring the CMB temperature in some fixed directions and counting the number of galaxies in a given patch of sky. In practice, we pixelise these maps using the HEALPix pixelization scheme [37], using a relatively coarse resolution: $N_{\text{side}} = 64$, corresponding to $N_{\text{pix}} = 49, 152$ pixels with dimensions $0.92^\circ \times 0.92^\circ$. This resolution is sufficient for large-scale correlations like the ISW and makes it tractable to perform large numbers of Monte Carlo simulations. A finer resolution ($N_{\text{side}} = 128$) was explored, but the results did not change significantly.

In making the maps, we assign the average temperature or the total number of galaxies to each pixel. The maps are masked according to the particular requirements for each catalogue and the most relevant foregrounds as discussed below. It is inevitable that some pixels are only partially covered in the original survey, either because only part of the area was observed, or because some of this area was masked out. In such cases, predominately occurring on the edge of the survey, the number of galaxies in a pixel is estimated as $n_i = n_i^{\text{obs}}/f_i$ where f_i is the fraction of the pixel observed. The mask was obtained through sampling all objects in each catalogue in a higher resolution ($N_{\text{high}} = 512$) as described in [20].

From these maps, both the auto- and cross-correlations were estimated, down weighting those pixels with partial coverage proportionally to f_i . For the autocorrelation functions (ACFs), we used the estimator,

$$\hat{C}(\vartheta) = \frac{1}{N_{\vartheta}} \sum_{i,j} f_i (n_i - \bar{n}) f_j (n_j - \bar{n}), \quad (13)$$

where \bar{n} is the average number of galaxies in a pixel for the survey of interest, and $N_{\vartheta} = \sum_{i,j} f_i f_j$ is the weighted number of pairs of pixels with separation ϑ . For the temperature maps, we simply replace n_i and \bar{n} with the pixel temperature and average temperature of the CMB maps.

More generally, we are interested in the cross-correlation function between the survey p and the survey q ; this is estimated similarly, accounting for the fact that the pixel weighting and mean number per pixel will depend on the survey

$$\hat{C}^{pq}(\vartheta) = \frac{1}{N_{\vartheta}^{pq}} \sum_{i,j} f_i^p (n_i^p - \bar{n}^p) f_j^q (n_j^q - \bar{n}^q). \quad (14)$$

The number of pairs of pixels at a given separation $N_{\vartheta}^{pq} = \sum_{i,j} f_i^p f_j^q$ will depend on both of the surveys under consideration. This again extends to the density-CMB CCFs in the obvious way.

We use $N_b = 13$ angular bins in the range $0^\circ < \vartheta < 12^\circ$. We use a linear binning, and have explored the dependence of our results on the choice of binning, changing both the number and trying a logarithmic binning; neither had significant impact on the results.

D. Covariance estimators

An important aspect of this calculation is the estimation of the covariance of the cross-correlation measurements. As described most recently by [38], there are a number of different ways to calculate the errors on this measurement, each with their own advantages and drawbacks. Here, we calculate our errors in three ways: a Monte Carlo method (MC1), where the covariance matrix is estimated by measuring the CCF between random CMB maps while keeping fixed the observed density map; a second Monte Carlo method (MC2), similar to MC1 but including also random

density maps, which are correlated at the expected level with the random temperature maps; and jackknife errors (JK) which are estimated by looking at the variance of the CCF when patches of the sky are removed.

The first approach is to generate random Monte Carlo maps of the CMB assuming the WMAP best-fit cosmology, and estimating the covariance matrix cross correlating these maps with the true density maps (MC1).

The WMAP third-year fiducial model we use throughout this paper has baryon density $\Omega_b = 0.04185$, matter density $\Omega_m = 0.2402$, Hubble constant $H_0 = 73.0$, scalar spectral index $n_s = 0.958$, optical depth $\tau = 0.092$, and amplitude of density fluctuations $A = 0.80$ at $k = 0.002 \text{ Mpc}^{-1}$.

The MC1 is the most widely used estimator in the literature, though here we extend the usual calculation to account for covariances between the CCFs of the CMB with different surveys. This method is reasonably fast to implement and accounts for the cosmic variance and the accidental correlations with the CMB, which are the primary source of error. However, it is asymmetrical, in that it does not account for the variance in the density maps or its Poisson noise; the MC1 method also assumes there are no cross correlations, though the expected (and observed) weakness of the cross correlation indicate that this should not introduce a large bias. Finally, like all Monte Carlo approaches, it is model dependent and could fail if the data model is poorly understood (e.g. foregrounds or non-Gaussianity of the maps).

However, some of these problems can be avoided if we also generate random density maps for each catalogue based on the WMAP cosmology and the redshift distribution, with the addition of Poisson noise to the maps (MC2). In this case, we have the ability to account for the expected correlations between the maps as described in the appendix. This method is more time demanding, in that it requires more random maps for each correlation measurement; it also retains the unwanted model dependence, and unlike the previous method has no explicit dependence on any of the observed maps.

To estimate the covariance between the different angular bins of a single CCF following the MC1 and MC2 methods for each catalogue k we use the following estimator of the full covariance matrix:

$$C_{ij} = \frac{1}{M} \sum_{k=1}^M [\hat{C}_k^{Tg}(\vartheta_i) - \bar{C}^{Tg}(\vartheta_i)][\hat{C}_k^{Tg}(\vartheta_j) - \bar{C}^{Tg}(\vartheta_j)], \quad (15)$$

where $\bar{C}^{Tg}(\vartheta_i)$ are the mean correlation functions in the i -th angular bin over M realizations; the diagonal part of these matrices gives the variance of the CCF in each bin $C_{ii}^k = \sigma_i^2$, while the off-diagonal part represents the covariance between the points.

The last method to estimate the covariance (jackknife) consists in estimating the variance by generating mock

density maps from the true ones, simply discarding a small patch of them. In practice, we can divide the original density map in M patches, which have roughly equal area, and discard in turn a different patch to calculate the CCF. The estimator for the covariance matrix is in this case

$$C_{ij} = \frac{M-1}{M} \sum_{k=1}^M [\hat{C}_k^{Tg}(\vartheta_i) - \bar{C}^{Tg}(\vartheta_i)] \times [\hat{C}_k^{Tg}(\vartheta_j) - \bar{C}^{Tg}(\vartheta_j)], \quad (16)$$

The advantage of this method is its model independence, but it has the big drawback of giving different answers depending on the size and number of the discarded areas. It also implicitly assumes independence of the various patches, which is not always the case.

Our ultimate goal is to measure the total covariance matrix between all the catalogues. To do so, we need to estimate the total covariance matrix C_{ij}^{pq} , as the matrix that has in the diagonal blocks the single catalogue C_{ij}^{pp} , and in the off-diagonal parts is

$$C_{ij}^{pq} = \frac{1}{M} \sum_{k=1}^M [\hat{C}_k^{Tp}(\vartheta_i) - \bar{C}^{Tp}(\vartheta_i)][\hat{C}_k^{Tq}(\vartheta_j) - \bar{C}^{Tq}(\vartheta_j)]. \quad (17)$$

For simplicity, we redefine the indexes i, j in a way that they run from 1 to $N_{\text{tot}} = N_{\text{bin}} \times N_{\text{cat}}$, i.e. redefining the data, theory, and mock arrays as the concatenation of all catalogues' CCFs with the CMB. In this way, the covariance matrix is simply the square matrix C_{ij} , identical to the Eq. (17) but now with dimension N_{tot} . A similar expression can be easily obtained for the JK case.

III. THE CATALOGUES

To best detect the ISW effect through the cross-correlation technique, we ideally require surveys covering large fractions of the sky, so that accidental correlations will cancel out. The surveys also need to be sufficiently deep, in order to probe the gravitational potentials where the ISW effect is being created. Ideally, we would like to span the redshift range $0 < z < 3$, separated into subsamples of different depths so as to measure the redshift dependence of the effect and get some handle on the evolution of the dark energy. However, only rather coarse redshift information is required, so redshift errors of $\Delta z \sim 0.1$ obtainable through photometric methods should be sufficient for these purposes. This is beyond the present state of the observations, but the differences in the redshift distributions of the various samples does provide some limited information on the dark energy evolution.

At present, the best surveys available for this purpose (and where ISW detections have previously been found) include the following: the optical Sloan Digital Sky Survey (SDSS), the infrared 2 Micron All-Sky Survey (2MASS),

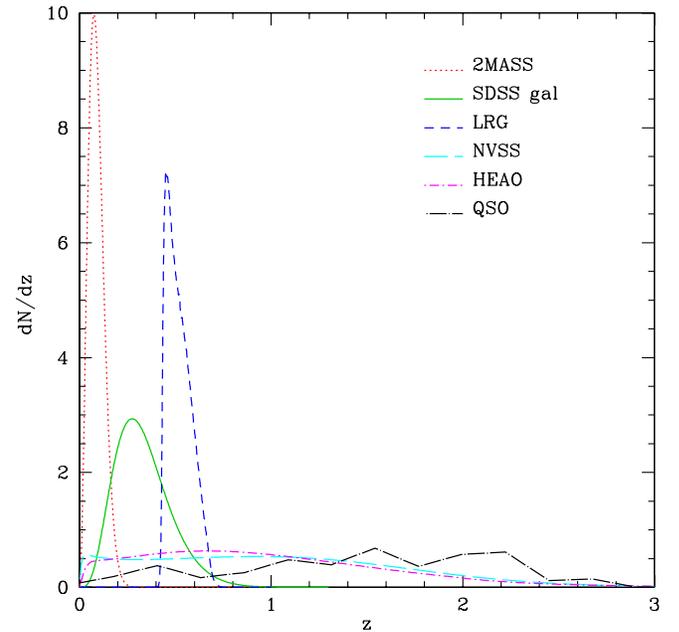


FIG. 1 (color online). The redshift distributions of all catalogues dN/dz normalized to unity. The significant overlap between redshift distributions (especially for the x-ray and radio surveys) results in a covariance matrix with significant non-diagonal elements.

the X-ray catalogue from the High Energy Astrophysical Observatory (HEAO) and radio galaxy catalogue from the NRAO VLA Sky Survey (NVSS). The high quality of the SDSS data allows us to extract some further subsamples from it, consisting of luminous red galaxies (LRG) and quasars (QSO) in addition to the main galaxy sample [39]. These are the samples we use in our analysis below, and include most of the significant reports of the ISW detection. Because the data has not been publicly released and since it is not significantly deeper than 2MASS, we omit the APM galaxy survey, which has also been reported to have evidence for ISW cross correlations [14].

We show in Fig. 1 the redshift distributions dN/dz of the catalogues we use, normalized to unity; we can see that they span a redshift range $0 < z < 2.5$, similar to the theoretical requirement, although the overlap between different samples is significant. This means that the covariance between the measures could be large: one of the goals of this paper is to quantify it.

In the rest of this section, we will present the characteristics of all the samples we use, in order of increasing redshift.

A. 2MASS

The 2 Micron All-Sky Survey (2MASS) is an infrared catalogue; its extended source catalogue [40] contains $\sim 800,000$ galaxies with median redshift $z \sim 0.1$ and, unlike the point source catalogue is almost free of stellar contamination. Some evidence for ISW cross correlations

has been seen in 2MASS previously [11,12], and we largely follow the galaxy selection of those previous analyses here.

Accordingly, we select galaxies according to their K_s -band isophotal magnitude K_{20} (`k_m_i_20_c`, -20 mag/arcsec²). These magnitudes are corrected for galactic extinction using the infrared reddening maps by [41], as $K'_{20} = K_{20} - A_K$, where the extinction is $A_K = 0.367(B - V)$. The requirement of completeness of the catalogue is satisfied by imposing a cut in magnitude $K'_{20} < 14.0$, while we can exclude low redshift sources with the condition $K'_{20} > 12.0$. We only include objects with a uniform detection threshold (`use_src = 1`), and remove known artifacts (`cc_flag ≠ a` and `cc_flag ≠ z`); we also exclude a small fraction of objects where the magnitude or its error were not recorded.

In addition to the pixelization geometry mask, we follow earlier analyses [11,12] excluding areas of the sky with high reddening, discarding pixels with $A_k > 0.05$; this leaves 69% of the sky and 718 000 galaxies after excluding artifacts. It is reported by [11,12] that the redshift distribution of these galaxies is well approximated by the function

$$\frac{dN}{dz} = \frac{1}{\Gamma(\frac{m+1}{\beta})} \beta \frac{z^m}{z_0^{m+1}} \exp\left[-\left(\frac{z}{z_0}\right)^\beta\right], \quad (18)$$

where the parameters are $z_0 = 0.072$, $\beta = 1.752$, and $m = 1.901$. This distribution is shown together with the others in Fig. 1.

To check the consistency of the dataset and its bias we calculate its autocorrelation function (ACF). The measure is in good agreement with the predictions for the best-fit WMAP model with a galactic bias $b_g = 1.4$ as found by [12], as we can see in Fig. 2.

B. SDSS galaxies

The SDSS Sixth Data Release (DR6) [42,43] is the largest wide optical galaxy survey available at the present for the northern hemisphere. From this catalogue we select a magnitude limited subsample $18 < r^* < 21$; this catalogue contains 30×10^6 galaxies. Here r^* is the extinction corrected r SDSS übercalibrated model magnitude, i.e. using the SDSS variables $r^* = \text{ubercal.modelMag}_r - \text{extinction}_r$: this corresponds to the procedure of [13], with the difference of using the data release DR6 and the

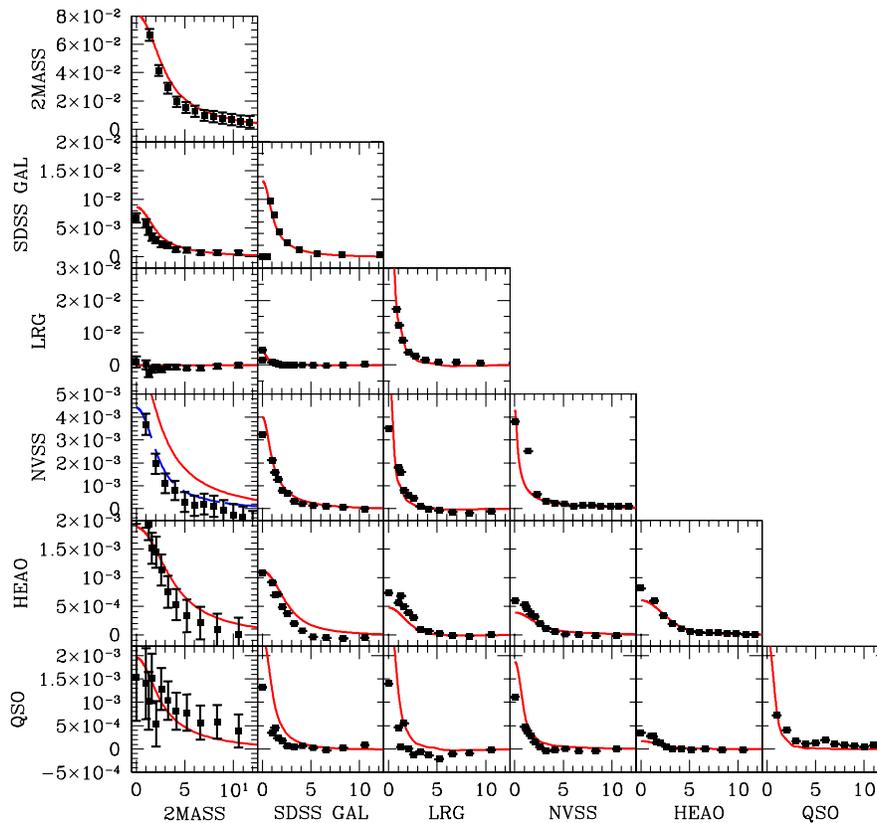


FIG. 2 (color online). Measures of the two-point correlation functions between all the combinations of catalogues, where the units in the x-axis are degrees. The autocorrelations are on the diagonal, and the solid (red) lines show the theory from WMAP best-fit cosmology and the galactic bias from the literature. The largest discrepancy with theory, in the NVSS-2MASS CCF, can be addressed by a small change in the assumed NVSS redshift distribution (blue dashed line).

übercalibrated model magnitude instead of the Petrosian magnitude, which is less reliable for faint objects. We apply the pixelization geometry mask and, in addition, we discard the pixels most affected by reddening, with $A_r > 0.18$. We also discard the southern stripes, since they are most affected by foregrounds and edge effects.

We select only objects with photometric redshifts between $0.1 < z < 0.9$ and with an error on the redshift $\sigma_z < 0.5z$, leaving 23.5×10^6 galaxies in the catalogue. We could use these photometric redshifts as the basis of the theoretical calculations; however, since the distribution of the photometric redshifts can be affected by singularities in the redshift determination procedure, we use instead a fit to their distribution with the smooth function of Eq. (18). The best-fit parameters are in this case $z_0 = 0.113$, $\beta = 1.197$, and $m = 3.457$, corresponding to a median redshift $z_{\text{med}} = 0.32$. (The results are actually independent of whether the fit or the actual redshift distribution is used). The fit is shown together with the others in Fig. 1.

The ACF is in agreement with the prediction for the WMAP best-fit cosmology and a bias $b_g = 1$, as we can see in Fig. 2.

C. SDSS LRG

LRGs from the SDSS have been used often to find evidence for the ISW effect, as they have a deeper redshift distribution than the ordinary galaxies, with a mean redshift of $z \sim 0.5$ [15–17]. In this analysis, we use the MegaZ LRG sample [44,45], which contains 1.5×10^6 objects from the SDSS DR6 selected with a neural network [46]. To ensure completeness, we require that $i < 20$. To reduce stellar contamination, we implement cuts on δ_{sg} , which is a variable of the MegaZ neural network estimator, defined such that $\delta_{sg} = 1$ if the object is a galaxy, and $\delta_{sg} = 0$ if it is a star [44]. Following the conservative suggestion by [44], we choose a cut $\delta_{sg} > 0.2$, which is reported to reduce stellar contamination below 2% while keeping 99.9% of the galaxies. Stricter cuts have been tried with no significant changes to the CCF.

The mask we apply to this catalogue is a combination of the pixel geometry mask and two foreground masks, to account for seeing (cutting pixels with median seeing in the red band greater than 1.4 arcsec) and reddening (cutting pixels with median extinction in the red band $A_r > 0.18$). The redshift distribution function in this case is found directly from the photometric redshifts that are given in the catalogue, and is shown in Fig. 1.

We show the autocorrelation function in Fig. 2, where we can see that this is in agreement with the theoretical prediction from the best-fit WMAP cosmology and a bias $b_g = 1.8$, which is compatible with the estimate $b_g = 1.7 \pm 0.2$ shown by [45], although some excess power at large scales is present, which might be explained as being produced by a residual stellar contamination.

D. NVSS

The NVSS is a flux limited radio survey at a frequency of 1.4 GHz, with a minimum flux of ~ 2.5 mJy. It is complete for declinations $\delta > -40^\circ$, covering roughly 80% of the sky and contains $1.8 \cdot 10^6$ sources. The mask to this catalogue is a combination of the most aggressive WMAP mask (*kp0*) plus a cut around point sources as described in [47], which also describes corrections made for a systematic in the mean density as a function of declination. The cross correlations between NVSS and WMAP have been observed by a number of groups, both in the correlation function [18,19] and using an array of wavelet techniques [21–24].

The redshift distribution is uncertain; we base ours on models of Dunlop and Peacock [48], which seem to be still widely accepted, and are largely consistent with observations of cross correlations with other surveys (though see below for further discussion). We calculate the autocorrelation function and present it in Fig. 2; there is good agreement with the theory from the WMAP best-fit model and a galactic bias $b_g = 1.5$, compatible with the result $b_g = 1.5 \pm 0.2$ by [47], although we see some excess power at small scales.

E. HEAO

The HEAO1-A2 data set is a full-sky flux map of hard x-rays counts in the 2–10 keV energy range [49]. We use the map and the mask determined by [50,51]: the map is masked for the galactic plane, a round area around the galactic center and patch areas around bright point sources. The redshift distribution is also uncertain and provided by modeling the x-ray background, as described in [50,52].

The modeling of the theoretical ACF for this catalogue is more complex than those considered above, in that we are looking at flux rather than number counts, and the experimental beam is large compared to the pixel size. (The point spread function of the beam is well modeled by a Gaussian with a full width, half maximum size of $\vartheta_{\text{FWHM}} = 3.04^\circ$ [51]). In addition, the number of photons is small, so there is an additional contribution from the photon shot noise. Thus, the observed correlation is the sum of three terms: the intrinsic correlations, the Poisson correlations due to finite numbers of sources, and shot noise due to the finite number of photons. The variance of the x-ray map is dominated by photon shot noise (41%) and Poisson correlations (45%), while intrinsic correlations are relatively small (14%). However, the shot noise contributes only to the 0° ACF, while the Poisson correlations fall off more quickly with angle than intrinsic correlations and become subdominant for $\theta > 4^\circ$. Consequently, the combination of shot noise and Poisson correlations are not the primary component of the total noise in the ISW signal. We can see in Fig. 2 that the total modeled ACF fits the observations on large angles, assum-

ing the WMAP best-fit model and a galactic bias $b_g = 1.06$, as found by [18].

F. SDSS QSO

The quasar survey we use comes from the SDSS DR6 through the NBC-KDE (non-Bayesian classifier–kernel density estimator) catalogue by [53,54], that contains over a million quasars. This new DR6 edition of the catalogue does not include as many parameter cuts as did the previous DR4 version. To obtain the cleanest possible data set, and for consistency with our previous measure of the cross correlation [20], we only used quasar candidates selected via the UV-excess-only criteria used in the previous version of this photometric quasar catalogue. In addition, we consider only objects with a *good* (positive) quality flag. Following our previous results [20], we impose a cut in reddening, discarding areas with $A_g > 0.18$. After these cuts, we are left with $N \approx 500,000$ quasars.

This catalogue comes with estimated photometric redshifts, upon which we base the redshift distribution shown in Fig. 1. There is evidence of some excess power in the ACF on large angular separations that indicate faint stars are still present in the catalogue after these cuts, as seen before in [20]. The amount of stellar contamination is $\sim 3\%$, as found by [54], from comparison with the ACF of a random sample of stars taken from the SDSS, and does not contribute to the correlation with the CMB, as expected. We can see in Fig. 2 the ACF for this sample; this is in good agreement with theoretical expectations and determines the bias of $b_g = 2.3$, as previously found in [20,55].

IV. RESULTS

In this section, we present the measurements of the all the correlation functions between the data sets we consider and their covariance.

A. Density-density cross correlations

We begin by examining the cross correlations between the different density maps. These measurements are shown in Fig. 2, with the autocorrelation measurements along the diagonal. This is the first measurement of the cross correlations between most of these data sets. The error bars are estimated by Monte Carlo realizations of all the data sets (MC2, as described above).

The measurements largely agree with their theoretical predictions, which are based on the WMAP best-fit model using the visibility functions in Fig. 1 and a linear bias for each. The agreement is to be expected for the autocorrelations, which were the basis for the estimates of the linear bias. However, the cross-correlation measurements provide a useful consistency check for our model, and, in particular, for the visibility functions, since the cross correlations

are most sensitive to the degree that the measurements overlap in redshift.

The largest discrepancy between the measurements and theory is in the NVSS-2MASS cross correlation, where the theory is roughly twice as large as expected. This is perhaps not unexpected, since the NVSS visibility function is known to be uncertain, and the overlap with 2MASS is in a narrow region of redshift. It does indicate that less of the NVSS correlations are arising from the 2MASS redshift range than expected in the model. This could be because either the low redshift tail of the NVSS visibility function is overestimated relative to the high redshift region, or because the bias of the radio galaxies increases as we move to higher redshift. This can be addressed by a small change in the visibility function, as demonstrated by the dashed (blue) line in the panel (in this case, we arbitrarily imposed a low redshift exponential damping in the visibility function, leaving the rest unchanged). Such a change does not significantly affect the expected CMB cross correlations considered here.

B. Temperature-density cross-correlations

We next turn our attention to the cross-correlation functions between each density map and the CMB maps from WMAP 3. We use the *internal linear combination* maps from WMAP, which are the cleanest data, although we have checked that the results do not depend on the frequency (see below), and we also apply the *kp0* mask to them, cutting the galactic plane region. As we can see in Fig. 3, the measures are again largely in agreement with the theoretical predictions for the WMAP best-fit model.

We have also checked the results obtained with the new WMAP 5 data, and we have not found any difference in the correlations. This is expected, since WMAP maps are already cosmic variance limited at large scales.

We now discuss the results obtained following the three methods of error estimation discussed in Sec. II above.

1. Temperature-only Monte Carlo errors

We generate 5000 Monte Carlo simulations of the CMB anisotropy map with the WMAP best-fit parameters. We estimate the covariance matrix for each catalogue using Eq. (15), and the total covariance matrix follows from its generalization.

These are the errors shown in the top panel of Fig. 3; as we can see, the errors are quite large, especially for the low redshift catalogues, and the significance is further decreased by the high correlation between the points. We have checked that these errors converge; the convergence is already good after ~ 700 Monte Carlos for each single catalogue, and after ~ 3000 Monte Carlos for the full covariance matrix. The covariance between the points is shown in Fig. 5.

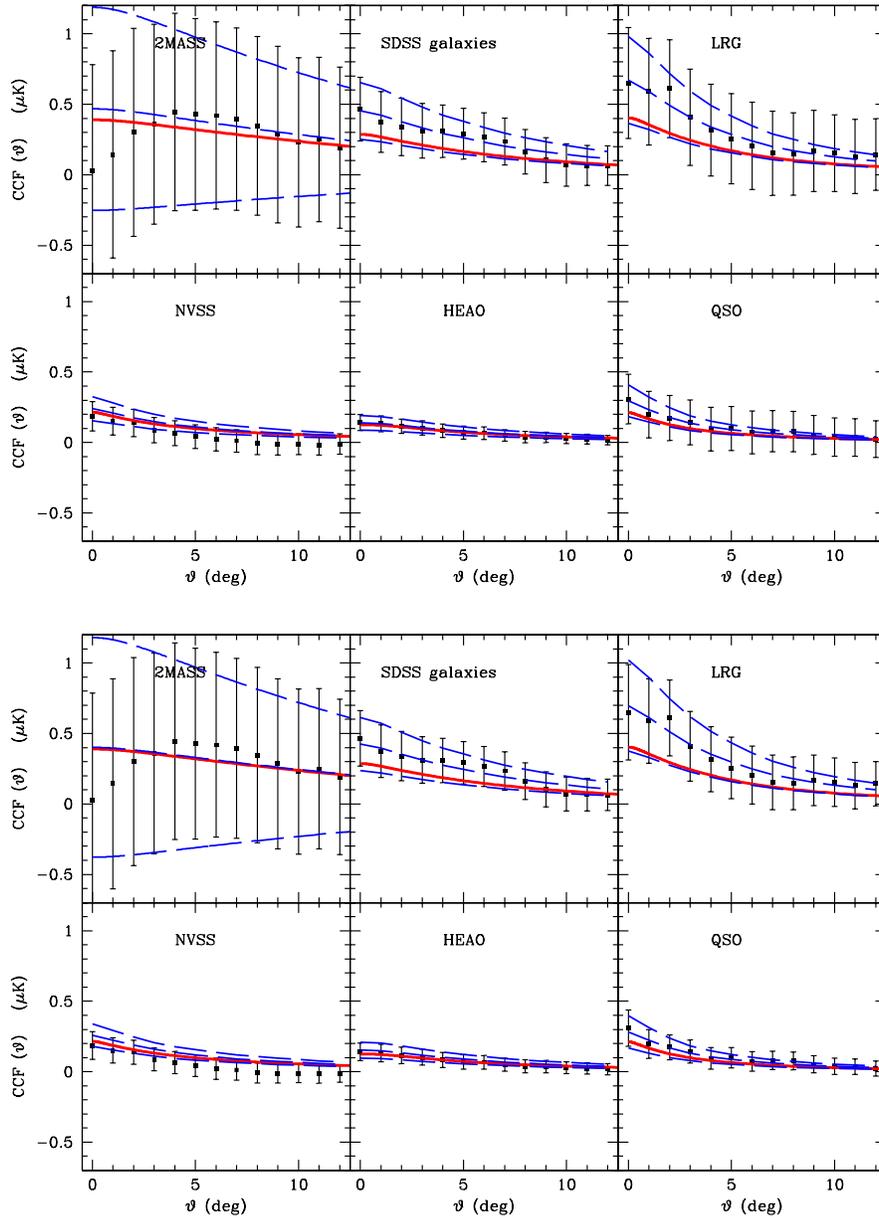


FIG. 3 (color online). Monte Carlo error estimation. Measurements of the cross-correlation functions between all the catalogues and the WMAP CMB maps (black points), compared with the theory from WMAP best-fit cosmology and the galactic bias from the literature (red solid lines). The best-fit amplitudes and their $1-\sigma$ deviations are shown in dashed (blue) lines. In the top panel, the errors are calculated with 5000 temperature-only Monte Carlos and, in the bottom panel, Monte Carlos for temperature and density including expected correlations. We see that the errors are comparable for individual observations. Because of known contamination from the Sunyaev-Zeldovich effect in the 2MASS data [11], the four smallest angle bins were excluded from the fits.

2. Full Monte Carlo errors

In this case, in addition to 5000 new mock CMB maps, we also generate 5000 mock density maps for each catalogue, correlated as expected theoretically, based on the WMAP best cosmology and their redshift distributions. In addition, the Poisson noise is added due to the expected number of objects per pixel.

The result calculated in this way is shown in the bottom panel of Fig. 3, and the relative full covariance matrix in

the bottom panel of Fig. 5. We can see that the errors estimated in this way are generally consistent with their MC1 counterparts.

The largest difference between the approaches is in the covariance between the cross correlations measured with different data sets (Fig. 5). Using the observed density maps yields both positive and negative covariance, while the covariance is only positive when all the maps are simulated. In the first approach, the strongest correlations are between the SDSS subsamples and 2MASS. In the

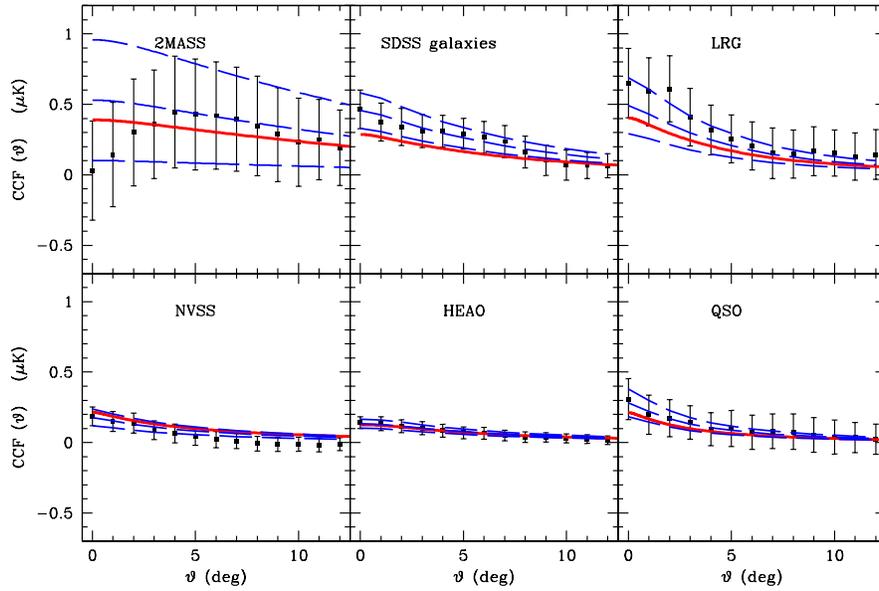


FIG. 4 (color online). Jackknife error estimation. The lines are the same as in Fig. 3. The errors are somewhat smaller than seen from the Monte Carlo estimates, possibly due to correlations between the jackknife subsamples.

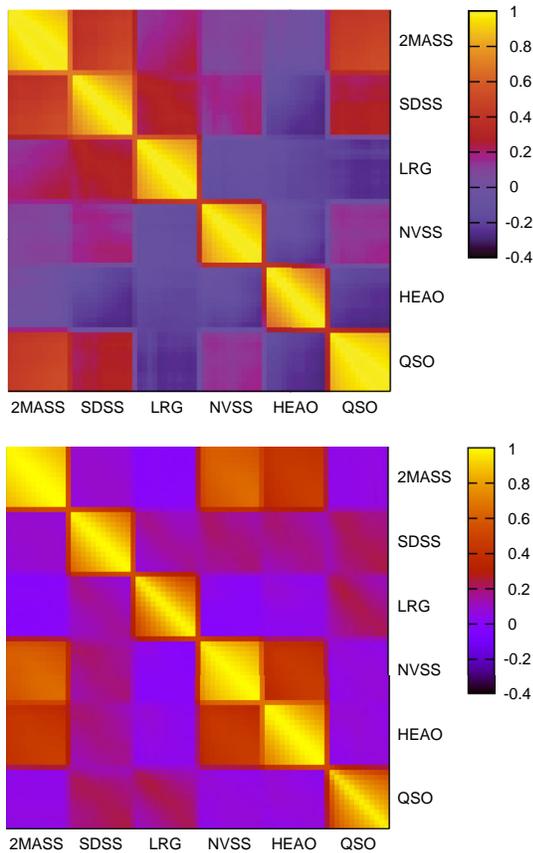


FIG. 5 (color online). The total covariance matrix obtained with 5000 Monte Carlos, normalized. The top panel shows the temperature-only Monte Carlos, while the bottom panel is the result of the full Monte Carlos. While the diagonal (single experiment) covariances are similar, those between experiments (off diagonal) are somewhat different.

second approach, which is purely theoretical, the largest covariances are between 2MASS, NVSS, and HEAO. The NVSS-HEAO covariance is expected to be large, since they are both essentially all-sky maps and have similar redshift coverage. The large covariance between 2MASS-HEAO and 2MASS-NVSS is more surprising given the differences in the redshift distributions, but seem to be driven by the low redshift tail of the NVSS and HEAO distributions. As noted above, the cross correlations are smaller than expected theoretically for 2MASS-NVSS (and to a lesser extent for 2MASS-HEAO). This indicates that the overlap of 2MASS with NVSS and HEAO is less than assumed, and that we have likely overestimated the covariance somewhat. However, the low significance of the 2MASS CCF means this has a small impact on the final result.

The differences between the two methods appear large for the off-diagonal elements. The reasons for these differences are unclear, but they suggest that the observed density maps are somewhat atypical of those simulated. However, it is not surprising that any particular realizations would appear atypical in some way. Despite these differences, these covariance matrices give comparable final significance, as is discussed below.

3. Jackknife errors

For completeness, we also present the errors estimated from a jackknife method. However, there is more ambiguity in implementing this method, leading to uncertainty in the resulting estimates in the errors.

One issue is what patch size to use. Ideally, for the jackknife approach one would like the cross-correlation

observations to be uncorrelated between patches. In reality, some correlation is inevitable. We also need enough patches to estimate the full covariance matrix without it becoming singular, which drives the size of the patches down. Thus, some kind of compromise is required.

Since we are interested in the CCF on scales of a few degrees, we choose a patch size of order 10 square degrees. Because the surveys have different geometry and masks, the number of sectors M will be different in each one. The number of patches we can have in this way is generally low (~ 100), so we cannot estimate the total covariance matrix which, having a dimension $N_{\text{tot}} = 78$, requires at least a few hundred independent random measures to be correctly estimated.

Cross-correlation measurements also introduce other issues, since the CMB and density maps are often covering different regions of the sky.

In the end, we tried to be conservative and ensure the most independence between the subsamples by only including data that were in the CMB and the density maps, and masking out both maps in the jackknife estimates. The results we obtained are shown in Table I, where we compare the results obtained with jackknife of the density map only and of both density and temperature maps using identical masks.

The jackknife ambiguities are even more problematic when calculating the full covariance between observations using different density maps, since the density maps often will not overlap on the sky. For this reason, and because such a large number of jackknives are required to estimate the total covariance matrix, we do not attempt to estimate it here.

The error bars we estimate using the jackknife method are of the same order of magnitude as those seen in the Monte Carlo approaches, but are somewhat smaller leading to higher significance in the detection. This could be due to the lack of independence of the jackknife patches, or

because some aspect of the data is missing from the Monte Carlo approach. We will use the Monte Carlo estimates below, focusing primarily on the results from MC2.

C. Foregrounds & systematics

Since the ISW effect is gravitational in origin, it is frequency independent as are the resulting CMB-density cross correlations. However, a frequency dependence may in principle be introduced by foregrounds and local contamination, such as the Sunyaev-Zel'dovich (SZ) effect. In Fig. 6, we compare the CCF obtained with the different frequency bands from WMAP (ILC, W, V, and Q bands), and we see that the result is substantially independent of frequency, with the exception of the 2MASS catalogue. However, the 2MASS CCF detection is of low significance, and our final answers are not greatly sensitive to its inclusion.

Foreground contamination of the ISW signal is generally produced at low redshifts. A good way to make sure that such effects are not dominating the measurement is to check for the sensitivity to the masking of these foregrounds (e.g. [28]). For samples derived from the SDSS (galaxies LRGs and QSO), we test for foreground effects by cutting the 20% of pixels with the highest reddening (extinction), seeing, sky brightness, and number of unresolved point sources. The most relevant masks are the reddening and seeing masks, which do not substantially change the results. For the other samples (2MASS, HEAO, and NVSS), we do not explore the masking, but we refer to the foreground analyses presented in earlier papers [11,12,47,51].

D. Comparison with previous measures

We briefly compare our CCF measurements to others in the literature.

TABLE I. The amplitudes and their significance for different methods of calculating the covariance. The left columns show the two Monte Carlo methods, while the right two show the jackknife method with equal area (10deg^2), in one case masking only patches of the density map, and in the other masking both density and temperature maps. We do not calculate the full covariance matrix or the total significance for the jackknife cases. For 2MASS, we have cut the first four angular bins because of their SZ contamination; the total significance is obtained discarding these bins.

Catalogue	5000 T-only Monte Carlos		5000 full Monte Carlos		JK - δ only		JK - δ and T	
	A	S/N	A	S/N	A	S/N	A	S/N
2MASS cut	1.22 ± 1.87	0.7σ	1.00 ± 1.96	0.5σ	0.66 ± 0.77	0.9σ	1.36 ± 1.10	1.2σ
SDSS	1.58 ± 0.70	2.2σ	1.48 ± 0.66	2.2σ	1.24 ± 0.42	3.0σ	1.59 ± 0.44	3.6σ
LRG	1.67 ± 0.76	2.2σ	1.73 ± 0.80	2.2σ	0.92 ± 0.50	1.8σ	1.22 ± 0.49	2.5σ
NVSS	1.12 ± 0.40	2.8σ	1.20 ± 0.37	3.3σ	0.68 ± 0.29	2.4σ	0.83 ± 0.27	3.1σ
HEAO	1.10 ± 0.41	2.7σ	1.22 ± 0.45	2.7σ	0.97 ± 0.26	3.7σ	1.00 ± 0.24	4.2σ
QSO	1.40 ± 0.53	2.6σ	1.33 ± 0.54	2.5σ	1.50 ± 0.58	2.6σ	1.33 ± 0.46	2.9σ
TOTAL	1.02 ± 0.23	4.4σ	1.24 ± 0.27	4.5σ	—	—	—	—

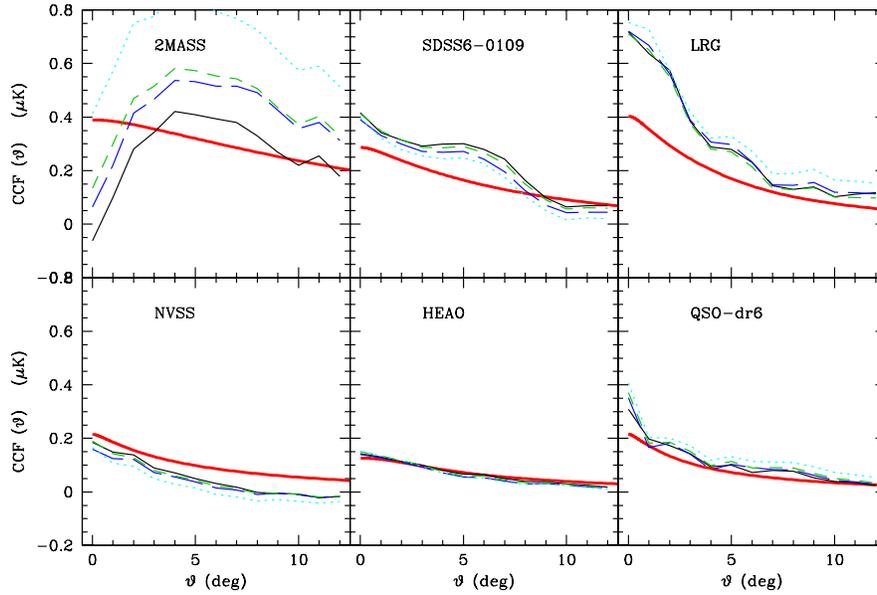


FIG. 6 (color online). Comparison of the CCF functions obtained with the different WMAP frequency bands. The black (solid) is using the internal linear combination map; the blue (long dashed) uses the W band, the green (short dashed) uses the V band and the cyan (dotted) uses the Q band. The thick (red) curves show the Λ CDM prediction.

1. 2MASS

From Fig. 3, it is clear that the CCF for the 2MASS survey is consistent with zero. Previous analyses of these data found some evidence for a positive correlation [11,12]; however, these were performed in Fourier space and included modeling of the SZ effect, which manifests itself with anticorrelations at small angular scales. Indeed, it appears in Fig. 3 that the observed CCF turns over at small angles. If the smallest four angular bins are removed, the fit to the CCF is consistent with the Λ CDM theory; however, it is only significant at the $\sim 1\sigma$ level. In any case, 2MASS appears to have the least significant evidence for cross-correlations.

2. SDSS galaxies

The main galaxy sample from the SDSS has a measured CCF, which is also in good agreement with the theory. In this case, we note that we do not find agreement with the previous result of [13], who reported a measured CCF of almost double the amplitude that we detect.

After discussions with the authors [13], we jointly found this discrepancy resulted from an additional cleaning cut, where they discarded all galaxies with a large error on their Petrosian r magnitude, imposing the condition $\text{petroMagErr}_r < 0.2$. Imposing this same condition, we found that we could reproduce their result. Further, masking those areas with high proportion of Petrosian error also gave similar results.

However, the motivation for such a cut is unclear. It is known that the Petrosian magnitudes are not accurate for faint objects, for which the best estimator is the `model`

magnitude [56]. While having objects with a well measured magnitude is desirable, we see no reason why cutting galaxies on the basis of a poor estimate of their magnitudes should double the correlation with the CMB. This could happen if it were produced by some foreground mechanism, such as seeing or reddening, but we checked that none of the possible foreground maskings raised the CCF in any way comparable to the aforementioned cut.

Therefore, lacking a valid reason to include this cut, and preferring to be conservative, we do not make the Petrosian error cut, and our CCF is thus lower than seen by Cabré *et al.* [13]. While it is worrying that a choice of masking has such a dramatic effect on the amplitude of the observed cross correlation, it should be noted that the cross correlation was largely independent of other masking choices.

3. SDSS MegaZ LRGs

The result for the LRG is the highest in comparison with the Λ CDM theory. It agrees with the result of [13]. A direct comparison with [16,17] is more difficult because these analyses use multiple photometric redshift bins. Concentrating on [17] (since it also does its analysis in physical space, rather than Fourier space), we find approximately the same detection significance as their single redshift bin measurements for similar data sets. An updated version of this paper (available on the astro-ph archive, but also unpublished) calculates a global χ^2 value using all four of their LRG samples, and detects a CCF with significance somewhat higher than we measure in this work. This is likely due in part to a somewhat larger redshift baseline for their measurement as well as the fact that they calculated their covariance matrix using a method similar to our

MC1 case. As one can see from Fig. 5, samples that cover very similar areas and have significant redshift overlap (as is the case with their LRG photometric redshift samples) can result in stronger anticorrelation between samples than one observes in covariance matrices generated with the MC2 method. This, in turn, would lead to a moderate overestimation of the detection significance.

4. Other measurements

Not surprisingly, since we use the same maps generated from HEAO and NVSS, our results are in agreement with previous measures by [18,19], and the amplitudes are consistent with the theoretical predictions. As discussed above, the new Monte Carlo approach gives consistent answers for individual experiments as the temperature-only Monte Carlo approach used in earlier analyses.

We found that the measured CCF for the quasars is consistent with the earlier measurement and the expectation from theory, and it is independent from the cleaning level of the catalogue.

In conclusion, all the measured CCF agree with the previous results and with the ISW theory for a Λ CDM model, although, they are in some cases marginally higher than theory predicts.

V. SIGNIFICANCE OF THE RESULT

Having established the measures of the CCFs and the total covariance matrix, we discuss the significance of this result and its consequences.

A. Single catalogue significance

Assuming that the detected cross correlations are due to the ISW effect, we can assign a significance value to the measure if the errors on the cross correlation are taken to be Gaussian. For each catalogue, we can compare the measured CMB-density cross correlation $\hat{C}(\vartheta_i)$ with the theoretical expectation obtained from the WMAP best-fit cosmological parameters with our modified version of the `cmbfast` code [34].

We perform the likelihood analysis first described in [50]. The shape of the CCF for each catalogue is assumed to follow the Λ CDM predictions. The theory template is

$$\bar{C}(\vartheta_i) = Ag(\vartheta_i), \quad (19)$$

where $g(\vartheta_i)$ is the theoretical prediction of the WMAP best-fit model and A is the fit amplitude, which will depend on the visibility function of the catalogue in question. Maximizing the likelihood of

$$\mathcal{L} = (2\pi)^{-N/2} [\det \mathcal{C}_{ij}]^{-1/2} \times \exp \left[-\sum_{ij} (\mathcal{C}_{ij})^{-1} (\hat{C}_i - \bar{C}_i)(\hat{C}_j - \bar{C}_j) / 2 \right], \quad (20)$$

we can find the best value for each A

$$A = \frac{\sum_{i,j=1}^N \mathcal{C}_{ij}^{-1} g_i \hat{C}_j}{\sum_{i,j=1}^N \mathcal{C}_{ij}^{-1} g_i g_j} \quad (21)$$

and the variance

$$\sigma_A^2 = \left[\sum_{i,j=1}^N \mathcal{C}_{ij}^{-1} g_i g_j \right]^{-1}. \quad (22)$$

We can also simply obtain the signal-to-noise ratio as $S/N = A/\sigma_A$.

The results obtained in this way with errors calculated with the three methods are summarized in Table I, and the resulting amplitudes and their errors can be seen in Figs. 3 and 4. Here we have allowed a separate amplitude A for each catalogue. Note that while the observed CCF is the same for the different methods, differences in the covariance matrices can result in different best-fit amplitudes.

It is possible to check that the Monte Carlo estimation has converged after N realizations by estimating the uncertainty on the errors. In detail, we use a jackknife approach consisting in observing the effect of removing $M = 10$ different subsets of the $N = 5000$ realizations of the MC2 method. The estimator of the uncertainty on S/N is

$$\sigma_{S/N}^2 = \frac{M-1}{M} \sum_{i=1}^M [(S/N)_i - \overline{S/N}]^2, \quad (23)$$

where $(S/N)_i$ are the signal-to-noise ratios obtained with each subset of $N - M$ Monte Carlos, and $\overline{S/N}$ is their average. We find in this way that the uncertainty on the S/N is less than 5%, indicating the level to which our Monte Carlos have converged.

B. Joint significance

We can easily generalize this to combine the different catalogues and obtain a single significance. Redefining the indexes i, j in a way that they now run from 1 to $N_{\text{tot}} = N_{\text{bin}} \times N_{\text{cat}}$, running over each of the bins of the observed (and theoretical) cross-correlation functions for each of the density catalogues. Using the full covariance matrix, we can follow again the same procedure, and find a single best-fit amplitude.

The results obtained in this way are shown at the bottom of Table I. The significance of the two different Monte Carlo methods, MC1 and MC2, are 4.4σ and 4.5σ , respectively. We also find that the uncertainty on the S/N for the joint amplitude is again less than 5%.

The two MC methods produce similar detection significances, but this could be a lucky coincidence, since the covariance matrices relating different surveys are much different. Both methods suggest some pairs of observations should be strongly correlated, but which pairs are strongly correlated depends on the method. If the covariance between surveys were ignored, the total significance would be about 5.8σ . Perhaps it is not surprising then that adding

similar levels of covariance between experiments with comparable individual detection levels would have a similar effect on the total significance.

As in the case of fits to individual correlation functions, strong covariances can have results that are counterintuitive. For example, the fit for the total amplitude using the MC1 approach is smaller than any single survey would suggest. Also, adding the small angle 2MASS CCF, believed to be suppressed by SZ, actually increases the fits by about 0.2σ , despite the points themselves being lower than the theory. These effects suggest that the degree of covariance between the different measurements might be overestimated, which would not be surprising given the much different systematics in each experiment. Even adding a small degree (5%) of diagonal noise is enough to increase the total MC1 amplitude to $A = 1.14 \pm 0.26$, with a corresponding $S/N = 4.4$, so that it is more consistent with the amplitudes of the individual experiments. The MC2 result is not affected by such a change, because the total amplitude is already consistent with the individual survey measurements.

Note that the theoretical model associated with a particular best-fit amplitude is not unique. While increasing the dark energy density will generally increase the ISW effect, the effect will generally be redshift dependent and could impact different catalogues differently. However, the Λ CDM model without any tweaking ($A = 1$) improves the likelihood at $\sim 4.5\sigma$ compared with the absence of cross correlations. Below we compare to specific alternative cosmologies without any scaling amplitude.

C. χ^2 Tests

Another way to assess the significance of the measure with respect to a theory is simply to look at the χ^2 , defined as

$$\chi^2 = \sum_{ij} \mathcal{C}_{ij}^{-1} (\hat{C}_i - \bar{C}_i)(\hat{C}_j - \bar{C}_j), \quad (24)$$

where the inverse covariance matrix and the data can be referred either to a single catalogue or to the total measure. Whereas the likelihood method discussed above looks at how well a model can reduce χ^2 , it is also worth simply looking at the magnitude of χ^2 for the null hypothesis test, where we calculate the χ_0^2 assuming the theoretical cross-correlation is zero.

In Table II, we show the χ^2 for the null hypothesis, as well as for the Λ CDM and best-fit models. We use the MC2 errors, dropping the first four bins of 2MASS, which appear to be affected by SZ. While there is much variation, in most cases there is not clear evidence against the null hypothesis, in that its χ_0^2 is not significantly greater than the number of data points. However, the χ^2 values are significantly reduced if one assumes one of the models, like Λ CDM, which predict a nonzero cross correlation.

TABLE II. A comparison of the absolute χ^2 for the various experiments.

Catalogue	f	χ_0^2	χ_{bestfit}^2	$\chi_{\Lambda\text{CDM}}^2$
2MASS	9	5.4	5.2	5.2
SDSS	13	17	11	12
LRG	13	9.6	4.9	5.7
NVSS	13	17	6.0	6.3
HEAO	13	18	10	10
QSO	13	9.7	3.7	4.0
TOTAL	74	67	47	48

The reasons for the particularly low χ^2 for the LRG case is unclear, and we investigate this more below. It might be an indication that the error estimates are in some sense too large, or that the covariance between angular bins is different than expected from the simple Monte Carlo simulations, perhaps as a result of foregrounds. However, it should be emphasized that the χ^2 for the null hypothesis is fairly conservative, and unlike the Bayesian likelihood approach, it fails to account for the fact that we have strong theoretical expectations for the signal we are looking for.

D. Eigenmode decomposition

To better understand the covariance of our data, and especially to understand the χ^2 , it is useful to study the eigenmode decomposition of the covariance matrix. As a worst-case example, we will use here the measurement and covariance matrix for the LRG sample calculated with the MC2 method (dimension $n = 13$).

We can factorize the covariance matrix into the form

$$\mathcal{C}_{ij} = \sum_{k,l=1}^n U_{ik}^T \Lambda_{kl} U_{lj}, \quad (25)$$

where $\Lambda_{ij} = \lambda_i \delta_{ij}$ is a diagonal matrix whose elements are the eigenvalues of \mathcal{C}_{ij} ; the rows of U_{ij} are the 13 eigenvectors $\hat{\mathbf{e}}_i$ of the covariance matrix. We plot the variances, $\lambda_i = \sigma_i^2$, in the top panel of Fig. 7, and some of the eigenvectors are shown in the bottom panel. There, we can see that the modes associated with the biggest variance are the low frequency ones, while the low variance modes oscillate significantly. This reflects the fact that the greatest differences between the Monte Carlo realizations is in the low frequency behavior of the cross-correlation functions.

Both the measured and theoretical CCFs can be decomposed into this eigenvector basis. In particular, any cross-correlation vector can be written as $\mathbf{v} = \sum_i A_i \hat{\mathbf{e}}_i$, where $A_i \equiv \mathbf{v} \cdot \hat{\mathbf{e}}_i$. We show in Fig. 8 the decomposition of the data and theory divided by the square root of the variance σ_i . For a typical CCF from the Monte Carlos, these amplitudes should be Gaussian distributed with unit variance. We can see how the smooth shape of the theoretical real space CCF is reflected in this eigenmode decomposition:

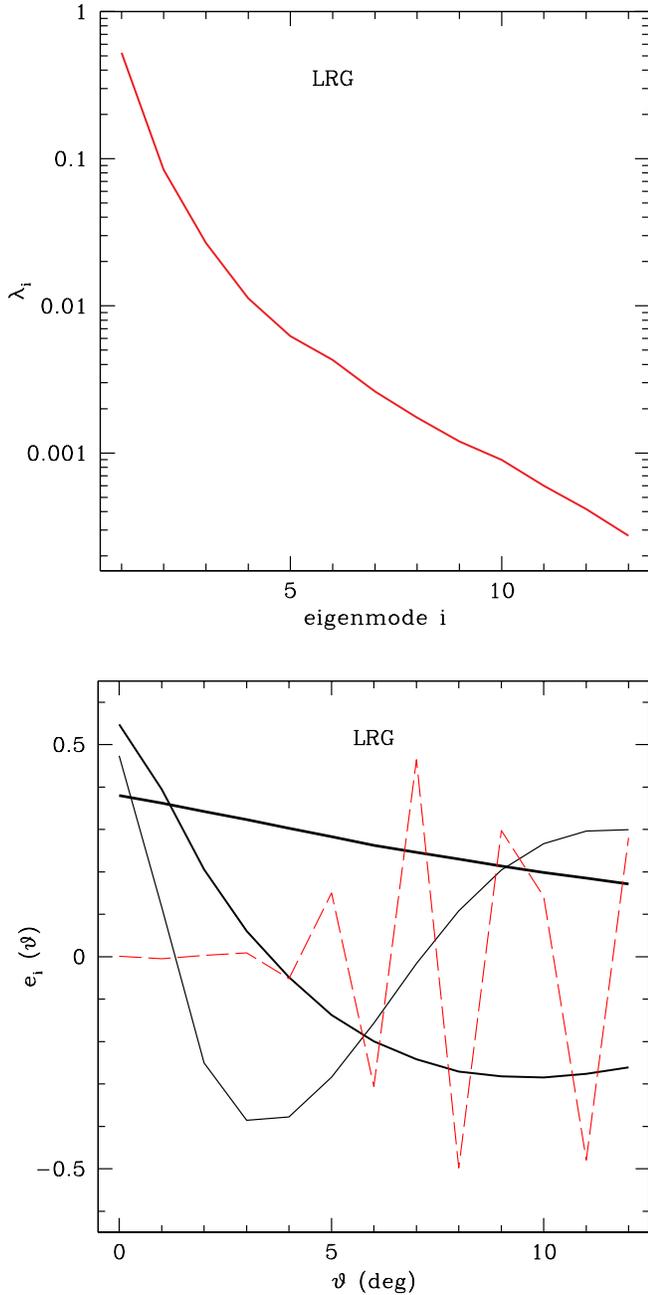


FIG. 7 (color online). Eigenvalues of the MC2 covariance matrices of the cross correlation between the LRG sample and the CMB (top panel), and first three eigenvectors (bottom panel). The dashed (red) line shows the highest frequency mode.

the theoretical amplitude is very well approximated by the first two modes only. However, this is not the case for the measured CCF, for which higher frequency modes are also significant.

We next look at the contributions to the χ^2 from the different eigenmodes. We show in Fig. 9 the evolution of the cumulative χ^2_i , i.e. the cumulative contribution to the χ^2 from each eigenmode. Here we compare the raw χ^2 from the observed cross-correlation function to that for the

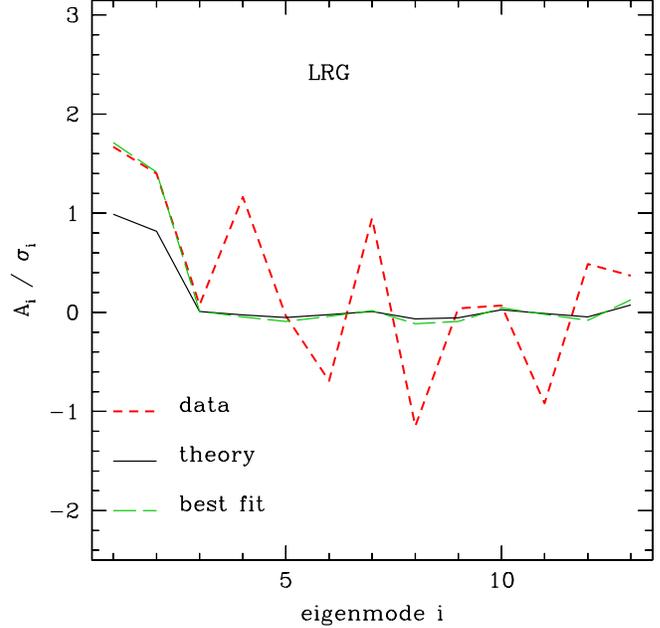


FIG. 8 (color online). Eigenmode decomposition of the amplitude of the measured (red dashed), theoretical (black solid) and best-fit (green long dashed) CCF.

residuals when the theoretical models (Λ CDM and the best-fit amplitude) are subtracted off. As expected, the theoretical models only impact the lowest two eigenmodes. The low χ^2 , however, is largely the result of the higher frequency modes, which seem to have slightly lower amplitudes than is seen in the Monte Carlos.

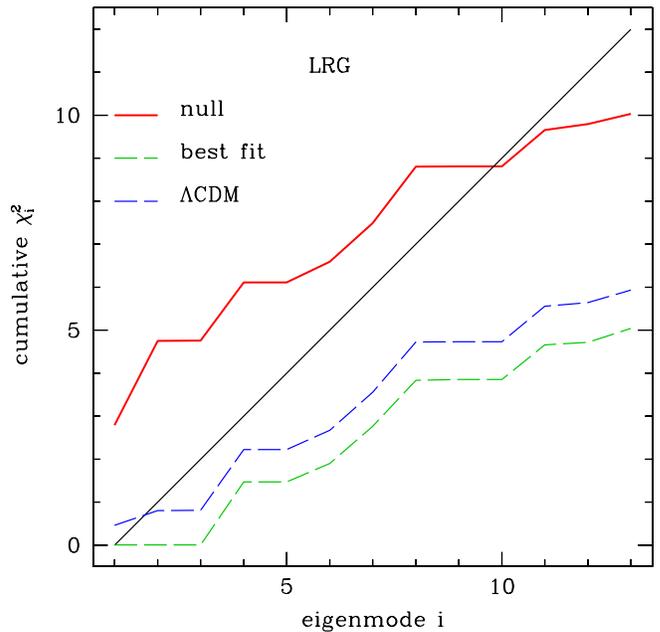


FIG. 9 (color online). Cumulative χ^2_i obtained summing the contribution up to the i -th eigenmode, for the three models: null hypothesis (red solid), best-fit and Λ CDM.

If we consider only those two modes that are expected theoretically, the χ^2 for the null hypothesis is actually fairly high: $\chi^2_2 = 4.8$. This would exclude the null hypothesis at more than the 90% level.

VI. COSMOLOGICAL CONSTRAINTS

Assuming the observed cross correlations are produced by the ISW effect, we can compare them with the theory predictions to obtain cosmological constraints. As described above, the ISW temperature anisotropies are produced as a result of time variation in the gravitational potential, and it is the evolution of the potential that our measurements constrain most directly. The cosmological parameters that impact the linear evolution of the potential are the dark energy density and its evolution, and the curvature of the Universe.

The actual cross-correlation measurements will also depend on the nature of the large-scale structure probe, its spectrum, and its bias. For example, if we normalize to the large-scale CMB, changing the shape of the power spectrum (e.g. by changing the Hubble constant or the dark matter density) will change the variance of the dark matter distribution on smaller scales, quantified by σ_8 . Since the ACFs of the surveys are fixed by observations, changing σ_8 effectively means a different bias will be inferred for each survey.

The cross correlations will rise and fall with the amount of structure in the probe. Thus, instead one could focus on the dimensionless cross correlation, effectively

$$r = \frac{C_\ell^{Tg}}{\sqrt{C_\ell^{TT} C_\ell^{gg}}} = \frac{C_\ell^{T\delta}}{\sqrt{C_\ell^{TT} C_\ell^{\delta\delta}}}, \quad (26)$$

which removes the dependence on bias (assuming it is linear) and probes more directly the ISW effect itself. The ISW effect arises on fairly large scales, e.g. $k \sim 0.01 h \text{ Mpc}^{-1}$, depending on the redshift distribution of the survey. Equivalently, for each model, we calculate the bias of each survey based on the observed ACF, and use this to find the predicted CCF for the model.

This makes our measurements largely independent of parameters other than Ω_{DE} , w , c_s , and Ω_k . In practice, we choose to keep the dark matter physical density fixed $\omega_m \equiv \Omega_m h^2 = 0.128$ to the WMAP best-fit value, but the constraints are largely independent of this assumption.

A. Models without dark energy

While many independent probes seem to indicate the existence of dark energy, it is worth exploring models which might account for the observations without dark energy; recently, an attempt has been made that does this, but which requires a significantly lower Hubble constant, modifications to the primordial power spectrum and other nonstandard features [57]. Such models would be dark matter dominated today, and have no late-time ISW

effect. Our observations of the ISW cross correlations rule out such models at the $\sim 4.5\sigma$ level, based on the difference in the χ^2 between the null hypothesis and the Λ CDM model in Table II. Such models also struggle to fit the recent observations of the angular scale of baryon oscillations [58].

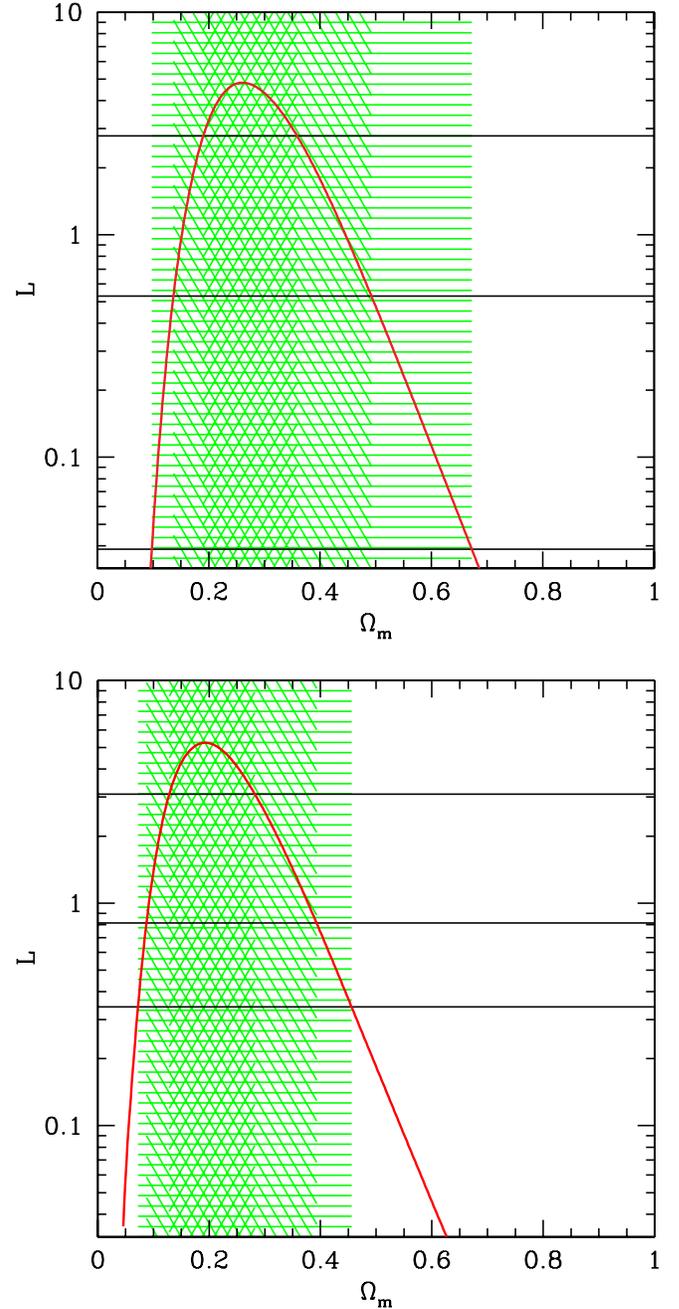


FIG. 10 (color online). Likelihood for flat models with varying Ω_m from the MC1 and MC2 errors. The shaded areas represent 1, 2, and 3 σ intervals for Ω_m . Λ CDM is a good fit to the data.

B. Flat Λ CDM models

Next, we study the likelihood of a family of flat models with varying Ω_m , $\Omega_\Lambda = 1 - \Omega_m$. As we can see in Fig. 10, the Λ CDM model is an excellent fit to our data: the 1σ interval for the parameter is $\Omega_m = 0.26^{+0.09}_{-0.07}$ using the MC1 covariance estimate. A higher ISW signal and slightly lower estimate for Ω_m results from the MC2 errors ($\Omega_m = 0.20^{+0.09}_{-0.07}$); this is due to the higher best-fit amplitude in this case. The error bars can be seen to be very

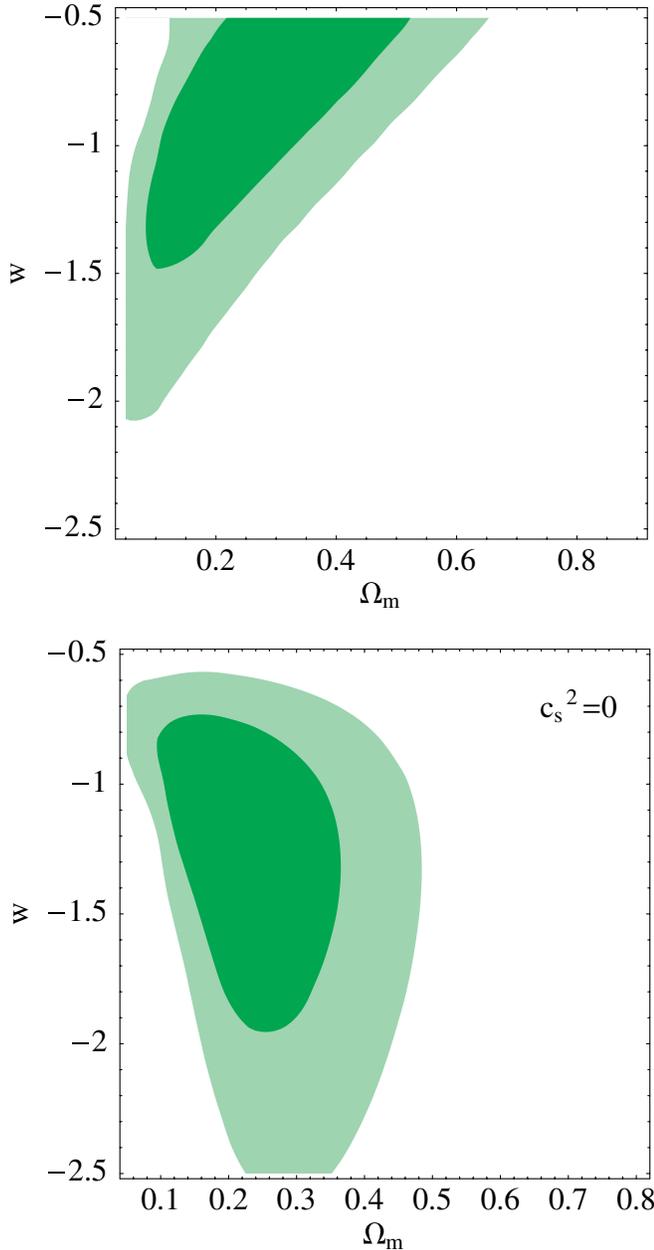


FIG. 11 (color online). Likelihood for flat models with varying Ω_m and w from the MC2 errors. The shaded areas represent 1 and 2σ intervals. The top panel assumes relativistic sound speed, such as would occur in a quintessence model, while the lower panel assumes the opposite extreme of zero sound speed.

asymmetric, as the ISW effect increases dramatically when the matter density becomes small. Models with $\Omega_m < 0.1$ would predict a much greater cross correlation than is observed.

C. Flat w CDM models

We next study the likelihood of a family of flat dark energy models, where we allow the dark energy density to evolve with equation of state w . The results are shown in Fig. 11, from which we can see that Λ CDM ($w = -1$) is very consistent with the measures. We can understand this if we observe that the measured excess in the ISW signal is largely redshift independent, while models along the same degeneracy line with a lower (higher) w would predict an excess at low (high) redshifts, respectively.

Initially, we assume that the dark energy sound speed is $c_s^2 = 1$, as is typical in scalar field models like quintessence. We also show the same range of models, but with a different dark energy sound speed $c_s^2 = 0$ in Fig. 11. We can see that in this case the degeneracy line changes direction due to the clustering of dark energy. Λ CDM is still a good fit to the data, as the cosmological constant likelihoods are not affected by the sound speed, and there is no clustering in that case.

The constraint on the sound speed itself is very weak. There are too many dark energy parameters (density, equation of state, sound speed) to expect any strong constraint. We reduce the numbers by assuming the CMB shift pa-

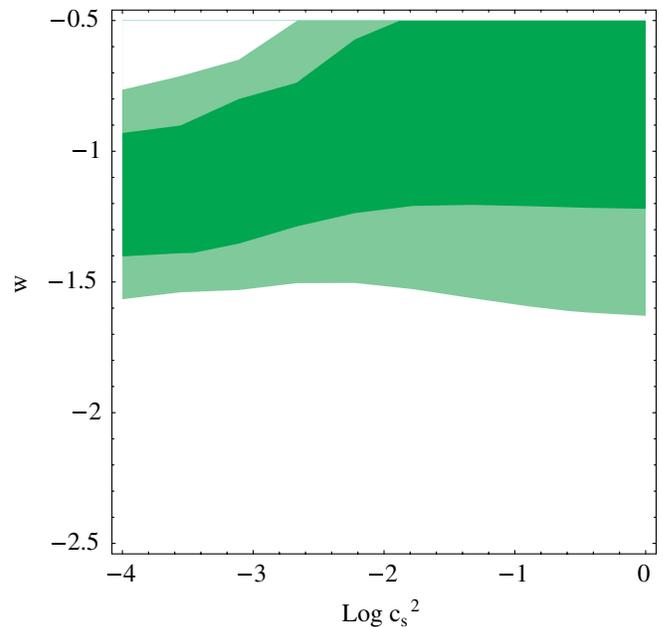


FIG. 12 (color online). Likelihood for flat models with dynamical dark energy as a function of the sound speed, where we fix the matter density based on the equation of state, assuming the CMB shift constraint. 1 and 2σ intervals are shown. No constraint is possible for the cosmological constant limit ($w = -1$).

parameter is fixed to the observed value, coupling the equation of state to the dark energy density. The results can be seen in Fig. 12. Even with this additional constraint, the sound speed is weakly constrained because the data are consistent with a Λ CDM model, where there is no dependence on the sound speed possible.

D. Curved Λ CDM models

Since curvature can also cause the gravitational potential to evolve, we explore the constraints if we relax the flatness condition. However, for simplicity we assume the dark energy is a cosmological constant. We study the likelihood of Ω_m , with a corresponding curvature $\Omega_k = 1 - \Omega_\Lambda - \Omega_m$. We explore the full $\Omega_m - \Omega_\Lambda$ space; we see the relative likelihoods in Fig. 13, which is obtained with MC2 errors.

From this figure, we see that Λ CDM is still a good fit to the data. An interesting feature of this figure is the degeneracy line between Ω_m and Ω_Λ : this is related to the relative efficiency of the curvature and dark energy as sources of ISW. Closed models (above the flat line) give negative ISW, and can cancel the effect of increasing the cosmological constant, while the opposite happens for open models (below the flat line).

E. Comparison with other constraints

Finally, we wish to compare the ISW constraints to those arising from other cosmological observations, including the CMB power spectrum, baryon oscillations, and Type Ia supernovae. For the latter, we use measurements

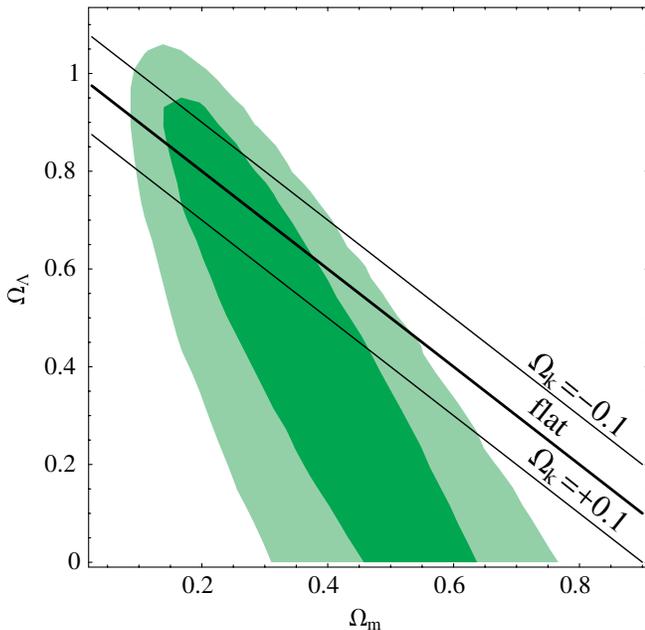


FIG. 13 (color online). Likelihood for curved models with varying Ω_m and Ω_Λ from the MC2 errors. The shaded areas represent 1 and 2σ intervals. Λ CDM is a good fit to the data.

of the luminosity distance from the Supernova Legacy Survey [1].

For the CMB observations, most of the dark energy information (at least that independent of the ISW effect) is distilled in the CMB *shift* parameter, defined as

$$R \equiv \sqrt{\Omega_m} H_0 \cdot (1 + z_\star) d_A(z_\star), \quad (27)$$

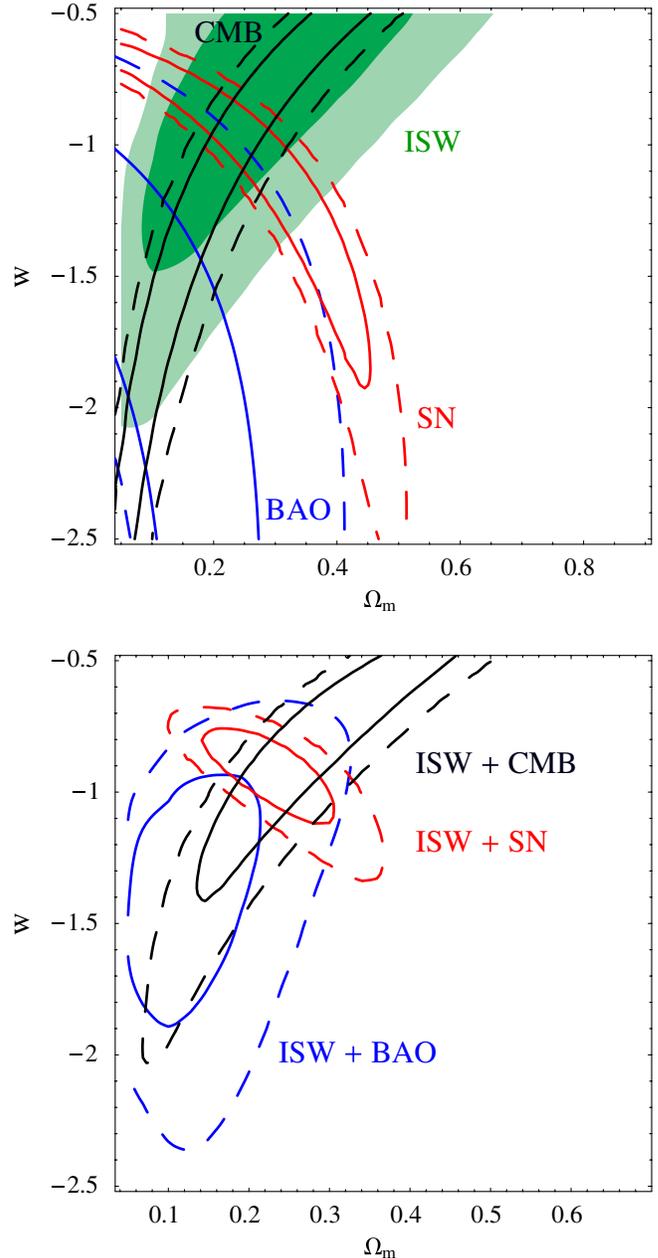


FIG. 14 (color online). Comparison with constraints from other observations, including CMB shift (black), SNe (red), and BAO (blue) (top panel), and combined likelihoods using the ISW + each one of these other constraints (bottom panel, same color coding). 1 and 2σ contours are shown (solid and dashed lines, respectively). The MC2 errors are used.

where $d_A(z)$ is the angular diameter distance and z_* is the redshift of the last scattering surface ($z_* = 1090$); this expression in the flat case reduces to

$$R = \sqrt{\Omega_m} H_0 \int_0^{z_*} \frac{dz'}{H(z')}. \quad (28)$$

R has been measured to be $R = 1.70 \pm 0.03$ [59]. We can see from Figs. 14 and 15 that this constraint has a degeneracy direction parallel to the ISW degeneracy in the flat case, but is less so in the general curved case.

Finally, for the BAO measurements [6] we use the constraint on the volume distance measure defined as

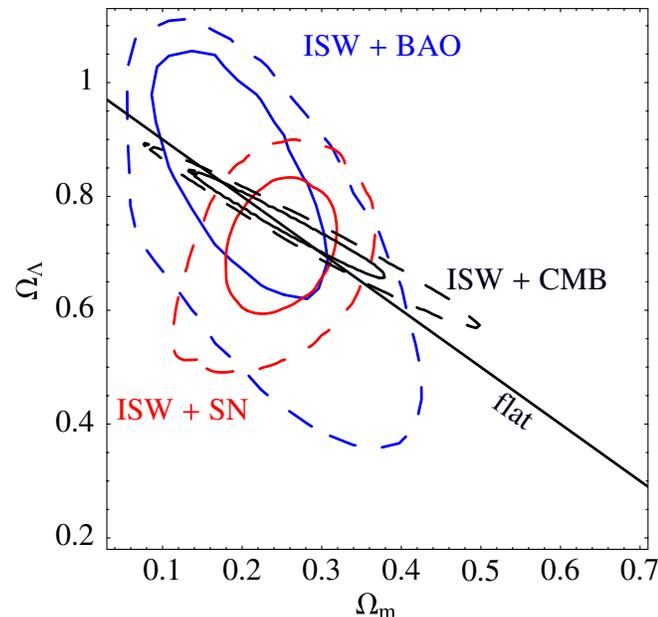
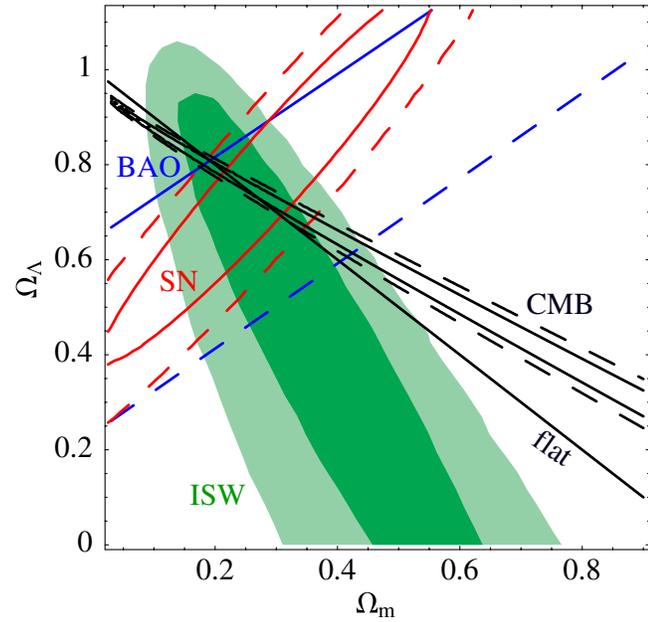


FIG. 15 (color online). Same as Fig. 14 for the curved case.

$$d_V(z) \equiv [(1+z)^2 d_A^2(z) z c / H(z)]^{1/3}. \quad (29)$$

The constraint on this parameter by [6] is $d_V(0.35)/d_V(0.2) = 1.812 \pm 0.060$.

The SN data is orthogonal to the ISW constraints, and jointly, they are consistent with the Λ CDM model; there is little evidence for additional curvature or evolving dark energy. The CMB shift constraint is similarly consistent with the cosmological constant concordance model, though the constraints are not as orthogonal to the ISW constraints. The Λ CDM model preferred by the SN and ISW measurements is consistent with the CMB shift combined with the measurements of the Hubble constant from the HST Key Project [60].

The exception to this concordance picture comes when the BAO data is considered. The BAO contours are similar to those from the SN, but shifted. In the flat dark energy case, the combination with the ISW prefers a larger dark energy density, which has increased with time (phantom). When all observations are combined, the BAO data are swamped by the SN data, and the result is fully consistent with the concordance model as found by [6].

VII. CONCLUSIONS

In this paper, we have measured the cross correlation between the CMB and a large range of probes of the density in a consistent way, and have calculated their covariance taking into account their overlapping sky coverage and redshift distributions. While individual measurements vary somewhat depending on how the data are cleaned and how the covariance is calculated, the overall significance of the detection of cross correlations is at the $\sim 4.5\sigma$ level.

These observations provide important independent evidence for the existence and nature of the dark energy. The observed cross correlations are consistent with the expected signal arising from the ISW effect in the concordance model with a cosmological constant. The observed signal is slightly higher than expected, higher than the expectation from WMAP best-fit model by about 1σ , thus favoring models with a lower Ω_m . However, we do not see any significant trend for the excess as a function of redshift, and so there is no indication of an evolving dark energy density. By combining these results with other cosmological data, we find a generally consistent picture of the behavior of the Universe, which is converging towards the Λ CDM model, although the uncertainties remain considerable. The only partial exception to this picture is the BAO result which, even when combined with our ISW measurement, is in slight tension with the Λ CDM model (at $\sim 1\sigma$).

The results of our analysis and the covariance matrices are available upon request from the authors (Giannantonio) and from [61].

ACKNOWLEDGMENTS

We thank Elisabetta Majerotto, Will Percival, Björn Malte Schäfer, and Jussi Väliiviita for useful conversations. We also thank Filipe Abdalla, Chris Blake, and Ofer Lahav for the use of the DR6 MegaZ data, and Anna Cabré and Enrique Gaztañaga for useful discussions relating to their earlier measurements. We also acknowledge Pier-Stefano Corasaniti and Alessandro Melchiorri who collaborated on the initial version of the modified Boltzmann code, that is publicly available at <http://www.astro.columbia.edu/~pierste/ISWcode.html>. We acknowledge support by a grant from STFC. While this work was in the final stages of preparation, a similar work appeared by Ho *et al.* [62], which reexamines some of the same data sets using a multipole space approach.

APPENDIX: CORRELATED MONTE CARLO MAPS

1. Basics

Here, we describe how to make Gaussian maps with a prescribed set of auto- and cross-correlation functions for use in the estimation of covariance matrices. Let us assume we have n maps, which could include temperature and various density maps at different redshifts or frequencies. Let us call these maps \mathbf{m}_i , where i ranges from 1 to n .

Any two maps, \mathbf{m}_i and \mathbf{m}_j , will be correlated, and these correlations will be described by a correlation function $C^{ij}(\vartheta)$ and associated multipole moments C_ℓ^{ij} . These correlations will be symmetric under interchange of the maps $C^{ij}(\vartheta) = C^{ji}(\vartheta)$, so we have $n(n+1)/2$ correlation functions or spectra, which describe the two maps.

Most map making algorithms, like `synfast` [37], work in Fourier or spherical harmonic space. Effectively, every mode is given a random amplitude ξ , which is a complex number with unit variance and zero mean $\langle \xi \xi^* \rangle = 1$ and $\langle \xi \rangle = 0$. These are then multiplied by the square root of the power spectrum in order to ensure the proper correlation functions. (There are additional constraints to preserve the reality of the fields on the lattice, e.g. $\xi_{\mathbf{k}} = \xi_{-\mathbf{k}}^*$, but it is not necessary to go through these here.)

It is sufficient to consider a single mode or harmonic amplitude of each map, as all the others will be similar but independent. Assuming we are working with spherical harmonics, we want to ensure that

$$\langle a_{\ell m}^i a_{\ell' m'}^{j*} \rangle = C_\ell^{ij} \delta_{\ell\ell'} \delta_{mm'}. \quad (\text{A1})$$

The δ functions follow simply from using uncorrelated random amplitudes for each harmonic mode. For a single map, the right power spectrum is ensured by simply using

$$a_{\ell m}^i = \sqrt{C_\ell^{ii}} \xi, \quad (\text{A2})$$

and this is effectively the prescription used by `synfast`.

When considering more maps, it is necessary to use more random phases, building the final maps from a com-

bination of different maps. With n maps, n different phases are required for each mode. Here, we denote the different phases with Latin letters, a, b, c, \dots . Different phases will be assumed uncorrelated, so that $\langle \xi_a \xi_{a'}^* \rangle = \delta_{aa'}$.

The simplest example is to consider two correlated maps, \mathbf{m}_1 and \mathbf{m}_2 . These are described by three spectra: C_ℓ^{11} , C_ℓ^{12} , and C_ℓ^{22} . These are made using the amplitudes

$$\begin{aligned} a_{\ell m}^1 &= \xi_a \sqrt{C_\ell^{11}} \\ a_{\ell m}^2 &= \xi_a C_\ell^{12} / \sqrt{C_\ell^{11}} + \xi_b \sqrt{C_\ell^{22} - (C_\ell^{12})^2 / C_\ell^{11}}. \end{aligned} \quad (\text{A3})$$

It is simple to verify that with these amplitudes $\langle a_{\ell m}^1 a_{\ell m}^{1*} \rangle = C_\ell^{11}$, $\langle a_{\ell m}^1 a_{\ell m}^{2*} \rangle = C_\ell^{12}$, and $\langle a_{\ell m}^2 a_{\ell m}^{2*} \rangle = C_\ell^{22}$.

This is simple to implement with `synfast`. First create a map with power spectrum C_ℓ^{11} , and then make a second map using the *same seeds* and power spectrum $(C_\ell^{12})^2 / C_\ell^{11}$. Add this second map to a third map made with a new seed and with power $C_\ell^{22} - (C_\ell^{12})^2 / C_\ell^{11}$. Note that this should never require taking the square root of a negative number; however, if it is very strongly correlated, numerical errors could cause problems. However, for the weak correlations considered here, this is never an issue.

The only difficulty is that this inherently produces positive correlations, as the default of the `synfast` code. This can be worked around simply. For example, if C_ℓ^{12} is always negative, one can simply flip the signs of the second map after it is produced. If instead C_ℓ^{12} changes sign, then break up the power spectrum into positive and negative pieces, making a map for each and subtracting the “negative” map from the “positive” map.

2. The general case

Next, we consider an arbitrary number of maps. For simplicity, we drop the ℓ and m subscripts where the meaning is unambiguous. Effectively, the challenge is to solve for a particular set of amplitudes \mathbf{T} , where

$$\begin{aligned} a^1 &= \xi_a T_{1a} & a^2 &= \xi_a T_{2a} + \xi_b T_{2b} \\ a^3 &= \xi_a T_{3a} + \xi_b T_{3b} + \xi_c T_{3c} \\ a^4 &= \xi_a T_{4a} + \xi_b T_{4b} + \xi_c T_{4c} + \xi_d T_{4d}, \end{aligned} \quad (\text{A4})$$

etc., subject to the constraints that $\langle a^i a^{j*} \rangle = C^{ij}$.

One thus has $n(n+1)/2$ equations with the same number of unknowns T . These begin as

$$\begin{aligned} C^{11} &= T_{1a}^2 & C^{12} &= T_{1a} T_{2a} & C^{22} &= T_{2a}^2 + T_{2b}^2 \\ C^{13} &= T_{1a} T_{3a} & C^{23} &= T_{2a} T_{3a} + T_{2b} T_{3b} \\ C^{33} &= T_{3a}^2 + T_{3b}^2 + T_{3c}^2 & C^{14} &= T_{1a} T_{4a} \\ C^{24} &= T_{2a} T_{4a} + T_{2b} T_{4b} \\ C^{34} &= T_{3a} T_{4a} + T_{3b} T_{4b} + T_{3c} T_{4c} \\ C^{44} &= T_{4a}^2 + T_{4b}^2 + T_{4c}^2 + T_{4d}^2 \end{aligned} \quad (\text{A5})$$

etc. While quadratic, these can be solved in stages linearly. Solve the first for $T_{1a} = \sqrt{C^{11}}$. Use the second to show $T_{2a} = C^{12}/\sqrt{C^{11}}$ and the third to get $T_{2b} = \sqrt{C^{22} - (C^{12})^2/C^{11}}$. This reproduces what was shown above.

After this, things continue similarly. At each point, we use the next equation to solve for the next missing variable

$$\begin{aligned}
 T_{1a} &= \sqrt{C^{11}} & T_{2a} &= C^{12}/\sqrt{C^{11}} \\
 T_{2b} &= \sqrt{C^{22} - (C^{12})^2/C^{11}} & T_{3a} &= C^{13}/\sqrt{C^{11}} \\
 T_{3b} &= (C^{23} - C^{12}C^{13}/C^{11})/\sqrt{C^{22} - (C^{12})^2/C^{11}} \\
 T_{3c} &= \left[C^{33} - (C^{13})^2/C^{11} - \frac{(C^{23} - C^{12}C^{13}/C^{11})^2}{C^{22} - (C^{12})^2/C^{11}} \right]^{1/2} \\
 T_{4a} &= C^{14}/\sqrt{C^{11}} T_{4b} \\
 &= (C^{24} - C^{12}C^{14}/C^{11})/\sqrt{C^{22} - (C^{12})^2/C^{11}}, \quad (\text{A6})
 \end{aligned}$$

etc. Things will take similar forms as one goes on, but become progressively more complicated.

It can also be programmed recursively, which may be simpler to implement. By this, we mean

$$\begin{aligned}
 T_{1a} &= \sqrt{C^{11}} & T_{2a} &= C^{12}/T_{1a} & T_{2b} &= \sqrt{C^{22} - T_{2a}^2} \\
 T_{3a} &= C^{13}/T_{1a} & T_{3b} &= (C^{23} - T_{2a}T_{3a})/T_{2b} \\
 T_{3c} &= \sqrt{C^{33} - T_{3a}^2 - T_{3b}^2} & T_{4a} &= C^{14}/T_{1a} \quad (\text{A7}) \\
 T_{4b} &= (C^{24} - T_{2a}T_{4a})/T_{2b} \\
 T_{4c} &= (C^{34} - T_{3a}T_{4a} - T_{3b}T_{4b})/T_{3c} \\
 T_{4d} &= \sqrt{C^{44} - T_{4a}^2 - T_{4b}^2 - T_{4c}^2}
 \end{aligned}$$

etc., with each step using only variables already solved. The general recursive expression for these spectra is

$$\begin{aligned}
 T_{ij} &= \sqrt{C^{ji} - \sum_{k=1}^{j-1} T_{ik}^2}, & \text{if } i = j \\
 T_{ij} &= \frac{C^{ji} - \sum_{k=1}^{j-1} T_{ik}T_{jk}}{T_{jj}}, & \text{if } i > j.
 \end{aligned} \quad (\text{A8})$$

These amplitudes are squared for the input spectra for `synfast`, but one must beware negative cross correlations as discussed above. A simple modification to a program like `synfast` could enable it to read in amplitudes rather than spectra, and this would be more efficient compared with reversing the sign of the maps after they are created.

-
- [1] P. Astier *et al.* *Astron. Astrophys.* **447**, 31 (2006)
- [2] A. G. Riess *et al.* (Supernova Search Team Collaboration), *Astrophys. J.* **607**, 665 (2004).
- [3] G. Hinshaw *et al.*, *Astrophys. J. Suppl. Ser.* **170**, 288 (2007).
- [4] D. J. Eisenstein *et al.* (SDSS Collaboration), *Astrophys. J.* **633**, 560 (2005).
- [5] S. Cole *et al.* (2dFGRS Collaboration), *Mon. Not. R. Astron. Soc.* **362**, 505 (2005).
- [6] W. J. Percival, S. Cole, D. J. Eisenstein, R. C. Nichol, J. A. Peacock, A. C. Pope, and A. S. Szalay, *Mon. Not. R. Astron. Soc.* **381**, 1053 (2007).
- [7] S. W. Allen, R. W. Schmidt, H. Ebeling, A. C. Fabian, and L. van Speybroeck, *Mon. Not. R. Astron. Soc.* **353**, 457 (2004).
- [8] S. W. Allen, D. A. Rapetti, R. W. Schmidt, H. Ebeling, G. Morris, and A. C. Fabian, arXiv:0706.0033.
- [9] R. K. Sachs and A. M. Wolfe, *Astrophys. J.* **147**, 73 (1967).
- [10] R. G. Crittenden and N. Turok, *Phys. Rev. Lett.* **76**, 575 (1996).
- [11] N. Afshordi, Y. S. Loh, and M. A. Strauss, *Phys. Rev. D* **69**, 083524 (2004).
- [12] A. Rassat, K. Land, O. Lahav, and F. B. Abdalla, *Mon. Not. R. Astron. Soc.* **377**, 1085 (2007).
- [13] A. Cabré, E. Gaztañaga, M. Manera, P. Fosalba, and F. Castander, *Mon. Not. R. Astron. Soc.* **372**, L23 (2006).
- [14] P. Fosalba and E. Gaztañaga, *Mon. Not. R. Astron. Soc.* **350**, L37 (2004).
- [15] P. Fosalba, E. Gaztañaga, and F. Castander, *Astrophys. J.* **597**, L89 (2003).
- [16] N. Padmanabhan, C. M. Hirata, U. Seljak, D. Schlegel, J. Brinkmann, and D. P. Schneider, *Phys. Rev. D* **72**, 043525 (2005).

- [17] R. Scranton *et al.* (SDSS Collaboration), arXiv:astro-ph/0307335.
- [18] S. Boughn and R. Crittenden, *Nature (London)* **427**, 45 (2004).
- [19] M. R. Nolta *et al.*, *Astrophys. J.* **608**, 10 (2004).
- [20] T. Giannantonio *et al.*, *Phys. Rev. D* **74**, 063520 (2006).
- [21] P. Vielva, E. Martinez-Gonzalez, and M. Tucci, arXiv:astro-ph/0408252.
- [22] J. D. McEwen, P. Vielva, M. P. Hobson, E. Martinez-Gonzalez, and A. N. Lasenby, *Mon. Not. R. Astron. Soc.* **376**, 1211 (2007).
- [23] D. Pietrobon, A. Balbi, and D. Marinucci, *Phys. Rev. D* **74**, 043524 (2006).
- [24] J. D. McEwen, Y. Wiaux, M. P. Hobson, P. Vanderghelynst, and A. N. Lasenby, arXiv:0704.0626.
- [25] A. Cooray, P. S. Corasaniti, T. Giannantonio, and A. Melchiorri, *Phys. Rev. D* **72**, 023514 (2005).
- [26] P. S. Corasaniti, T. Giannantonio, and A. Melchiorri, *Phys. Rev. D* **71**, 123521 (2005).
- [27] E. Gaztanaga, M. Manera, and T. Multamaki, *Mon. Not. R. Astron. Soc.* **365**, 171 (2006).
- [28] T. Giannantonio and A. Melchiorri, *Classical Quantum Gravity* **23**, 4125 (2006).
- [29] M. J. Rees and D. W. Sciama, *Nature (London)* **217**, 511 (1968).
- [30] W. Hu and R. Scranton, *Phys. Rev. D* **70**, 123002 (2004).
- [31] A. Kinkhabwala and M. Kamionkowski, *Phys. Rev. Lett.* **82**, 4172 (1999).
- [32] M. Blanton, R. Cen, J. P. Ostriker, and M. A. Strauss, *Astrophys. J.* **522**, 590 (1999).
- [33] W. J. Percival *et al.*, *Astrophys. J.* **657**, 645 (2007).
- [34] U. Seljak and M. Zaldarriaga, *Astrophys. J.* **469**, 437 (1996).
- [35] N. Afshordi, *Phys. Rev. D* **70**, 083536 (2004).
- [36] L. Pogosian, P. S. Corasaniti, C. Stephan-Otto, R. Crittenden, and R. Nichol, *Phys. Rev. D* **72**, 103519 (2005).
- [37] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, *Astrophys. J.* **622**, 759 (2005).
- [38] A. Cabre, P. Fosalba, E. Gaztanaga, and M. Manera, arXiv:astro-ph/0701393.
- [39] H. V. Peiris and D. N. Spergel, *Astrophys. J.* **540**, 605 (2000).
- [40] T. H. Jarrett, T. Chester, R. Cutri, S. Schneider, M. Skrutskie, and J. P. Huchra, *Astron. J.* **119**, 2498 (2000).
- [41] D. J. Schlegel, D. P. Finkbeiner, and M. Davis, *Astrophys. J.* **500**, 525 (1998).
- [42] J. K. Adelman-McCarthy *et al.* (SDSS Collaboration), *Astrophys. J. Suppl. Ser.* **175**, 297 (2008).
- [43] D. G. York *et al.* (SDSS Collaboration), *Astron. J.* **120**, 1579 (2000).
- [44] A. Collister *et al.*, *Mon. Not. R. Astron. Soc.* **375**, 68 (2007).
- [45] C. Blake, A. Collister, S. Bridle, and O. Lahav, arXiv:astro-ph/0605303.
- [46] While the DR4 MegaZ data is public, the DR6 data used here was provided thanks to F. Abdalla, C. Blake, and O. Lahav (personal communication).
- [47] S. P. Boughn and R. G. Crittenden, *Phys. Rev. Lett.* **88**, 021302 (2001).
- [48] J. S. Dunlop and J. A. Peacock, *Mon. Not. R. Astron. Soc.* **247**, 19 (1990).
- [49] E. A. Boldt, *Phys. Rep.* **146**, 215 (1987).
- [50] S. P. Boughn, R. G. Crittenden, and N. G. Turok, *New Astron. Rev.* **3**, 275 (1998).
- [51] S. P. Boughn, R. G. Crittenden, and G. P. Koehrsen, *Astrophys. J.* **580**, 672 (2002).
- [52] S. Boughn and R. Crittenden, *Astrophys. J.* **612**, 647 (2004).
- [53] G. T. Richards *et al.* (SDSS Collaboration), *Astrophys. J. Suppl. Ser.* **155**, 257 (2004).
- [54] G. T. Richards *et al.* (unpublished).
- [55] A. D. Myers *et al.*, *Astrophys. J.* **638**, 622 (2006).
- [56] SDSS Collaboration, *Astrophys. J. Suppl. Ser.* **172**, 634 (2007).
- [57] A. Blanchard, M. Douspis, M. Rowan-Robinson, and S. Sarkar, *Astron. Astrophys.* **412**, 35 (2003).
- [58] A. Blanchard, M. Douspis, M. Rowan-Robinson, and S. Sarkar, *Astron. Astrophys.* **449**, 925 (2006).
- [59] Y. Wang and P. Mukherjee, *Astrophys. J.* **650**, 1 (2006).
- [60] W. L. Freedman *et al.*, *Astrophys. J.* **553**, 47 (2001).
- [61] <http://www.dsg.port.ac.uk/~giannant/iswdata.php>.
- [62] S. Ho, C. M. Hirata, N. Padmanabhan, U. Seljak, and N. Bahcall, arXiv:0801.0642.