Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **434,** 3590–3596 (2013)
Advance Access publication 2013 August 1

doi:10.1093/mnras/stt1287

Large-angle cosmic microwave background suppression and polarization predictions

Craig J. Copi,[1⋆] Dragan Huterer,[2⋆] Dominik J. Schwarz[3⋆] and Glenn D. Starkman[1,4⋆]

[1]*CERCA/Department of Physics/ISO, Case Western Reserve University, Cleveland, OH 44106-7079, USA*
[2]*Department of Physics, University of Michigan, 450 Church St, Ann Arbor, MI 48109-1040, USA*
[3]*Fakultät für Physik, Universität Bielefeld, Postfach 100131, 33501 Bielefeld, Germany*
[4]*Physics Department, Theory Unit, CERN, CH-1211 Genève 23, Switzerland*

Accepted 2013 July 11. Received 2013 June 27; in original form 2013 May 6

ABSTRACT

The anomalous lack of large-angle temperature correlations has been a surprising feature of the cosmic microwave background (CMB) since first observed by *COBE-DMR* and subsequently confirmed and strengthened by the *Wilkinson Microwave Anisotropy Probe*. This anomaly may point to the need for modifications of the standard model of cosmology or may indicate that our Universe is a rare statistical fluctuation within that model. Further observations of the temperature auto-correlation function will not elucidate the issue; sufficiently high-precision statistical observations already exist. Instead, alternative probes are required. In this work, we explore the expectations for forthcoming polarization observations. We define a prescription to test the hypothesis that the large-angle CMB temperature perturbations in our Universe represent a rare statistical fluctuation within the standard cosmological model. These tests are based on the temperature-Q Stokes parameter correlation. Unfortunately, these tests cannot be expected to be definitive. However, we do show that if this TQ-correlation is observed to be sufficiently large over an appropriately chosen angular range, then the hypothesis can be rejected at a high confidence level. We quantify these statements and optimize the statistics we have constructed to apply to the anticipated polarization data. We find that we can construct a statistic that has a 25 per cent chance of excluding the hypothesis that we live in a rare realization of Λ cold dark matter at the 99.9 per cent confidence level.

Key words: cosmic background radiation – large-scale structure of Universe.

1 INTRODUCTION

In the two decades since the *Cosmic Background Explorer* (*COBE*) first detected the primordial fluctuations in the cosmic microwave background (CMB) temperature (Wright et al. 1992), and perhaps even more so in the past decade over which the *Wilkinson Microwave Anisotropy Probe* (*WMAP*) has provided ever more accurate full-sky maps of those fluctuations (see Komatsu et al. 2011, for example), the CMB has become a keystone in the remarkable transition of cosmology from a qualitative to a precision science.

An important element of the role of the CMB in precision cosmology has been that the canonical theory of cosmology, inflationary Λ cold dark matter (ΛCDM), makes clear predictions for the statistical properties of the spherical harmonic coefficients of the temperature fluctuations,

$$a_{\ell m} \equiv \int Y^*_{\ell m}(\theta,\phi) T(\theta,\phi)\, \mathrm{d}(\cos\theta)\, \mathrm{d}\phi, \tag{1}$$

which are predicted to be statistically isotropic realizations of independent Gaussian random variables of zero mean and with variance C_ℓ depending only on ℓ,

$$C_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2. \tag{2}$$

In modern discussions of the CMB, the two-point angular power spectrum, embodied in these C_ℓ, plays a central role and is the source of the remarkable precision of the cosmological parameters (Komatsu et al. 2011).

Before the *COBE* era, it was the two-point angular correlation function of the fluctuations,

$$C(\theta) \equiv \overline{T(\hat{\boldsymbol{n}}_1) T(\hat{\boldsymbol{n}}_2)}|_{\hat{\boldsymbol{n}}_1\cdot\hat{\boldsymbol{n}}_2=\cos\theta}, \tag{3}$$

rather than the angular power spectrum that was of primary interest to astronomers. Statistically, the two-point angular correlation function is an ensemble average but, in practice, this must be replaced by an average over pairs of points separated by an angle θ, as denoted by the bar over the expression. In fact, the *COBE* differential microwave radiometer (*COBE-DMR*) did report $C(\theta)$,

⋆E-mail: cjc5@cwru.edu (CJC); huterer@umich.edu (DH); dschwarz@physik.uni-bielefeld.de (DJS); glenn.starkman@case.edu (GDS)

though only in their final, four-year paper (Bennett et al. 1996). When extracted from a full-sky map, both $C(\theta)$ and the C_ℓ contain the same information, albeit in different forms. The same is true for a function and its Fourier transform; signals are typically most easily seen in one or the other forms but not both. In the case of the CMB, $C(\theta)$ and the C_ℓ are related by a Legendre series. The C_ℓ most easily show the small angular scale behaviour, microphysics at last scattering, whereas the $C(\theta)$ most easily shows the large angular scale behaviour.

As observed by the *COBE-DMR*, $C(\theta)$ had an unexpected property – it was consistent with zero for angular separations between approximately 60° and 160°. This was duly noted at the time but mainly remembered today as a low quadrupole. The *WMAP* team confirmed the *COBE-DMR* observation of a lack of large-angle correlation with significantly smaller error bars. In their initial, one-year release Spergel et al. (2003) phrased the anomaly in terms of a statistic

$$S_{1/2} \equiv \int_{-1}^{1/2} [C(\theta)]^2 \, \mathrm{d}(\cos\theta). \tag{4}$$

In the best-fitting ΛCDM model, the expected value of this statistic is approximately $50\,000\,(\mu\mathrm{K})^4$, whereas the observed value is approximately $8500\,(\mu\mathrm{K})^4$ on the full sky, e.g. from the *WMAP* independent linear combination (ILC) map,[1] with a p-value of approximately 0.05. Even more striking is that if one considers only the part of the sky outside a conservative Galaxy cut, then $S_{1/2} \simeq 1000$–$1150\,(\mu\mathrm{K})^4$ and is only $\sim 1300\,(\mu\mathrm{K})^4$ in each of the V and W frequency bands, which are expected to be dominated by the CMB signal. The cut-sky $S_{1/2}$ has a p-value of about 2.5×10^{-4}, depending on the precise map (Copi et al. 2009).

We have argued that such absence of the two-point angular correlation is unlikely to result solely from a small quadrupole, or even a small quadrupole and octopole, and that it instead requires a 'conspiracy' among the first several multipoles (Copi et al. 2009). Such covariance among the C_ℓ is likely contrary to the fundamental prediction of the canonical cosmological model that the $a_{\ell m}$ underlying the C_ℓ are independent Gaussian random variables with variances depending only on ℓ. This is one of a number of large-scale anomalies that suggest that modifications of the standard models are required on large angular scales (see Bennett et al. 2011; Copi et al. 2010, and references therein for further details).

One possible explanation of the absence of large-angle correlations among the CMB temperature fluctuations is that it is merely a statistical fluke. In this paper, we explore the consequences of this hypothesis. In particular, since the fluctuations in the CMB temperature and in its polarization arise largely from the same source – the gravitational potential – one might have hoped that a small temperature–temperature (*TT*) correlation function on large angular scales would predict a similarly small cross-correlation between CMB temperature and CMB Stokes parameter Q, or in the polarization–polarization (*QQ*) correlation.

Unfortunately, as we shall see, the connection between temperature and polarization fluctuations is too weak for a general definitive test of the origin of the vanishing correlation function. However, we do find that if the $S_{1/2}$ is small because of a statistical fluke within ΛCDM cosmology, then the cross-correlation between temperature and polarization is unlikely to be large on large angular scales. Therefore, were we to infer a large value of this cross-correlation from future data, that would be evidence against a statistical fluke as an explanation of the vanishing *TT* correlation.

In this paper, we provide a prescription to follow in order to test this hypothesis. In Section 2, we describe the construction of an ensemble of realizations of ΛCDM that are constrained to resemble our observed Universe in the properties of their *TT* angular power spectrum and full-sky and cut-sky two-point angular correlation functions. In Section 3, we construct statistics that, like $S_{1/2}$ for the *TT* correlations, can be used to quantify the smallness of the correlations between temperature and polarization fluctuations. Section 4 contains a discussion of the results of applying the temperature–polarization cross-correlation statistics to the ensemble of constrained realizations and looks forward to what might be learned by applying them to future polarization data. Given the constructed realizations, the values of the statistics are fixed and the optimal application of them to polarization data can be determined as discussed in this section. This optimization is *independent* of polarization observations. Finally, Section 5 contains the conclusions.

2 CONSTRAINED REALIZATIONS

To study the signature of the lack of large-angle correlations in the *WMAP* temperature data on upcoming polarization measurements, such as from *Planck*, we require realizations of ΛCDM consistent with these large-angle results. For this purpose, we have generated 300 000 such realizations as follows.

(i) In the standard, ΛCDM model, our Universe is a realization from an ensemble, the width of which (the cosmic variance) for low-ℓ information is quite large. However, once measured our realization can and has been precisely determined. In the work reported here this information is given in the *WMAP* reported C_ℓ^{TT}. We are interested in producing realizations of *our* Universe as represented by the *WMAP* observations, *not* realizations of the full ΛCDM model. For this purpose, we treat the observational errors in the *WMAP* reported C_ℓ^{TT} as Gaussian distributed and generate realizations accordingly. Thus, we generate random C_ℓ^{TT} from Gaussian distributions centred on the *WMAP*-reported values. This produces a power spectrum consistent with that reported by *WMAP*. Again it is important to stress that this is a power spectrum consistent with the observation of our particular realization of the Universe as measured by *WMAP*, not a general realization of the best-fitting ΛCDM model. For this reason, cosmic variance is not relevant nor do we generate C_ℓ from a χ^2 distribution. It is also true that on a partial sky the C_ℓ are slightly correlated. To correct for this we actually use the Fisher matrix from the *WMAP* likelihood code *without* the contribution from cosmic variance in drawing these C_ℓ^{TT}. In practice, this is a small correction but has been included for completeness.

(ii) The C_ℓ^{TT} generated in the previous step contain the statistical information about the power in each mode consistent with the *WMAP* observations. For further analysis we need a map, not just the power spectrum. A map is a particular realization of this power spectrum. As in ΛCDM we assume that the modes in the map have random phases. In practice, this means we choose the $a_{\ell m}^T$ randomly on the 2ℓ-sphere such that

$$\frac{1}{2\ell+1} \sum_{\ell=-m}^{m} \left|a_{\ell m}^T\right|^2 = C_\ell^{TT}, \tag{5}$$

where the C_ℓ^{TT} in this expression are *exactly* the values generated from the previous step. By this construction the resulting sky realization is guaranteed to have a $S_{1/2}$ consistent with the small value in the full-sky *WMAP* ILC map. For each set of C_ℓ^{TT} from the

[1] The ILC map and all data from the *WMAP* mission used in this work are freely available from http://lambda.gsfc.nasa.gov/.

previous step, we generate a single complete sky realization, i.e. a single map.

(iii) To further be consistent with *WMAP* observations the $S_{1/2}$ on the cut sky must also be small. For our realizations we require $S_{1/2}^{\rm cut} \leq 1292.6\,(\mu{\rm K})^4$, the value from the *WMAP* seven year, KQ75y7 masked ILC map. This $S_{1/2}$ value is calculated for a realization by first constructing a map at NSIDE $= 64$ from the $a_{\ell m}^T$ generated above. The pseudo-C_ℓ are extracted from this map based on the region outside the KQ75y7 mask using SPICE (Chon et al. 2004). Finally, $S_{1/2}$ is calculated using these C_ℓ up to $\ell_{\rm max} = 100$.

(iv) For temperature realizations that satisfy the cut sky constraint, we also generate realizations of the $a_{\ell m}^E$. See Appendix A for a review of the process.

The construction of a set of constrained realizations is the basis for the prescription we are describing. It will next provide predictions for the expectations from the observations of the CMB polarization.

3 STATISTICS

For the temperature auto-correlation, the $S_{1/2}$ statistic was defined a posteori (Spergel et al. 2003) to be

$$S_{1/2} \equiv \int_{-1}^{1/2} \left[C^{TT}(\theta)\right]^2 {\rm d}(\cos\theta). \tag{6}$$

Inspired by this we define a comparable statistic for $C^{TQ}(\theta)$, the two-point angular correlation function between fluctuations in the temperature and the Stokes parameter Q.

Observable properties of photons can be characterized by the Stokes parameters. For the CMB the relevant quantities are the intensity, conventionally represented by the temperature, T, and the linear polarization given by the Q and U parameters. For the CMB the circular polarization, represented by the V Stokes parameter, is expected to be zero and not considered further. When working in real space the natural correlations to construct are amongst these observables, T, Q and U. These correlations are constructed such that they only depend on the angular separation along the great circle connecting each pairs of point on the sky and are thus rotationally invariant despite the fact that the definition of Q and U depend on the choice of coordinate axes (Kamionkowski, Kosowsky & Stebbins 1997). When working in harmonic space, it is natural to decompose the polarization into 'gradient' and 'curl' modes (Kamionkowski et al. 1997) alternatively called E and B modes (Zaldarriaga & Seljak 1997), which are similarly rotationally invariant quantities. These latter names will be used throughout. Thus, in real space we will work with the TQ two-point angular correlation function, $C^{TQ}(\theta)$, which may be written in terms of the two-point angular power spectrum coefficients, C_ℓ^{TE}.

3.1 S^{TQ} statistic

In the case of polarization a priori the optimal range over which to integrate the correlation function is unknown and will be explored below so we define the general statistic

$$S^{TQ}(\theta_1, \theta_2) \equiv \int_{\cos\theta_2}^{\cos\theta_1} \left[C^{TQ}(\theta)\right]^2 {\rm d}(\cos\theta). \tag{7}$$

As with $S_{1/2}$ we may calculate this easily in terms of the power spectrum coefficients, C_ℓ^{TE}. Using (Kamionkowski et al. 1997)

$$C^{TQ}(\theta) = \sum_{\ell=2}^{\infty} \frac{2\ell+1}{4\pi} \sqrt{\frac{(\ell-2)!}{(\ell+2)!}} C_\ell^{TE} P_\ell^2(\cos\theta), \tag{8}$$

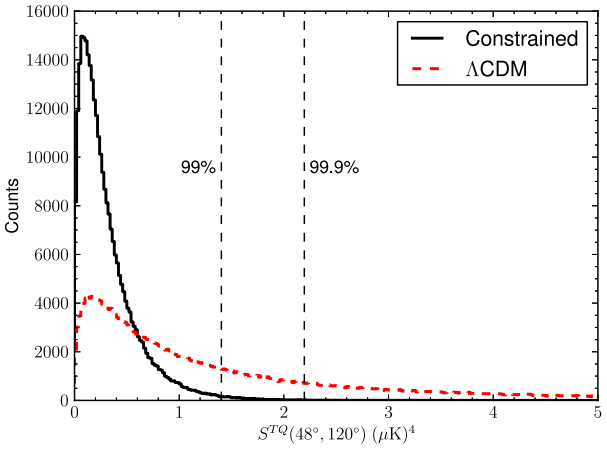

Figure 1. Example histogram of the S^{TQ} statistic, defined in equation (7), for constrained (solid, black line) and ΛCDM (dashed, red line) realizations. We note that the constrained realizations are more sharply peaked at low S^{TQ} than ΛCDM which, though peaked at approximately the same value, has a long tail. The dashed, vertical lines represent the values of S^{TQ} for which 99 and 99.9 per cent, respectively, of the constrained realizations have smaller values.

we may show that

$$S^{TQ}(\theta_1, \theta_2) = \sum_{\ell,\ell'} C_\ell^{TE} I_{\ell,\ell'}^{TQ}(\theta_1, \theta_2) C_{\ell'}^{TE}, \tag{9}$$

where $I_{\ell,\ell'}^{TQ}(\theta_1, \theta_2)$ are components of a known matrix calculated in Appendix B. A histogram of the S^{TQ} statistic for a particular choice of θ_1 and θ_2 is shown in Fig. 1 comparing the constrained realizations to ΛCDM.

3.2 s^{TQ} Statistic

Motivated solely by its simplicity and ease of computation, we also define a new statistic which is linear, rather than quadratic, in the TQ correlation function

$$s^{TQ}(\theta_1, \theta_2) \equiv \int_{\cos\theta_2}^{\cos\theta_1} C^{TQ}(\theta)\, {\rm d}(\cos\theta). \tag{10}$$

As with $S^{TQ}(\theta_1, \theta_2)$ we may calculate this easily in terms of the C_ℓ^{TE},

$$s^{TQ}(\theta_1, \theta_2) = \sum_{\ell=2}^{\infty} C_\ell^{TE} i_\ell^{TQ}(\theta_1, \theta_2), \tag{11}$$

where $i_\ell^{TQ}(\theta_1, \theta_2)$ are the components of a known vector calculated in Appendix B. A histogram of the s^{TQ} statistic for a particular choice of θ_1 and θ_2 is shown in Fig. 2 comparing the constrained realizations to ΛCDM.

4 RESULTS

The S^{TQ} and s^{TQ} statistics defined above have been calculated for the constrained realizations discussed in Section 2 and for a comparable number of realizations of ΛCDM. In both cases, these have been calculated from maps produced at NSIDE $= 64$. Data from the temperature map outside the KQ75y7 mask and from the seven-year polarization analysis mask, both provided by *WMAP*, have been used. As shown in Figs 1 and 2, the constrained and ΛCDM

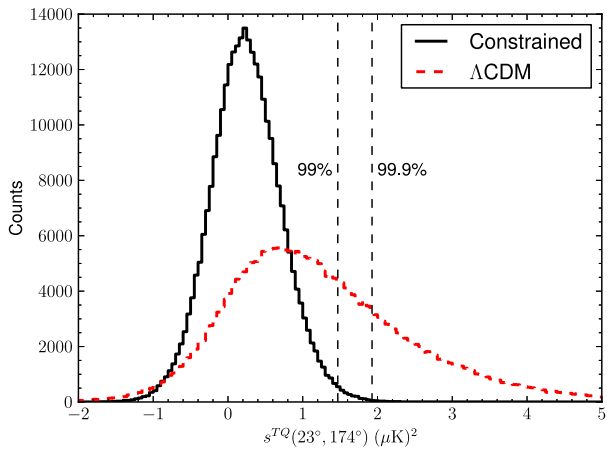

Figure 2. Example histogram of the s^{TQ} statistic, defined in equation (10), for constrained (solid, black line) and ΛCDM (dashed, red line) realizations. We note that the constrained realizations are more sharply peaked near zero than ΛCDM which, though also peaked near zero, has a long tail particularly to large, positive values. The dashed, vertical lines represent the values of s^{TQ} for which 99 and 99.9 per cent, respectively, of the constrained realizations have smaller values.

realizations have most likely values for the S^{TQ} and s^{TQ} statistics at nearly the same value. However, as we also see ΛCDM predicts much broader distributions for the two statistics. In particular, there is a significant probability of producing values larger than the constrained realizations. This provides a means of testing the hypothesis that our Universe is just a rare realization of ΛCDM. This results in a simple but not definitive test.

Consider the case of the S^{TQ} statistic as represented in Fig. 1. For the constrained realizations 99 per cent of them have $S^{TQ}(48^\circ, 120^\circ) \leq 1.403\,(\mu\mathrm{K})^4$ and 99.9 per cent have $S^{TQ}(48^\circ, 120^\circ) \leq 2.195\,(\mu\mathrm{K})^4$. Unconstrained ΛCDM (with the best-fitting values of cosmological parameters) randomly generates realizations with values larger than these 38.6 and 25.6 per cent of the time, respectively. If observations of the polarization show our Universe to have a $S^{TQ}(48^\circ, 120^\circ)$ value larger than these values, then we can reject the random ΛCDM realization hypothesis at the appropriate confidence level. Alternatively, if the $S^{TQ}(48^\circ, 120^\circ)$ value is smaller, then no definitive statement can be made; the polarization fluctuations would be consistent with the hypothesis that we live in a rare ΛCDM realization but do nothing to advance that hypothesis. This is the main point of the paper.

Similar statements may be made about the s^{TQ} statistic shown in Fig. 2. It provides similar information as S^{TQ}.

We still have freedom to choose the optimal range of angles over which to evaluate the statistics. We define optimal to mean the maximum discriminatory power between the distribution of the statistic in the constrained versus ΛCDM realizations. For a given per cent cut-off from the constrained realizations, we wish to find the angle range $[\theta_1, \theta_2]$ that has the maximum fraction of ΛCDM realizations above this value.

The results of such a study are shown in Figs 3 and 4 for the $S^{TQ}(\theta_1, \theta_2)$ statistic and in Figs 5 and 6 for the $s^{TQ}(\theta_1, \theta_2)$ statistic. The statistics are non-zero only up to the diagonal $\theta_1 = \theta_2$, shown as the dashed, black line in the figures, but not along it. For this reason, the contours are truncated at the diagonal. They have been made symmetric in θ_1 and θ_2 (by taking $|S^{TQ}|$ and $|s^{TQ}|$) so the results are shown as identical when reflected through the diagonal. The optimal

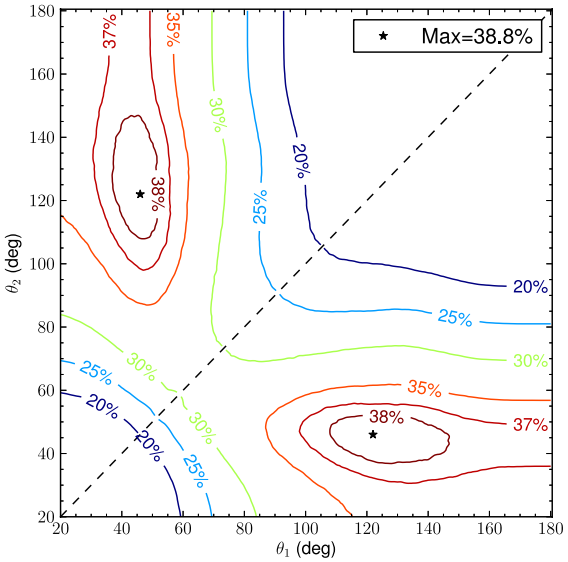

Figure 3. Contours for the fraction of ΛCDM realizations above the 99 per cent value of the constrained realizations from the S^{TQ} statistic, defined in equation (7). In the optimal case 38.8 per cent of the ΛCDM realizations have a larger value. In this figure, the results are not defined along the diagonal (black, dashed line) and have been made symmetric about it by taking the absolute value of the statistic.

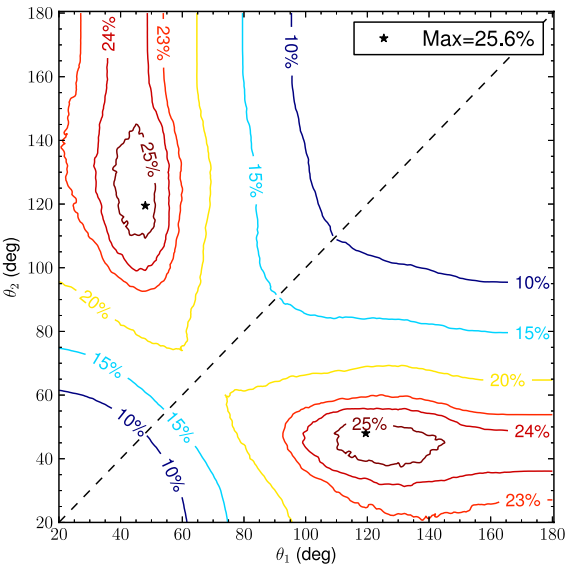

Figure 4. Same as Fig. 3 now for the 99.9 per cent case. Here in the optimal case 25.6 per cent of the ΛCDM realizations have a larger value and the full histograms are shown in Fig. 1.

ranges and fractions of ΛCDM realizations are listed in Table 1. We note that the optimal surfaces represented by the contours seen in the figures are relatively broad, at least in one direction. Due to this the values of θ_1 and θ_2 in a neighbourhood of those listed in Table 1 can be employed with nearly the same efficacy.

5 CONCLUSIONS

The absence of two-point angular correlations on large angular scales in the CMB temperature data is, by now, well established. It was first measured by *COBE-DMR*, but became more significant

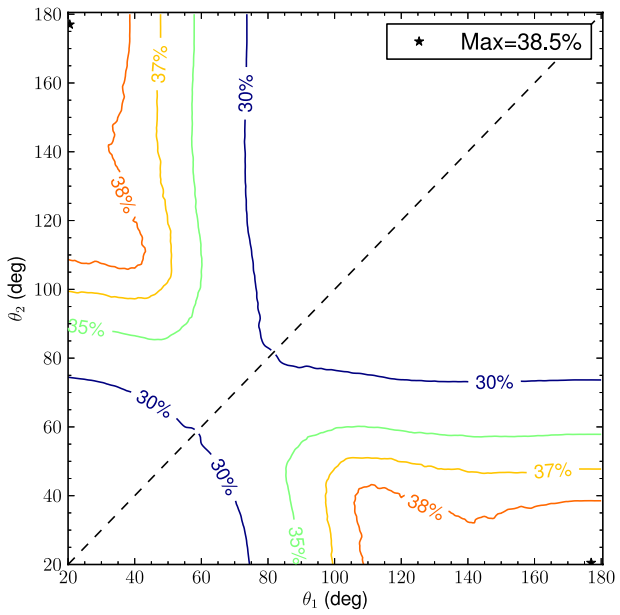

Figure 5. Same as Fig. 3 but now for the s^{TQ} statistic, defined in equation (10). Here in the optimal case 38.5 per cent of the ΛCDM realizations have a larger value at the 99 per cent level.

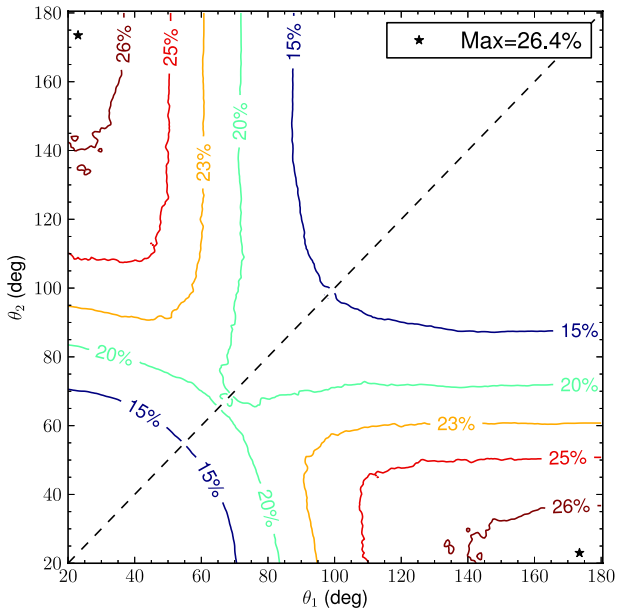

Figure 6. Same as Fig. 5 now for the 99.9 per cent case. Here in the optimal case 26.4 per cent of the ΛCDM realizations have a larger value and the full histograms are shown in Fig. 2.

in the *WMAP* temperature maps. This absence of correlation is difficult to accommodate within the standard cosmological model, especially since it seems to imply covariance among low-ℓ multipoles of the CMB. A simple explanation that has been proffered is that we just live in a rare realization of ΛCDM that happens to have a lack of large-angle *TT* correlations. If so, one might hope that constraining ΛCDM realizations to have low *TT* correlations at large angles would have observable consequences for other correlation functions, such as *TQ*. These would be the basis for an observational test of this statistical fluke hypothesis.

Table 1. Optimal angle ranges for the S^{TQ} statistic (7) and s^{TQ} statistic (10). The optimal ranges are determined by finding the maximum fraction of ΛCDM realizations with the appropriate statistic above the 99 or 99.9 per cent level of the constrained realizations. The histograms for the optimal 99.9 per cent ranges are shown in Figs 1 and 2. The full contours are shown in Figs 3–6.

Statistic	C.L. (per cent)	θ_1 (deg)	θ_2 (deg)	Fraction (per cent)
S^{TQ}	99	46	122	38.8
	99.9	48	120	25.6
s^{TQ}	99	20	177	38.5
	99.9	23	174	26.4

In this paper, we have discussed a prescription to follow in order to test this hypothesis of our Universe being a statistical fluke. The prescription may be simply stated as follows.

(i) Construct realizations of our Universe consistent with the observed temperature fluctuations. This means construct sets of $a^T_{\ell m}$ and $a^E_{\ell m}$ consistent with the observed C^{TT}_ℓ and full and cut sky $S_{1/2}$ as discussed in Section 2.

(ii) Apply the S^{TQ} and s^{TQ} statistics as defined in Section 3 to these constrained realizations.

(iii) Also apply the S^{TQ} and s^{TQ} statistics to a comparable number of best-fitting ΛCDM realizations and use these to find the optimal ranges $[\theta_1, \theta_2]$ for each statistic. Optimal here means that the maximum fraction of ΛCDM realizations fall above the value at some confidence level in the constrained realization, e.g. the 99 or 99.9 per cent level.

(iv) Given the optimal ranges from the previous step now apply these particular cases to the observed polarization signal. If the observations produce values for these statistics larger than that expected from the constrained realizations, then the statistical fluke hypothesis can be rejected at the appropriate confidence level. Alternatively, if the values are smaller, then the hypothesis remains consistent but unproven.

We further note that the optimization in this prescription is independent of the polarization observations, or, in fact, whether the polarization has been observed or not.

Our work is, in spirit, related to Dvorkin, Peiris & Hu (2008). While they consider observables in the polarization signal for 'models' of three dimensional primordial power modulation that might explain the breaking of statistical isotropy in the temperature field, we predict the polarization statistics starting directly from realizations of ΛCDM models that are constrained to show the suppressed *TT* correlation at large angular scales.

In the work reported here we have generated realizations and performed the optimization based on the *WMAP* seven-year data release. The prescription could be applied to the *WMAP* nine-year data release and the results are not expected to differ significantly. We have also said nothing about applying the statistics to the *WMAP* reported polarization observations. Unfortunately, the signal-to-noise ratio in the polarization observations is not yet sufficient to make meaningful statements. To see this, using the *WMAP* nine-year reported C^{TE}_ℓ (Hinshaw et al. 2012), we find for the optimal ranges

$$S^{TQ}_{WMAP}(48^\circ, 120^\circ) = (1.0 \pm 0.8)\,(\mu\mathrm{K})^4 \qquad (12)$$

and

$$s^{TQ}_{WMAP}(23^\circ, 174^\circ) = (0.8 \pm 0.8)\,(\mu\mathrm{K})^2. \quad (13)$$

Here the error bars are crude estimates assuming that the reported C_ℓ^{TE} are statistically independent and the noise is Gaussian. These assumptions are not justified and a more careful assessment could be performed using the Fisher matrix. However, given the large estimated errors such an assessment is not warranted.

We have shown that the prescription described in this work is not a definitive test of the statistical fluke hypothesis for our Universe. Nevertheless, by carefully optimizing the statistical measure of large-angle *TQ* correlations, we were able to demonstrate that once good *TQ* correlation data are available there is a reasonable probability (over 25 per cent) to reject the statistical fluke hypothesis at the 99.9 per cent confidence level. *WMAP* data are not up to this task; however, *Planck* data should be.

ACKNOWLEDGEMENTS

GDS and CJC are supported by a grant from the US Department of Energy to the Particle Astrophysics Theory Group at CWRU. DH has been supported by the DOE, NASA and NSF. DJS is supported by the DFG grant RTG 1620 'Models of gravity'. GDS thanks the Theory Unit at CERN for their hospitality. This work made extensive use of the HEALPIX package (Górski et al. 2005). The numerical simulations were performed on the facilities provided by the Case ITS High Performance Computing Cluster.

REFERENCES

Bennett C. L. et al., 1996, ApJ, 464, L1
Bennett C. L. et al., 2011, ApJS, 192, 17
Chon G., Challinor A., Prunet S., Hivon E., Szapudi I., 2004, MNRAS, 350, 914
Copi C. J., Huterer D., Schwarz D. J., Starkman G. D., 2009, MNRAS, 399, 295
Copi C. J., Huterer D., Schwarz D. J., Starkman G. D., 2010, Adv. Astron., 2010, 78
Dvorkin C., Peiris H. V., Hu W., 2008, Phys. Rev. D, 77, 063008
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Hinshaw G. et al., 2012, preprint (arXiv:1212.5226)
Kamionkowski M., Kosowsky A., Stebbins A., 1997, Phys. Rev. D, 55, 7368
Komatsu E. et al., 2011, ApJS, 192, 18
Spergel D. N. et al., 2003, ApJS, 148, 175
Wright E. L. et al., 1992, ApJ, 396, L13
Zaldarriaga M., Seljak U., 1997, Phys. Rev. D, 55, 1830

APPENDIX A: POLARIZATION GAUSSIAN RANDOM REALISATIONS

The generation of correlated Gaussian random variables is a well-known topic. For use in the CMB this is implemented in HEALPIX (Górski et al. 2005), for example. Here we review the details relevant for the generation of our constrained realisations.

In ΛCDM the temperature and *E*-type polarization are correlated as encoded in the power spectrum coefficients C_ℓ^{TT}, C_ℓ^{TE} and C_ℓ^{EE} from the best-fitting ΛCDM model. Working in the real spherical harmonic basis we may generate the spherical harmonic coefficients as

$$a_j^T = \sqrt{C_\ell^{TT}}\,\zeta_1, \quad (A1)$$

$$a_j^E = \frac{C_\ell^{TE}}{\sqrt{C_\ell^{TT}}}\zeta_1 + \sqrt{C_\ell^{EE} - \frac{(C_\ell^{TE})^2}{C_\ell^{TT}}}\,\zeta_2, \quad (A2)$$

where ζ_1 and ζ_2 are Gaussian random variables drawn from a distribution with zero mean and unit variance and the index j refers to the pair of indices (ℓ, m). To be precise, j takes the values 0 to 2ℓ and the complex coefficients are constructed as

$$a_{\ell m}^T = \begin{cases} a_0^T, & m = 0 \\ \frac{1}{\sqrt{2}}\left(a_{2m-1}^T + \mathrm{i}\, a_{2m}^T\right), & m > 0 \end{cases}. \quad (A3)$$

For our purposes we need to generate constrained realizations of ΛCDM so the above procedure must be modified. The steps discussed in Section 2 lead to the generation of constrained $a_{\ell m}^T$. In other words, we have determined a_j^T which by equation (A1) means we have also determined ζ_1. That is, instead of choosing ζ_1 as a Gaussian random variable we have used observational constraints to determine its value and find it by inverting that equation. Since the temperature and *E*-type polarization are correlated, this constrained temperature realization affects a_j^E. The real and imaginary components of $a_{\ell m}^E$ may now be generated from (A2) as

$$a_j^E = \frac{C_\ell^{TE}}{C_\ell^{TT}}a_j^T + \sqrt{C_\ell^{EE} - \frac{(C_\ell^{TE})^2}{C_\ell^{TT}}}\,\zeta_2, \quad (A4)$$

where ζ_2 is still to be chosen as a Gaussian random variable.

APPENDIX B: DERIVATION OF STATISTICS FORMULAS

B1 $S^{TQ}(\theta_1, \theta_2)$

We wish to evaluate $S^{TQ}(\theta_1, \theta_2)$ as discussed in the text (7). Consider the simpler case

$$\begin{aligned} S^{TQ}(x) &\equiv \int_{-1}^{x} \left[C^{TQ}(\theta)\right]^2 \mathrm{d}(\cos\theta) \\ &= \sum_{\ell,\ell'} \frac{(2\ell+1)(2\ell'+1)}{(4\pi)^2}\sqrt{\frac{(\ell-2)!(\ell'-2)!}{(\ell+2)!(\ell'+2)!}}C_\ell^{TE}C_{\ell'}^{TE} \\ &\quad \times \int_{-1}^{x} P_\ell^2(\cos\theta)P_{\ell'}^2(\cos\theta)\,\mathrm{d}(\cos\theta). \end{aligned} \quad (B1)$$

To evaluate this expression we need to perform the integral of two associated Legendre functions, P_ℓ^m, of order $m = 2$,

$$\tilde{I}^{TQ}_{\ell,\ell'}(x) \equiv \int_{-1}^{x} P_\ell^2(x)P_{\ell'}^2(x)\,\mathrm{d}x. \quad (B2)$$

Note that this integral is only defined for $\ell, \ell' \geq 2$.

For $\ell \neq \ell'$, we may proceed by directly integrating the associated Legendre differential equation to find

$$\begin{aligned} \tilde{I}^{TQ}_{\ell,\ell'}(x) &= \frac{1-x^2}{\ell(\ell+1) - \ell'(\ell'+1)} \\ &\quad \times \left[P_\ell^2(x)\frac{\mathrm{d}P_{\ell'}^2(x)}{\mathrm{d}x} - P_{\ell'}^2(x)\frac{\mathrm{d}P_\ell^2(x)}{\mathrm{d}x}\right]. \end{aligned} \quad (B3)$$

For $\ell = \ell'$ more care is required. Starting from the Rodriguez formula

$$P_\ell^2(x) = (1-x^2)\frac{\mathrm{d}^2 P_\ell(x)}{\mathrm{d}x^2} \quad (B4)$$

and the recursion relation

$$x\frac{\mathrm{d}P_\ell(x)}{\mathrm{d}x} = \frac{\mathrm{d}P_{\ell-1}(x)}{\mathrm{d}x} + \ell P_\ell(x), \tag{B5}$$

we can show that

$$\tilde{I}^{TQ}_{\ell,\ell}(x) = 4J^{(2)}_{\ell-1}(x) - 4\ell(\ell-1)J^{(1)}_\ell(x) + \ell^2(\ell-1)^2\tilde{I}_{\ell,\ell}(x). \tag{B6}$$

Here, $\tilde{I}_{\ell,\ell}(x)$ is the equivalent integral over Legendre polynomials encountered in the definition of $S_{1/2}$; see appendix A of Copi et al. (2009) for details. The remaining quantities, $J^{(1)}_\ell(x)$ and $J^{(2)}_\ell(x)$, are calculated through integration by parts and use of the recursion relation to find

$$J^{(1)}_\ell(x) = P_{\ell-1}(x)P_\ell(x) + \frac{1}{2}\left\{1 - x\left[P_{\ell-1}(x)\right]^2 - (2\ell-1)\tilde{I}_{\ell-1,\ell-1}(x)\right\}, \tag{B7}$$

and

$$J^{(2)}_\ell(x) = J^{(2)}_{\ell-1}(x) + \ell\left[P_{\ell-1}(x)P_\ell(x) + 1\right]. \tag{B8}$$

Note that $J^{(2)}_\ell(x)$ is defined recursively. We can directly show that $J^{(2)}_0(x) = 0$.

With these expressions for the integrals we may write

$$S^{TQ}(\theta_1,\theta_2) = \sum_{\ell,\ell'} C^{TE}_\ell I^{TQ}_{\ell,\ell'}(\theta_1,\theta_2)C^{TE}_{\ell'}, \tag{B9}$$

where

$$I^{TQ}_{\ell,\ell'}(\theta_1,\theta_2) = \frac{(2\ell+1)(2\ell'+1)}{(4\pi)^2}\sqrt{\frac{(\ell-2)!(\ell'-2)!}{(\ell+2)!(\ell'+2)!}} \times \left[\tilde{I}^{TQ}_{\ell,\ell'}(\cos\theta_1) - \tilde{I}^{TQ}_{\ell,\ell'}(\cos\theta_2)\right]. \tag{B10}$$

This matrix may be precomputed for rapid evaluation of $S^{TQ}(\theta_1,\theta_2)$.

B2 $s^{TQ}(\theta_1,\theta_2)$

We wish to evaluate $s^{TQ}(\theta_1,\theta_2)$ as discussed in the text (10). We proceed as above and consider the simpler case

$$s^{TQ}(x) \equiv \int_{-1}^{x} C^{TQ}(\theta)\,\mathrm{d}(\cos\theta) = \sum_{\ell=2}^{\infty}\frac{(2\ell+1)}{4\pi}\sqrt{\frac{(\ell-2)!}{(\ell+2)!}}C^{TE}_\ell \times \int_{-1}^{x} P^2_\ell(\cos\theta)\,\mathrm{d}(\cos\theta). \tag{B11}$$

To evaluate this expression we need to perform the integral

$$\tilde{\imath}^{TQ}_\ell(x) \equiv \int_{-1}^{x} P^2_\ell(x)\,\mathrm{d}x. \tag{B12}$$

This integral is straightforward to evaluate starting from the Rodriguez formula (B4) and integrating by parts to find

$$\tilde{\imath}^{TQ}_\ell(x) = \ell P_{\ell-1}(x) - (\ell-2)xP_\ell(x) + 2(-1)^\ell - \frac{2}{2\ell+1}\left[P_{\ell+1}(x) - P_{\ell-1}(x)\right]. \tag{B13}$$

With this we may write

$$s^{TQ}(\theta_1,\theta_2) = \sum_{\ell=2}^{\infty} C^{TE}_\ell i^{TQ}_\ell(\theta_1,\theta_2), \tag{B14}$$

where

$$i^{TQ}_\ell(\theta_1,\theta_2) = \frac{2\ell+1}{4\pi}\sqrt{\frac{(\ell-2)!}{(\ell+2)!}}\left[\tilde{\imath}^{TQ}_\ell(\cos\theta_1) - \tilde{\imath}^{TQ}_\ell(\cos\theta_2)\right]. \tag{B15}$$

This vector may be precomputed for rapid evaluation of $s^{TQ}(\theta_1,\theta_2)$.

This paper has been typeset from a TEX/LATEX file prepared by the author.