Is cosmology consistent?

Xiaomin Wang[1], Max Tegmark[1] & Matias Zaldarriaga[2]
[1] *Dept. of Physics, Univ. of Pennsylvania, Philadelphia, PA 19104; xiaomin@physics.upenn.edu*
[2] *Dept. of Physics, New York University, 4 Washington Pl., New York, NY 10003*
(May 10 2001. Submitted to Phys. Rev. D.)

We perform a detailed analysis of the latest CMB measurements (including BOOMERaNG, DASI, Maxima and CBI), both alone and jointly with other cosmological data sets involving, e.g., galaxy clustering and the Lyman Alpha Forest. We first address the question of whether the CMB data are internally consistent once calibration and beam uncertainties are taken into account, performing a series of statistical tests. With a few minor caveats, our answer is yes, and we compress all data into a single set of 24 bandpowers with associated covariance matrix and window functions.
We then compute joint constraints on the 11 parameters of the "standard" adiabatic inflationary cosmological model. Out best fit model passes a series of physical consistency checks and agrees with essentially all currently available cosmological data. In addition to sharp constraints on the cosmic matter budget in good agreement with those of the BOOMERaNG, DASI and Maxima teams, we obtain a tau(?) neutrino mass range $0.04 - 4.4$ eV and the sharpest constraints to date on gravity waves which (together with preference for a slight red-tilt) favor "small-field" inflation models.

I. INTRODUCTION

The recent discovery [1–3] of multiple peaks in the cosmic microwave background (CMB) power spectrum marks a major milestone in cosmology. Confirming 1970 predictions of Peebles & Yu [4] and Sunyaev & Zeldovich [5], it greatly bolsters the credibility of the emerging standard model of cosmology, and allows many of its free parameters to be measured with a precision that cosmologists have yet to get accustomed to.

This new precision also offers new opportunities for consistency tests, both for systematic errors that might affect individual measurements and for incorrect assumptions about the underlying physical processes. The goal of the present paper is to carry our these two types of tests.

We begin in Section II by testing measurements of the CMB power spectrum for consistency with a series of quantitative statistical tests, including the effects of calibration and beam uncertainties. Since the customary plot of available measurements has now evolved into a bewildering zoo of over 100 band power measurements which is increasingly difficult to interpret visually because of calibration and beam effects, we perform an essentially lossless data compression of all data into a single set of 24 bandpowers with associated covariance matrix and window functions, useful as a starting point for parameter fitting.

In Section III, we compute quantitative constraints on the 11 parameters of the "standard" adiabatic inflationary cosmological model in a variety of different ways, using data from, *e.g.*, the CMB, galaxy clustering, the Lyman Alpha Forest, Big Bang nucleosynthesis (BBN) and Hubble constant measurements in various combinations. This enables us to identify a number of parameter constraints that are robust and consistent with all data, as well as areas where there is slight tension between data sets pulling in different directions. Although numerous such studies have been performed in the recent literature, *e.g.*, [6–17], the dramatically improved precision allowed by new BOOMERaNG [1], DASI [2], Maxima [3] and CBI [18] data makes it worthwhile and timely to revisit this issue. The present work extends the solid recent analyses of the experimental teams [1,19,20] mainly in the following ways:

1. Inclusion of more parameters allows us to place constraints on neutrinos and gravity waves and to quantify the additional degeneracies that they introduce.

2. Joint analysis of all CMB data allows us to place stronger constraints and perform consistency tests.

3. Inclusion of explicit power spectrum modeling for the galaxy clustering and for the Lyman Alpha Forest allows stronger constraints and important new cross-checks.

II. IS THE CMB STORY CONSISTENT? COMPARING AND COMBINING MEASUREMENTS OF THE ANGULAR POWER SPECTRUM

In this section we test the CMB data for internal consistency and combine them into a single set of band powers, calibrating the experiments against each other.

A. CMB data

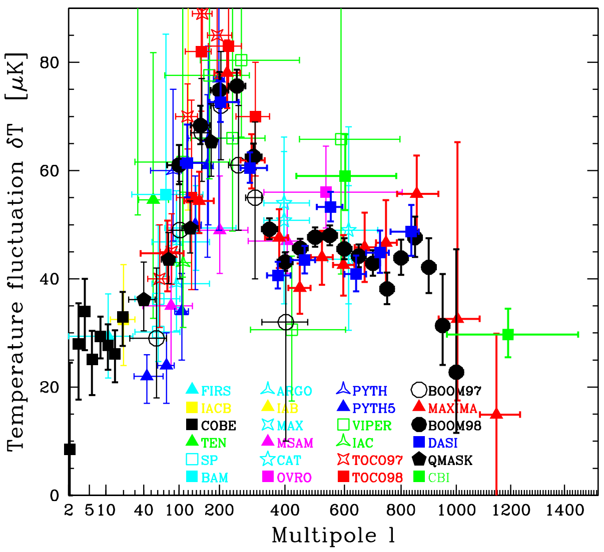

FIG. 1. CMB data used in our analysis. Error bars do not include calibration or beam errors which allow substantial vertical shifting and tilting for some experiments.

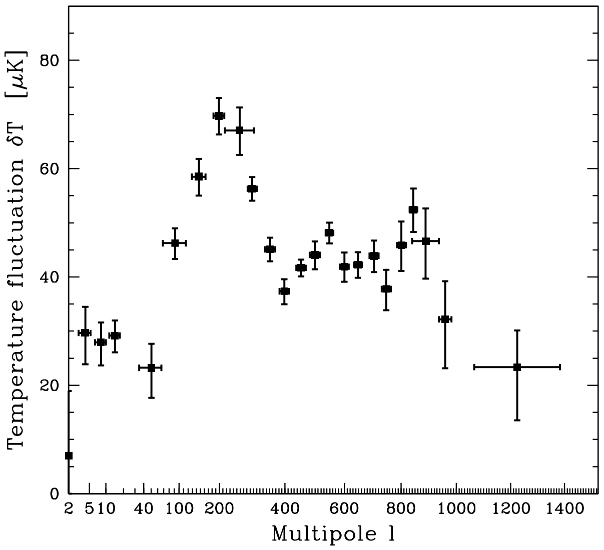

FIG. 2. Combination of all data from Figure 1. These error bars include the effects of beam and calibration uncertainties, which cause long-range correlations of order 5%-10% over the peaks. In addition, points tend to be anti-correlated with their nearest neighbors, typically at the level of a few percent. The horizontal bars give the rms widths of the window functions.

Figure 1 shows the 105 band power measurements used in our analysis. Starting with the data tabulated in [21], we have added the new measurements from CBI [18], QMASK [22], BOOMERaNG [1], DASI [2] and Maxima [3]. Since QMASK combines the Saskatoon [23] and QMAP [24–26] datasets, these have been omitted. A recent data review is given in [27].

The success of experimental CMB work has made data plots such as Figure 1 increasingly bewildering and difficult to interpret. Not only are do many data points with widely different error bars overlap, but important correlations due to calibration and beam uncertainties are difficult to visualize graphically and tend not to be included in the plotted error bars. The obvious solution to this problem is some sort of data compression.

A radical but common example of this is to simply throw away most of the data and show/analyze only one or two experiments, often the most recent. This is not ideal, however, since it both wasteful of information and lacks a objective criterion for data culling. Moreover, among the most accurate and thoroughly systematics-tested measurements on large scales still come from older maps (COBE DMR [28] and QMASK [36]).

A more desirable alternative is to average the data together somehow, into a single measurement of power on each angular scale. Such CMB data compression has been performed by many authors, *e.g.*, [10,21,29–31], and can retain all cosmological information provided that the new power bins are narrow enough to not smooth out important power spectrum features. However, such compression throws away any evidence for discrepancies between experiments that may have been present in the full data set, so it is important to complement the averaging by consistency checks.

B. Combining experiments

Let us group the power measurements at hand (say the 105 measurements of δT^2 from Figure 1) into a vector $\mathbf{y}$. We will model this as

$$\mathbf{y} = \mathbf{W}\mathbf{x} + \mathbf{n}, \tag{1}$$

where $\mathbf{x}$ is a vector containing the true power spectrum coefficients δT_ℓ^2 up to some sufficiently large multipole $\ell_{\max}$, the window function matrix $\mathbf{W}$ encodes the angular sensitivity of the measurements (the rows of $\mathbf{W}$ sum to unity) and the noise vector $\mathbf{n}$ represents all forms of measurement error. We model the errors as random with zero mean ($\langle\mathbf{n}\rangle = \mathbf{0}$) and with a covariance matrix $\mathbf{N} \equiv \langle\mathbf{n}\mathbf{n}^t\rangle$ that is the sum of four terms,

$$\mathbf{N} = \mathbf{N}^{(\mathrm{meas})} + \mathbf{N}^{(\mathrm{scal})} + \mathbf{N}^{(\mathrm{ical})} + \mathbf{N}^{(\mathrm{beam})}, \tag{2}$$

corresponding to basic measurement errors, source calibration errors, instrument calibration errors and beam errors, respectively. In general, all of these errors depend on the actual power spectrum $\mathbf{x}$, either through sample variance [32] or because calibration and beam errors are multiplicative rather than additive. Below we will make the approximation that the relative errors are small ($\ll 1$). In this limit, $\mathbf{N}$ reduces to a known matrix independent of $\mathbf{x}$. Explicit expressions for the four matrices in equation (2) are given in Appendix A.

Given $\mathbf{W}$, $\mathbf{N}$ and $\mathbf{y}$, it is straightforward to invert equation (1) to compute an estimate of the underlying power spectrum $\mathbf{x}$. This problem is mathematically identical to that involved in CMB mapmaking [33,34] except that the matrices involved are small enough to be trivial to manipulate numerically. As our estimator of $\mathbf{x}$ we use

$$\tilde{\mathbf{x}} \equiv \left[\mathbf{W}^t\mathbf{N}^{-1}\mathbf{W}\right]^{-1}\mathbf{W}^t\mathbf{N}^{-1}\mathbf{y}, \qquad (3)$$

which can be shown to be unbiased ($\langle\tilde{\mathbf{x}}\rangle = \mathbf{x}$), to minimize the rms noise in each power band and, if the noise properties are Gaussian, to retain all information about the true power spectrum $\mathbf{x}$ from the original data $\mathbf{y}$ [33]. The corresponding covariance matrix of the noise $\epsilon \equiv \tilde{\mathbf{x}} - \mathbf{x}$ is

$$\mathbf{\Sigma} \equiv \langle\varepsilon\varepsilon^t\rangle = \left[\mathbf{W}^t\mathbf{N}^{-1}\mathbf{W}\right]^{-1}. \qquad (4)$$

The resulting power spectrum $\tilde{\mathbf{x}}$ is shown in Figure 2 and listed in Table 1. The corresponding covariance matrix $\mathbf{\Sigma}$ is available at
www.hep.upenn.edu/ $\sim$ max/cmb/experiments.html.

When computing this spectrum, we did not treat the power as an independent parameter at each multipole. Rather, we treated the power spectrum as piecewise constant, parametrized by its height x_i in the 24 bands listed in Table 1. Since our compressed band powers $\tilde{\mathbf{x}}$ are simply linear combinations of the original measurements $\mathbf{y}$, we are able to compute their window functions exactly by taking the same linear combinations of the rows of $\mathbf{W}$ from equation (1). Our compressed data set can therefore be analyzed ignoring the details of how it was constructed (ignoring the first column of Table 1), viewed as simply the window matrix times the true power.

The rms widths of these window functions are reflected by the horizontal bars in Figure 2 and listed in the Table; the exact windows are available on the above-mentioned web site. This correlation matrix includes the residual effects of calibration uncertainty and beam errors. These long-range correlations are found to be moderate, typically of order 5-10% over the peaks, which shows that the experiments have to a certain extent been calibrated off of each other. In addition, there tends to be a slight anti-correlation between neighboring points, typically at the level of a few percent, as the power spectrum inversion performs a slight deconvolution of the input window functions from the experiments used.

One interesting feature of Figure 2 is that it shows both the first and second peak* somewhat lower than a large fraction of the data. Indeed, the preferred recalibrations for all five multiband experiments flagged in the following subsection are downward. To understand the origin of this effect, we performed a series of tests where the combined spectrum was recomputed with one or more experiments omitted. The explanation centers around the BOOM98 data, which combines a sharp constraint on the relative heights of the first two peaks (even taking the beam uncertainty into account) with a relatively large overall calibration uncertainty. Since both Maxima and DASI (with one exception) have points with small error bars below BOOM98 around the second peak, they pull the BOOM98 calibration down. QMASK also pulls BOOM98 down because of its statistical weight around $\ell \sim 100$. Although Figure 1 suggests that influence from, *e.g.*, TOCO around the first peak might raise BOOM98, this pull is weaker because of TOCO calibration uncertainties and since error bars are overall smaller out at the second peak. This somewhat low normalization persists even if any one of Maxima, DASI or QMASK is excluded from the analysis.

Table 1 – Band powers combining the information from all 105 CMB data points from Figure 1. The 2nd column gives the means and rms widths of the window functions. The spectrum was computed treating δT^2 as constant in the bands listed in the first column. The error bars in the 3rd column include the effects of calibration and beam uncertainty. The full 24×24 correlation matrix and 24×2000 Window matrix are available at at
www.hep.upenn.edu/ $\sim$ max/cmb/experiments.html.

ℓ-Band	ℓ-window	δT^2 [μK^2]
$2-2$	2 ± 0	50 ± 319
$3-5$	4 ± 1	880 ± 308
$6-10$	8 ± 2	780 ± 219
$11-30$	14 ± 3	853 ± 171
$31-75$	50 ± 16	540 ± 227
$76-125$	92 ± 23	2137 ± 263
$126-175$	146 ± 17	3423 ± 394
$176-225$	199 ± 15	4865 ± 469
$226-275$	258 ± 43	4496 ± 584
$276-325$	296 ± 14	3169 ± 248
$326-375$	352 ± 16	2035 ± 195
$376-425$	398 ± 16	1396 ± 172
$426-475$	453 ± 15	1737 ± 129
$476-525$	499 ± 17	1942 ± 226
$526-575$	549 ± 13	2319 ± 185
$576-625$	600 ± 15	1755 ± 225
$626-675$	648 ± 14	1787 ± 200
$676-725$	704 ± 16	1928 ± 254
$726-775$	747 ± 16	1427 ± 280
$776-825$	802 ± 16	2106 ± 418
$826-875$	845 ± 14	2750 ± 420
$876-925$	889 ± 49	2173 ± 598
$926-1025$	959 ± 22	1036 ± 501
$1026-\infty$	1224 ± 159	546 ± 362

C. Comparing experiments

Since the combined power spectrum presented above is only as reliable as the data that went into it, let us now test this data for internal consistency.

*One of the Argentinian authors feels that there is still no conclusive evidence for multiple peaks in the power spectrum, especially in regard to any bets or wagers that may or may not be outstanding.

As mentioned above, equation (1) is analogous to the CMB mapmaking problem, which means that all methods developed for comparing and combining maps can be applied to comparing and combining power spectra as well. Given two power spectrum measurements $\mathbf{y}_1$ and $\mathbf{y}_2$ modeled as in equation (1), *e.g.*, as

$$\mathbf{y}_1 = \mathbf{W}_1\mathbf{x} + \mathbf{n}_1, \quad \mathbf{y}_1 = \mathbf{W}_2\mathbf{x} + \mathbf{n}_2, \tag{5}$$

we wish to know whether they are consistent or display evidence of systematic errors. Specifically, is there some underlying power spectrum $\mathbf{x}$ such that the data sets $\mathbf{y}_1$ and $\mathbf{y}_2$ are consistent with equation (5)? Let us consider the simple case where the two experiments have identical window functions, that is, $\mathbf{W}_1 = \mathbf{W}_2$. The general case can be reduced to this one: In practice, we start by reducing all experiments to the simple form $\mathbf{W} = \mathbf{I}$ using the deconvolution method described in Appendix D of [22].

Consider two hypotheses:

- H_0: The null hypothesis H_0 that there are no systematic errors, so that the difference spectrum $\mathbf{z} \equiv \mathbf{y}_1 - \mathbf{y}_2$ consists of pure noise with zero mean and covariance matrix $\langle \mathbf{z}\mathbf{z}^t \rangle = \mathbf{N} \equiv \mathbf{N}_1 + \mathbf{N}_2$.
- H_1: The alternative hypothesis that the difference spectrum $\mathbf{z}$ has the same covariance matrix $\mathbf{N}$ but a non-zero mean $\mathbf{m}$.

A straightforward variation of the derivation in [35] shows that the "null-buster" statistic [35]

$$\nu \equiv \frac{\mathbf{z}^t\mathbf{N}^{-1}\mathbf{Q}\mathbf{N}^{-1}\mathbf{z} - \mathrm{tr}\,\{\mathbf{N}^{-1}\mathbf{Q}\}}{[2\,\mathrm{tr}\,\{\mathbf{N}^{-1}\mathbf{Q}\mathbf{N}^{-1}\mathbf{Q}\}]^{1/2}},$$
$$\mathbf{Q} \equiv \mathbf{m}\mathbf{m}^t, \tag{6}$$

can be shown to rule out the null hypothesis H_0 with the largest average significance $\langle\nu\rangle$ if H_1 is true, and can be interpreted as the number of "sigmas" at which H_0 is ruled out [35]. The case derived in [35] differed in that the mean vanished under H_1 but that the covariance matrix contained extra signal $\mathbf{S}$ — the result was of the same for as above, but with $\mathbf{Q} = \mathbf{S}$. Note that for the special case $\mathbf{Q} \propto \mathbf{N}$, it reduces to simply $\nu = (\chi^2 - n)/\sqrt{2n}$, where $\chi^2 \equiv \mathbf{z}^t\mathbf{N}^{-1}\mathbf{z}$ is the standard chi-squared statistic. The null-buster test can therefore be viewed as a generalized χ^2-test which places more weight on those particular modes where the expected signal-to-noise is high. It has proven successful comparing both microwave background maps [24–26,22,36] and galaxy distributions [37,38]. Tips for rapid implementation in practice are given in [22].

Equation (6) shows that in our case, all weight is placed on a single mode $\mathbf{m}$. More generally, the test pays the greatest attention to those eigenvectors of $\mathbf{Q}$ whose eigenvalues are large. Consider first the case of calibration errors. Then the two power measured spectra are generically expected to have the same shape but different normalizations, so that the vectors $\langle\mathbf{y}_1\rangle$ and $\langle\mathbf{y}_2\rangle$ are parallel but with different lengths. In other words, $\mathbf{S} = \mathbf{0}$ and $\mathbf{m} = \langle\mathbf{z}\rangle \propto \mathbf{W}\mathbf{x}$, so the mode $\mathbf{m}$ that we want our test to be sensitive to is shaped like the the expected sky power spectrum itself. Similarly, beam errors show up in a different mode, which (as discussed in Appendix A) is shaped like the sky power spectrum times ℓ^2 to first order.

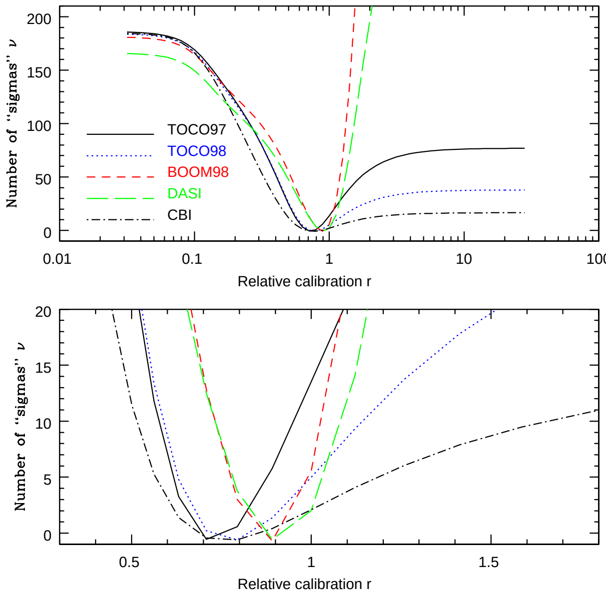

FIG. 3. Each curve shows the number of standard deviations ("sigmas") at which a given experiment is inconsistent with all others when its power spectrum δT is multiplied by a constant r.

D. Test results

As emphasized by, *e.g.*, Press [40], it is important to check any data collection for statistically significant outliers. We do this separately for each of the 24 experiments from Figure 1 as follows. We form the difference spectrum

$$\mathbf{z} \equiv \mathbf{y}_2 - r\mathbf{y}_1, \tag{7}$$

where $\mathbf{y}_1$ and $\mathbf{y}_2$ are the deconvolved power spectra produced from the experiment under study and from all other experiments, respectively, and vary the normalization parameter r. We take calibration and beam errors into account in computing the "all other" spectrum, but not for the experiment under consideration. The resulting significance level ν at which the difference spectrum is inconsistent with zero is plotted as a function of r in Figure 3.

All experiments taken together have now detected CMB fluctuations at about the 190σ level, and the fact that all curves asymptote to just under that value as $r \to 0$ shows that no single experiment dominates all others in statistical weight. To the right, as $r \to \infty$, each curve asymptotes to the significance level at which the experiment in question detects signal. If the experiment under study has no calibration or beam errors and everything is consistent, ν should be near zero for $r = 1$, where it has vanishing mean and unit variance ($\langle\nu\rangle = 0$, $\Delta\nu = 1$). Only five of the 24 experiments show a significant difference at the 2σ-level ($\nu(1) > 2$). Their $\nu(r)$-curves are plotted in Figure 3, are seen to be perfectly consistent as well — the relative calibration r simply has to be shifted to a different value, which is in all cases lower, for which ν drops below 2.

Above we saw that the experiments could be effectively recalibrated off of each other. We also computed ν for the latest BOOMERaNG data on a 2-dimensional grid, varying both the calibration and the beam size. However, beam information from inter-experiment comparison is substantially less useful at the present time than extracting the corresponding calibration information — we found the BOOMERaNG beam constraints obtained in this way to be weaker than those measured from the instrument directly.

In conclusion, the quantitative tests described above show no evidence of inconsistency between the 24 CMB experiments when beam and calibration uncertainties are taken into account, and the power spectrum shown in Figure 2 is consistent with all of them. The dip around $\ell = 50$ may warrant further investigation to study if a smoother underlying spectrum can be consistent with all the data in that range.

III. IS THE COSMOLOGY STORY CONSISTENT? COMPARING AND COMBINING DIFFERENT COSMOLOGICAL DATASETS

In this section we confront the CMB data with other cosmological observations, with the goal being both accurate constraints on cosmological parameters and various cross-checks on the underlying physical assumptions. We first map out the subset of the 11-dimensional cosmological parameter space from [39] that is consistent with CMB, large scale structure (LSS) and Lyman Alpha Forest (LyαF) power spectra, with Big Bang nucleosynthesis (BBN) and with direct Hubble constant determinations, included and excluded in various combinations. We pay particular attention to whether these priors are mutually consistent or pull in different directions, both in terms of which parts of parameter space they pull towards and in terms of how much they increase the overall χ^2. We will return to the consistency issue in Section IV, comparing our "concordance model" with various other cosmological constraints.

A. Analysis method

We employ the multiparameter analysis method described in [39] and [21] with the following modifications: optional inclusion of Lyα power spectra, optional discretization of the gravity wave contribution to allow explicit limits on this parameter, inclusion of CMB beam uncertainties as described in Appendix A and refined parameter grids to reflect the improved data accuracy. It consists of the following steps:

1. Compute CMB, LSS and LyαF power spectra for a grid of models in our 11-dimensional parameter space.
2. Compute a likelihood for each model that quantifies how well it fits the data.
3. Perform 11-dimensional interpolation and marginalization to obtain constraints on individual parameters and parameter pairs.

Our 11 parameters are

$$\mathbf{p} \equiv (\tau, \Omega_{\mathrm{k}}, \Omega_\Lambda, \omega_{\mathrm{dm}}, \omega_{\mathrm{b}}, f_\nu, n_s, n_t, A_s, r, b). \quad (8)$$

These are the reionization optical depth τ, the primordial amplitudes A_s, rA_s and tilts n_s, n_t of scalar and tensor fluctuations, a bias parameter b defined as the ratio between rms galaxy fluctuations and rms matter fluctuations on large scales, and five parameters specifying the cosmic matter budget. The various contributions Ω_i to critical density are for curvature Ω_{k}, vacuum energy Ω_Λ, cold dark matter Ω_{cdm}, hot dark matter (neutrinos) Ω_ν and baryons Ω_{b}. The quantities $\omega_{\mathrm{b}} \equiv h^2\Omega_{\mathrm{b}}$ and $\omega_{\mathrm{dm}} \equiv h^2\Omega_{\mathrm{dm}}$ correspond to the physical densities of baryons and total (cold + hot) dark matter ($\Omega_{\mathrm{dm}} \equiv \Omega_{\mathrm{cdm}} + \Omega_\nu$), and $f_\nu \equiv \Omega_\nu/\Omega_{\mathrm{dm}}$ is the fraction of the dark matter that is hot. We assume that the galaxy bias b is constant on large scales but make no assumptions about its value, and therefore marginalize (minimize) over this parameter before quoting constraints on the other ten. In the adaptive mesh spirit, we iteratively refined our parameter grid to adequately resolve the peak of the likelihood function. Our final parameter grids were as follows:

- $\tau = 0, 0.05, 0.1, 0.2, 0.3, 0.5, 0.8$
- $\Omega_\Lambda = 0, 0.1, ..., 1.0$
- $\Omega_{\mathrm{k}} = \pm1.0, \pm0.8, \pm0.6, \pm0.5, \pm0.4, \pm0.3, \pm0.2,$ $\pm0.1, \pm0.03, 0,$ truncated so that $\Omega_{\mathrm{m}} \equiv 1 - \Omega_{\mathrm{k}} - \Omega_\Lambda \in [0.05, 1]$
- $\omega_{\mathrm{dm}} = .02, .05, .08, .11, , .13, .16, .20, .50$
- $\omega_{\mathrm{b}} = .003, .015, .018, .020, .022, .025, .03, .04, .07$
- $f_\nu = 0, 0.05, 0.1, 0.2, 0.3, 0.4, 0.6, 0.8, 1.0$
- $n_s = 0.5, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.4, 1.7$

- $n_t = -1.0, -0.7, -0.4, -0.2, -0.1, 0$
- $r = 0, 0.1, 0.2, 0.3, 0.4, 0.6, 0.8, 1.0, 1.4, 1.8, 2.5$
- A_s is not discretized
- b is not discretized

The parameter r was only discretized when computing constraints involving r — it was treated as continuous and marginalized over analytically as in [39] when constraining other parameters. We used the original CMB data set rather than the compressed one to compute the CMB likelihood, since step 2 above is the fastest of the three in our analysis pipeline anyway.

B. Non-CMB data used

As our LSS data, we follow [39] in using the power spectrum measured from by the IRAS PSCz survey [41] by [42], discarding all measurements on scales smaller than $20h^{-1}$Mpc to be conservative.

As our LyαF data, we use the 13 recent power spectrum measurements of [43], with an additional 27% calibration uncertainty common to all points. We compute theoretical predictions for these 13 numbers corresponding to each of the $\sim$ 300 million models in our database by first computing the matter power spectrum $P(k)$ as in [39], then shifting it vertically and horizontally (on a log-log scale) to account for the fluctuation growth and Hubble parameter shift between $z \sim 2.72$ and the present epoch. We used the growλ package [44] for computing the growth factors.

We quantitatively explore how our results are affected by adding various other constraints ("priors"). For BBN, we test the prior $\omega_{\rm b} = 0.02$ for simplicity, since the error bars on the most recent BBN estimates $\omega_{\rm b} = 0.02 \pm 0.002$ (95%) from [45] are smaller than our grid spacing. For the Hubble parameter, we test the prior $h = 0.72 \pm 0.08$ from the HST Hubble Key Project [46], assuming a Gaussian distribution for simplicity. We also try the priors $\tau = 0$, $\Omega_{\rm k} = 0$, $f_\nu = 0$, $r = 0$ and $n_s = 1$ in various combinations.

C. Basic results

Our constraints on individual cosmological parameters are listed in Table 2 for four cases. Constraints are plotted in figures 4 and 5 for cases 2 and 3. All tabulated and plotted bounds are 95% confidence limits†. The first case uses constraints from CMB alone, which are still rather weak because of a one-dimensional degeneracy coupling curvature, baryons, tilt, tensors, dark matter and dark energy as described in the following subsection. The second case breaks this degeneracy by combining the CMB information with the power spectrum measurements from PSCz, and is seen to give rather interesting constraints on most parameters. The third case adds the prior assumptions that the Hubble parameter is $h = 0.72 \pm 0.08$ [46], which tightens up many constraints, in particular that on Ω_Λ. The fourth case adds the assumption that the neutrino contribution is cosmologically negligible ($f_\nu \sim 0$). This is equivalent to assuming that there is no strong mass-degeneracy between the relevant neutrino families, and that the Super-Kamiokande atmospheric neutrino data therefore set the scale of the neutrino density to be $\omega_\nu \lesssim 10^{-3}$ [70,71].

Table 2 – Best fit values and 95% confidence limits on cosmological parameters. z_{ion} is the redshift of reionization and t_0 is the present age of the Universe. For the numbers below the horizontal line, the limits were computed from moments as described in the text, so the central values are means rather than those for the best fit model. If the reader wishes to use some of these model parameters for other purposes, the values to use are thus those inferred from the numbers above the line. For instance, the Hubble parameters for the best fit models are h =0.51, 0.48, 0.64 and 0.64 for the four columns, respectively.

	CMB alone	+ PSCz	+ $h = .72 \pm .08$	+ $f_\nu \sim 0$
τ	$0.00^{+.30}$	$0.00^{+.17}$	$0.00^{+.24}$	$0.00^{+.23}$
$\Omega_{\rm k}$	$-0.06^{+.13}_{-.59}$	$-0.05^{+.08}_{-.15}$	$0.00^{+.06}_{-.06}$	$0.00^{+.05}_{-.06}$
Ω_Λ	$0.60^{+.32}_{-.55}$	$0.53^{+.17}_{-.40}$	$0.66^{+.10}_{-.17}$	$0.66^{+.10}_{-.13}$
$h^2\Omega_{\rm dm}$	$0.10^{+.07}$	$0.10^{+.04}_{-.04}$	$0.12^{+.05}_{-.03}$	$0.12^{+.05}_{-.03}$
$h^2\Omega_{\rm b}$	$0.02^{+.06}_{-.01}$	$0.02^{+.01}_{-.01}$	$0.02^{+.01}_{-.01}$	$0.02^{+.01}_{-.01}$
f_ν	0.07	$0.05^{+.30}$	$0.04^{+.17}$	*0*
n_s	$0.93^{+.75}_{-.16}$	$0.91^{+.13}_{-.10}$	$0.93^{+.13}_{-.10}$	$0.93^{+.12}_{-.10}$
h	$0.42^{+.23}_{-.24}$	$0.57^{+.31}_{-.30}$	$0.71^{+.12}_{-.12}$	$0.73^{+.11}_{-.10}$
z_{ion}	$5.7^{+12.9}_{-5.7}$	$7.2^{+13.8}_{-7.2}$	$7.6^{+14.2}_{-7.6}$	$6.8^{+13.1}_{-6.8}$
t_0 [Gyr]	$20.5^{+9.0}_{-9.0}$	$14.2^{+4.3}_{-4.3}$	$12.3^{+1.6}_{-1.6}$	$12.7^{+1.5}_{-1.5}$

For the first 8 parameters listed in Table 2, the numbers were computed from the corresponding 1-dimensional likelihood functions (plotted in Figure 4 and Figure 5 for the second and third cases). The best fit value corresponds to the peak in the likelihood function

† Bayesean 95% confidence limits are in general those that enclose 95% of the area. In this paper, we make the approximation that the boundary of the confidence region is that where the likelihood has fallen by a factor e^{-2} from its maximum for 1-dimensional cases (such as the numbers in Table 2) and by a factor $e^{-6.18}$ for 2-dimensional cases (such as figures 4 and 5). As shown in Appendix A of [21], this approximation becomes exact only for the case when the likelihood has a multivariate Gaussian form. We make this approximation to be consistent with the multidimensional marginalization algorithm employed here (and by most other authors), which is equivalent to the integration technique only for the Gaussian case.

and the 95% limits correspond to where the likelihood function drops below the dashed line at e^{-2} of the peak value. For the remaining parameters listed, which are not fundamental parameters in our 11-dimensional grid, the numbers were computed as in [10] by calculating the likelihood-weighted means and standard deviations over the multidimensional parameter space. Here the tabulated limits are the mean $\pm 2\sigma$.

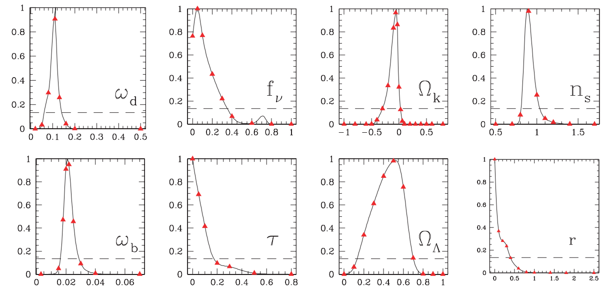

FIG. 4. Constraints on individual parameters using only CMB and LSS information. The quoted 95% confidence limits are where each curve drops below the dashed line.

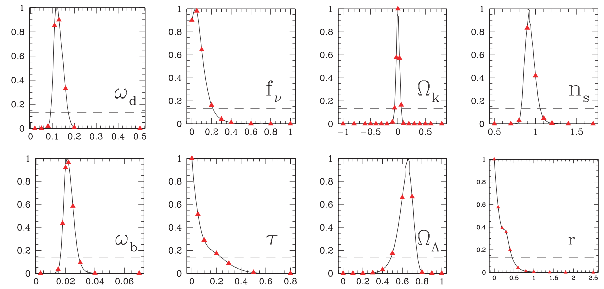

FIG. 5. Like the previous figure, but adding the prior $h = 0.74 \pm 0.08$.

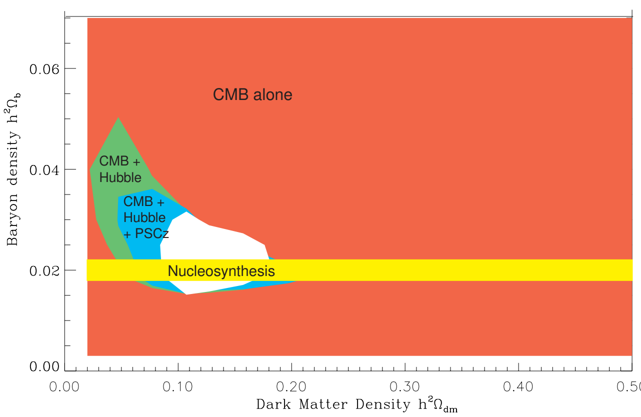

FIG. 6. Constraints in the $(\omega_{\rm dm}, \omega_{\rm b})$-plane. The large red/dark grey region is ruled out by CMB data along at 95% confidence. The two smaller shaded regions become excluded when imposing additional constraints from Hubble constant measurements ($h = 0.74 \pm 0.08$) and PSCz galaxy clustering. The yellow/light gray band shows the BBN constraints from [45]. The model best fitting the CMB+PSCz+h constraints has $\chi^2 \approx 126$ for 127 degrees of freedom.

D. Matter budget

We will now investigate the parameter constraints in more detail, exploring which conclusions come from what assumptions. This subsection is centered around the cosmic matter budget (the densities of baryons, cold dark matter, hot dark matter, dark energy and curvature) — we return to the inflationary parameters n_s, n_t and r in the next subsection.

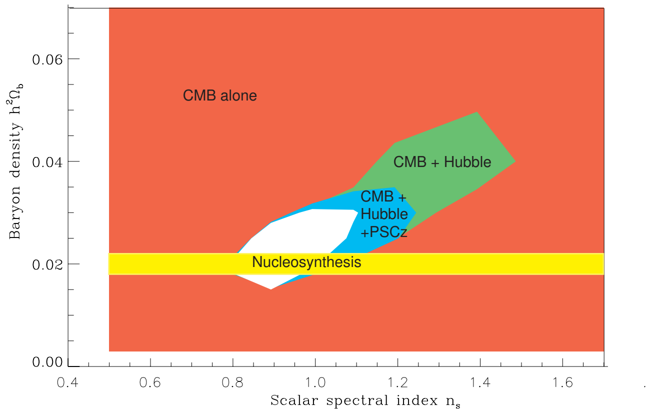

FIG. 7. Constraints in the $(n_s, \omega_{\rm b})$-plane. The yellow/light gray band shows the BBN constraints from [45].

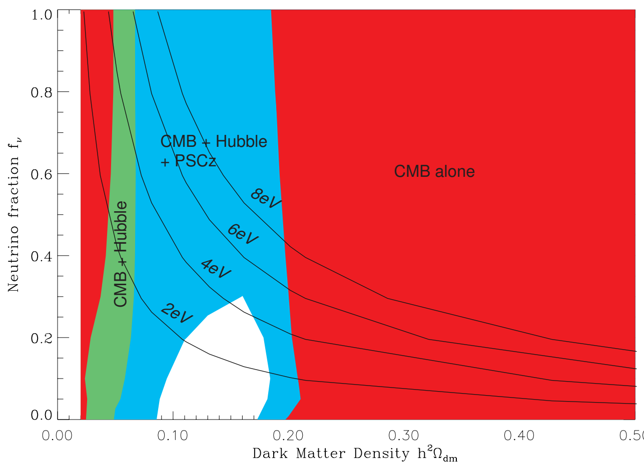

FIG. 8. Constraints in the $(\omega_{\rm dm}, f_\nu)$-plane. The four curves show contours of constant neutrino mass.

1. Constraints from CMB alone

Figures 6, 7 and 8 summarize some of the key 2-dimensional constraints on the matter budget. It is seen that the CMB data alone are now for the first time (compare, eg, [39]) powerful enough to close off bounded regions in these planes. The low second acoustic peak inferred from the first Antarctic BOOMERaNG results [48] was still consistent with a purely baryonic Universe, prompting speculation [49] that an alternative theory of gravity might be able to eliminate the need to for dark matter altogether. Since the more accurate measurements of the second peak height from the new Boomerag, DASI and Maxima data give a higher value, even the CMB alone now requires a non-zero amount of dark matter, at least at modest significance.

The increased second peak height is also seen to resolve a second hot discussion topic of the past year: the CMB lower limit on the baryon density has now dropped down in beautiful agreement with the BBN prediction.

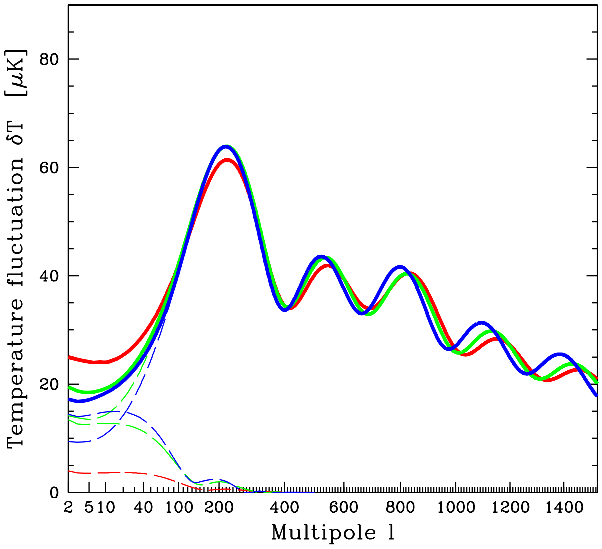

FIG. 9. The main remaining CMB degeneracy. The three solid curves show the best fit models for CMB alone subject to the constraint $\omega_{\rm b} = 0.02$ (red), $\omega_{\rm b} = 0.03$ (green) and $\omega_{\rm b} = 0.04$ (blue), spanning the "banana" allowed by the CMB in Figure 7. Dashed curves show the scalar and tensor contributions. All models have $\tau = f_\nu = 0$. The parameters $\mathbf{p} \equiv (\Omega_{\rm k}, \Omega_\Lambda, \omega_{\rm dm}, \omega_{\rm b}, n_s, n_t)$ are $(-0.043, 0.461, 0.116, 0.020, 0.924, 0.982)$, $(-0.134, 0.799, 0.050, 0.030, 1.200, 0.884)$ and $(-0.160, 0.840, 0.043, 0.040, 1.401, 0.955)$, respectively.

A third noteworthy result is that the allowed region in our 11-dimensional space is shaped like a long and skinny rather one-dimensional tube. This is seen clearly in the projection onto the $(n_s, \omega_{\rm b})$ plane in Figure 7. The physics underlying this CMB degeneracy is illustrated in Figure 9. Starting in the generally favored (white) region of Figure 7 and moving up to the right, the Universe becomes more and more closed ($\Omega_{\rm k}$ decreases to negative values), which would on its own shift the acoustic peaks too far to the left. This is compensated by reducing the density of dark matter $\omega_{\rm dm}$ and increasing the dark energy density Ω_Λ and the power spectrum tilt n_s so that the peak location stays essentially fixed. Through a rather spurious coincidence, the resulting changes in the various peak heights relative to the COBE scale can be almost perfectly reversed by increasing the baryon density and adding substantial amounts of gravity waves. This is in essence the familiar $(\Omega_{\rm k}, \Omega_\Lambda)$ angular distance degeneracy described in many parameter forecasting papers [50,51], with the extra twist that changing $\omega_{\rm b}$, $\omega_{\rm dm}$, n_s and r as well helps minimize the change in the late integrated Sachs-Wolfe effect on COBE scales.

2. Breaking the CMB degeneracy

Since this degeneracy involves so many parameters, prior constraints on any one of them will help break it. In particular, the reason that this degeneracy is not as prominent in the recent analyses by the BOOMERaNG [1], DASI [19] and Maxima [20] teams is because they all assumed negligible gravity waves, $r = 0$.

Since the Hubble parameter is given by

$$h = \sqrt{\frac{\omega_{\rm dm} + \omega_{\rm b}}{1 - \Omega_{\rm k} - \Omega_\Lambda}}, \tag{9}$$

it decreases sharply as one moves along the degeneracy track, dropping as low as 0.3 at the upper right endpoint in Figure 7. Imposing the prior $h = 0.72 \pm 0.08$ therefore shrinks the allowed 11-dimensional region substantially, among other things tightening the lower limit on dark matter in Figure 6 and the upper limit on baryons in Figure 7.

Because of its sensitivity to n_s and $\omega_{\rm dm}$ in particular, our galaxy clustering data (the PSCz power spectrum) is seen to break this CMB degeneracy completely, shrinking the allowed CMB "bananas" to almost round regions in Figure 6 and Figure 7. The effect of adding PSCz is particularly dramatic in Figure 8, since CMB alone has almost no sensitivity to the neutrino fraction, whereas increasing f_ν provides a strong suppression of the galaxy power spectrum on small scales.

The model best fitting the CMB+PSCz+h constraints has $\chi^2 \approx 126$ for $105 + 32 + 1 - 11 = 127$ degrees of freedom. The effective number of degrees of freedom might be a few larger than this, since some of the 11 parameters had little effect, but even if we ignore this, all fits are good in the sense of giving reduced χ^2-values of order unity.

Adding our additional priors $\omega_{\rm b} = 0.02$, $\tau = 0$, $\Omega_{\rm k} = 0$ and $r = 0$ cause little further change, since they are all consistent with the favored results and there are no major degeneracies left to break.

Adding our LyαF data produced effects quite similar to those of adding PSCz: breaking the CMB degeneracy and favoring a flat, roughly scale-invariant Universe. We have chosen to highlight the effects of the

PSCz data rather those of the LyαF data in the figures since we found its overall constraining power on parameters to be slightly stronger. For instance, the CMB+LyαF constraints on tilt, curvature and baryon density are $0.91^{+0.14}_{-0.10}$, $-0.03^{+0.07}_{-0.22}$ and $0.021^{+0.008}_{-0.005}$, respectively. Overall, the PSCz and LyαF were found to be strikingly consistent in pulling the CMB in the same direction, as was also found in [12]. Augmenting the CMB+PSCz data with the 13 LyαF points increased the χ^2 for the best fit model by only 10, and the good agreement between the galaxy and LyαF can also be seen visually, in Figure 10.

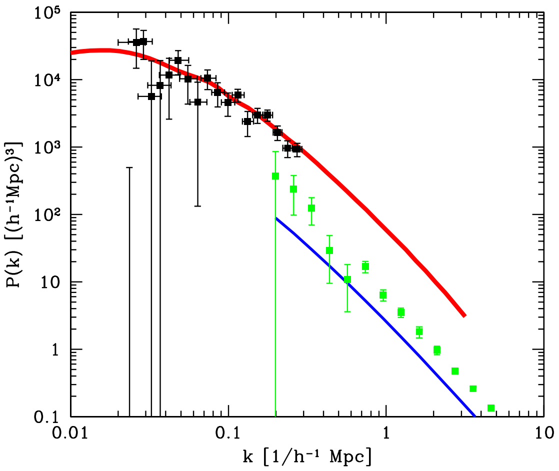

FIG. 10. The model best fitting the CMB, LSS and LyαF data with h and f_ν priors is shown together with the two data sets. The red and blue curves are the model predictions for the LSS and LyαF data, respectively. The blue curve is simply the red one shifted vertically and horizontally to account for the fluctuation growth and Hubble parameter shift between $z \sim 2.72$ and the present epoch. The overall calibration uncertainty in the LyαF data is seen to let the model prediction lie below the data by a constant factor.

We conclude this subsection by summarizing what is obtained by adding successively stronger assumptions as in Table 2.

1. CMB alone now gives constraints on most parameters, but they are generally weak because of the above-mentioned degeneracy.
2. Adding the PSCz galaxy clustering constraints breaks this degeneracy, resulting in tight constraints on everything except the dark energy density Ω_Λ.
3. Adding the constraint $h = 72 \pm 8$ finally nails down Ω_Λ. It also sharpens the limits on Ω_k and f_ν. Adding this one constraint raises χ^2 by as much as 4 (from 122 to 126), which reflects a slight tension between the h-prior and the CMB+PSCz data, which alone prefer the lower range $h = 0.57^{+.31}_{-.30}$.
4. Adding an $f_\nu = 0$ prior boosts χ^2 only by 0.1, further sharpening the Ω_k, Ω_Λ and τ constraints slightly, and none of our additional priors (including LyαF) have much of an effect. The only exception is imposing $n_s = 1$, which is slightly disfavored by the data and raises χ^2 by 4.

E. Inflationary parameters

We now turn our attention to the parameters associated with inflation. Since space remains perfectly consistent with the inflationary flatness prediction despite the sharp error bar reduction caused by the new latest CMB data ($\Omega_k = 0.00 \pm 0.06$ at 95% for CMB+PSCz+h), it is interesting to quantify the constraints on the parameters n_s, n_t and r to see how they compare with the predictions of various models, as has previously been done using earlier CMB data [11,12].

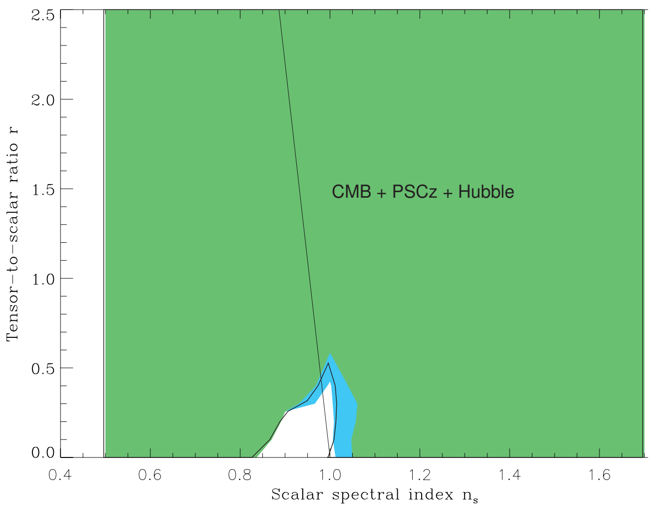

FIG. 11. Constraints in the (n_s, r) plane. The green region shows the joint constraints from CMB, LSS and the priors $h = 0.72 \pm 0.08$ and $f_\nu = 0$. The blue region shows the effect of adding the BBN prior $\omega_b = 0.02$, and the black curve shows the effect of dropping the neutrino prior. The straight line shows the prediction $r = (200/9)(1 - n_s)$ from a power law class of inflation models.

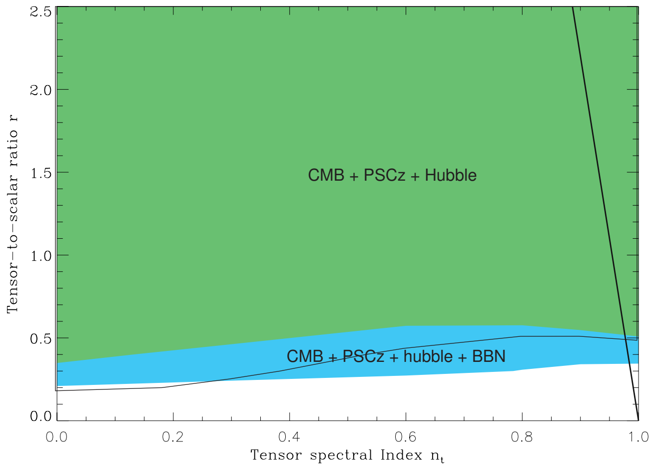

FIG. 12. Same as previous figure but for the (n_t, r) plane. The straight line shows the "inflationary consistency relation" $r = -(200/9)n_t$.

As discussed in [52], there has been a fair amount of notational confusion in the literature surrounding the tensor-to-scalar ratio r. We define

$$r \equiv \frac{A_t}{A_s}, \tag{10}$$

where A_s and A_t are the scalar and tensor fluctuation amplitudes as defined in [52]. For inflation models where the slow-roll approximation is valid, this ratio is related to the tensor tilt n_t by the so-called inflationary consistency condition [53,52]

$$r = -\frac{200}{9} n_t. \tag{11}$$

For models with $\Omega_k = \Omega_\Lambda = 0$, this r has often been related to the tensor-to-scalar quadrupole ratio with the approximation

$$r \approx a \frac{C_2^{\rm tensor}}{C_2^{\rm scalar}}, \tag{12}$$

with a proportionality constant a typically between 3 and 5, but we find this approximation to be highly inaccurate for many models in our grid. In particular, the ISW effect is non-negligible now that models with $\Omega_\Lambda > 0$ are favored, so equation (12) is best avoided. A power law class of inflation models (see [11] for a review) make the additional prediction that

$$r \approx \frac{200}{9}(1 - n_s). \tag{13}$$

i.e., that $n_s = n_t + 1$. (Although the quantity $n_t + 1$ would be a more natural definition for the tensor spectral index, we will stay true to the astronomical tradition of clinging to silly notation for historical reasons.)

As shown in Table 2, Figure 4 and Figure 5, the data prefers no tensor contribution at all, placing a 95% upper limit $r < 0.35$ for the CMB+PSCz+h case. For comparison with prior work, this corresponds to a quadrupole ratio of 0.2 in the sense that this is the quadrupole ratio for the best fit model in our grid with this r-value (which is by definition ruled out at exactly the 95% level). For comparison, previous multiparameter analyses incorporating gravity waves [11,12] found $r < 0.3$ using older CMB data and stronger priors.

Figures 11 and 12 show the joint constraints on r with the scalar and tensor tilts, respectively, and allows comparison with the predictions of equations (11) and (13). The constraints in the (n_s, r) have becomes progressively sharper during the past year [11,12], and are now gradually becoming quite interesting. In particular, Figure 11 shows that the preference for a slight red-tilt ($n_s \sim 0.9$) and low r in the data favors so-called "weak-field" models [11].

Figure 12 shows that the holy grail of testing the inflationary consistency relation (11) is still a ways off. As expected, the constraints on the tensor tilt n_t are stronger for high tensor normalizations r and vanish completely when r does.

IV. DISCUSSION

In Section II, we compared and combined the different CMB experiments, finding that a consistent picture of the angular power spectrum emerged. In Section III, we compared and combined a limited number of cosmological data sets (power spectra from CMB, LSS, LyαF and various priors), finding that a consistent set of cosmological parameters emerged that provided a good fit to this data.

We conclude with some remarks on how these parameter measurements match up with the many other cosmological observations that probe these parameters, focusing on the cosmic matter budget. In light of the checkered history of many cosmological parameters, where tiny quoted error bars have repeatedly masked larger uncertainties about underlying assumptions, such end-to-end consistency checks are crucial. For instance, the Hubble parameter h has dropped by a factor of eight since Hubble's original paper, the BBN baryon density ω_b has risen by 50% in less than a decade, and the favored value of the cosmological constant Ω_Λ has fluctuated wildly at the hands on both theorists and observers.

A. Baryon density ω_b

The rise of the second peak in the new data has completely eliminated the tension between BBN and CMB, and they two are now in beautiful agreement that the baryon density $\omega_b \approx 0.02$. This agreement was noted in the latest team papers as well [1,19,20]. That one method involving nuclear physics when the Universe was seconds old and another involving plasma physics more than 100,000 years later give the same answers, despite involving completely different systematics, is a landmark achievement for cosmology. It greatly boosts the credibility of the basic cosmological storyline since the Universe was a split second old. This sudden agreement is all the more impressive given the lack thereof during the past year and the ado this generated.

Consensus has yet to be reached on the second decimal of the BBN predictions, with the value $\omega_{\rm b} = 0.023$ from a recent VLT deuterium study [54] lying above the 95% range of [45]. As a reality check, our baryon value also agrees with slightly less accurate estimates of the abundance in the local Universe — *e.g.*, the range $0.007 \lesssim \Omega_{\rm b} \lesssim 0.041$ inferred from a low-redshift inventory [55] and the range $0.015 \lesssim \omega_{\rm b} \lesssim 0.03$ at redshifts of a few from the Lyα forest [57,58], although the latter tend to be on the high side. Similarly, the inferred baryon fraction $\omega_{\rm b}/\omega_{\rm dm} \sim 18\%$ agrees with that inferred from galaxy clusters [59,16], although this match weakly prefers lower h-values.

B. Dark matter and dark energy density

By now, there are a large number of independent methods for probing the dark matter density $\omega_{\rm dm}$ or the (almost identical) total density, including studies of the cluster abundance at various redshifts [64–66,59,60], mass-to-light ratios [61,62], the baryon fraction $\Omega_{\rm b}/\Omega_{\rm m}$ from cluster studies [63], cosmic velocity fields, redshift space distortions, mass-to-light ratio techniques. Most of these probes do not measure $\omega_{\rm dm}$ directly, but $\Omega_{\rm dm} = \omega_{\rm dm}/h^2$ or some combination of this with the fluctuation normalization parameter σ_8. Although each comes with potential systematic errors of its own, and there is still some internal controversy in the velocity and cluster areas [67], they are all broadly consistent with the range $\omega_{\rm dm} = 0.09 - 0.18$ listed in Table 2. Indeed, it is interesting that the some of the last evidence supporting $\Omega_{\rm m} \sim 1$ came from velocity fields and redshift distortion studies, and that improved data now gives values as low as in Table 2 from both velocity fields [68] and redshift space distortions [42,69].

With CMB strongly favoring an essentially flat Universe, the low matter density automatically implies a high dark energy density. As has been discussed in numerous papers (*e.g.*, [14,39,1]), there are now independent ways of reaching this conclusion, making it robust to throwing away information on either supernovae 1a [72,73] or large-scale structure. The Boomerang team [1] argue that even CMB and Hubble constant measurements now favor $\Omega_\Lambda > 0$ — we have seen that this conclusion hinges on the the important additional assumption of negligible gravity ways.

C. Neutrino density

For the neutrino density ω_ν, we are still a far cry from the grand goal of a precision cross-check between cosmological and laboratory measurements, since the two have so far provided only upper and lower limits, respectively. However, the two limits are steadily creeping closer. Atmospheric neutrino oscillations [70] show that there is at least one neutrino (presumably ν_τ) whose mass exceeds a lower limit somewhere in the range 0.04-0.08 eV [71]. Since

$$\omega_\nu \equiv \frac{1}{93.8\mathrm{eV}} \sum_i m_i, \tag{14}$$

where m_i is the mass of the ith neutrino, this corresponds to a lower limit $\omega_\nu \gtrsim 0.0004-0.0008$, or $f_\nu \gtrsim 0.003-0.01$, Our constraints (see Figure 8) give $m_\nu < 4.4$eV, further sharpening the 5.5 eV limit from a careful analysis of previous cosmological data [74]. If atmospheric neutrino oscillations convert ν_μ into ν_τ, the mass of the τ neutrino is thus in the range $0.04 - 4.4$ eV.

How do our results compare with those of other recent analyses? The analysis most similar to ours is that of the BOOMERaNG team [1]. A detailed numerical comparison of their results with our Table 2 is very encouraging. Despite major differences in both analysis technique (priors, parameter space, marginalization method, *etc.*) and data used (that analysis was limited to BOOMERAnG and DMR data), both the central values and the error bars are very similar for most parameters when taking into account that they and we quote 1σ and 2σ errors, respectively. This indicates that what is being measured is really borne out loud and clear by the data in a way that is robust towards data selection or analysis details. Perhaps we are inevitably approaching the dreaded day when not only cosmology is consistent, but cosmologists are as well.

The authors wish to thank Angélica de Oliveira-Costa, Alexander Friedland, David Hogg, Andrew Jaffe, Josh Klein, Adam Riess, Dominik Schwarz, George Smoot, David Weinberg, and especially Andrew Hamilton for helpful comments and suggestions. Support for this work was provided by NSF grant AST00-71213, NASA grant NAG5-9194 and the University of Pennsylvania Research Foundation.

APPENDIX A: LINEARIZED MODELING OF CALIBRATION AND BEAM ERRORS

This appendix describes our modeling of the four terms of the error matrix from equation (2), assuming that all relative errors are small ($\ll 1$).

$\mathbf{N}^{(\rm meas)}$ reflects the part of the errors which are uncorrelated between the different experiments and is due to detector noise and sample variance. We approximate it by

$$\mathbf{N}_{ij}^{(\mathrm{meas})} \equiv \delta_{ij}\sigma_i^2, \tag{A1}$$

where σ_i is defined as the average of the upper and lower error bars quoted for $d_i \equiv \delta T^2$.

The last three terms in equation (2) reflect correlations between measurements due to calibration and beam errors. As in [21,31], $\mathbf{N}^{(\mathrm{ical})}$ is the part specific to a single multi-band experiment and $\mathbf{N}^{(\mathrm{scal})}$ is the part that is correlated with other experiments that are calibrated off of the same (slightly uncertain) source. TOCO, MSAM, CBI and BOOMERaNG all calibrate off of Jupiter. To be conservative, we assume that the full 5% calibration uncertainty from Jupiter's antenna temperature is shared by these experiment. The true correlation should be lower, since the four experiments observed Jupiter at different frequencies. The remaining multi-band experiments in Figure 1 should not have any such inter-experiment correlations. This contribution to the noise matrix is therefore

$$\mathbf{N}_{ij}^{(\mathrm{scal})} \equiv (2s_{ij})^2 y_i y_j, \tag{A2}$$

where

$$s_{ij} = \begin{cases} 5\% & \text{if } i \text{ and } j \text{ refer to a Jupiter-calibrated point,} \\ 0 & \text{otherwise.} \end{cases} \tag{A3}$$

The factor of 2 in equation (A2) stems from the fact that the percentage error on δT_i^2 is roughly twice that for δT_i as long as it is small. We simply use the observed values for y_i in this expression.

Similarly, the remaining calibration term is

$$\mathbf{N}_{ij}^{(\mathrm{ical})} \equiv (2r_{ij})^2 y_i y_j, \tag{A4}$$

where $r_{ij} = 0$ if i and j refer to different experiments. If band powers i and j are from the same experiment, then r_{ij} is the quoted calibration error with the source contribution s_{ij} subtracted off in quadrature. We use 10% for QMASK, 14% for Python V, 8% for Viper, 8.7% for Toco 97, 6.2% for Toco 98, 10% for QMASK, 6.4% for BOOM97, 10% for BOOM98, 8% for Maxima and 8% for DASI.

A Gaussian beam of standard deviation θ suppresses power by a factor well approximated by $e^{(\ell\theta)^2}$. Taylor expanding this expression shows that a small beam error $\Delta\theta$ causes the power to be mis-estimated by a percentage $2\ell^2\theta\Delta\theta$. Defining

$$b_i \equiv 2\frac{\Delta\theta_i}{\theta_i}\left(\ell\theta_i\right)^2 y_i, \tag{A5}$$

we can therefore write the beam error as

$$\mathbf{N}_{ij}^{(\mathrm{beam})} = b_i b_j \tag{A6}$$

when i and j refer to the same experiment, zero otherwise. We use this approximation for BOOM98, where $\theta = 12.9'/c$, $\delta\theta = 1.4'/c$, and $c = \sqrt{8\ln 2}$ is the familiar FWHM conversion factor. The other experiment reporting important beam uncertainties is Maxima, for which we use the b_i-coefficients published in [3] in place of the approximation of equation (A5).

---

[1] C. B. Netterfield, astro-ph/0104460 (2001).
[2] N. W. Halverson, astro-ph/0104489 (2001).
[3] A. T. Lee *et al.*, astro-ph/0104459 (2001).
[4] P. J. E Peebles and J. T. Yu, ApJ **162**, 815 (1970).
[5] R. Sunyaev and Ya. Zeldovich, Astrophys. & Space Sci. **7**, 3 (1970).
[6] A. E. Lange *et al.*, Phys. Rev. D **63**, 042001 (2001).
[7] M. Tegmark and M. Zaldarriaga, Phys. Rev. Lett. **85**, 2240 (2000).
[8] A. Balbi *et al.*, ApJL **545**, L1 (2000).
[9] W. Hu, M. Fukugita, M. Zaldarriaga, and M. Tegmark, ApJ **549**, 669 (2001).
[10] A. Jaffe *et al.*, Phys. Rev. Lett. **86**, 3475 (2000).
[11] W. H. Kinney, A. Melchiorri, and A. Riotto, Phys.Rev. D **63**, 23505 (2001).
[12] S. Hannestad, S. H. Hansen, F. L. Villante, and A. J. S Hamilton, astro-ph/0103047 (2001).
[13] J. R. Bond *et al.*, astro-ph/0011378 (2000).
[14] N. Bahcall, J. P. Ostriker, S. Perlmutter, and P. J. Steinhardt, Science **284**, 1481 (1999).
[15] B. Novosyadlyj, R. Durrer, S. Gottlöber, V. N. Lukash, and S. Apunevych, A&A **356**, 418 (2000).
[16] R. Durrer and B. Novosyadlyj, astro-ph/0009057 (2000).
[17] S. L. Bridle *et al.*, MNRAS **321**, 333 (2001).
[18] S. Padin *et al.*, ApJL **549**, L1 (2001).
[19] C. Pryke *et al.*, astro-ph/0104490 (2001).
[20] R. Stompor *et al.*, astro-ph/010506 (2001).
[21] M. Tegmark and M. Zaldarriaga, ApJ **544**, 30 (2000).
[22] Y. Xu, M. Tegmark, A. de Oliveira-Costa, M. Devlin, T. Herbig, A. D. Miller, C. B. Netterfield, and L. A. Page 2001, astro-ph/0010552, *Phys. Rev. D*, in press.
[23] C. B. Netterfield, M. J. Devlin, N. Jarosik, L. A. Page, and E. J. Wollack; ApJ, **474**, 47 (1997)
[24] M. Devlin, A. de Oliveira-Costa, T. Herbig, A. D. Miller, C. B. Netterfield, L. A. Page, and M. Tegmark, ApJL **509**, L77 (1998).
[25] T. Herbig *et al.*, ApJL **509**, L73 (1998).
[26] A. de Oliveira-Costa, M. Devlin, T. Herbig, A. D. Miller, C. B. Netterfield, L. A. Page, and M. Tegmark, ApJL **509**, L77 (1998).
[27] E. Gawiser and J. Silk, Phys. Rept. **333-334**, 245 (2000).
[28] C. L. Bennett *et al.*, ApJ **464**, L1 (1996).
[29] D. Scott, J. Silk, and M. White, Science **268**, 829 (1995).
[30] E. Pierpaoli, D. Scott, and M. White, Mod. Phys. Lett. A **15**, 1357 (2000).
[31] L. Knox and L. Page, Phys. Rev. Lett. **85**, 1366 (2000).
[32] J. R. Bond, A. H. Jaffe, and L. E. Knox, ApJ **533**, 19

(2000).
[33] M. Tegmark, ApJ **480**, L87 (1997).
[34] E. L. Wright, astro-ph/9612006 (1996).
[35] M. Tegmark, ApJ **519**, 513 (1999).
[36] Y. Xu, M. Tegmark, and A. de Oliveira-Costa, astro-ph/0104419 (2001).
[37] M. Tegmark and B. C. Bromley, ApJL **518**, L69 (1999).
[38] M. Seaborne *et al.*, MNRAS **309**, 89 (1999).
[39] M. Tegmark, M. Zaldarriaga, and A. J. S Hamilton, Phys. Rev. D **63**, 43007 (2001).
[40] W. H. Press, astro-ph/9604126: (1996).
[41] W. Saunders *et al.*, MNRAS **317**, 55 (2000).
[42] A. J. S Hamilton, M. Tegmark, and N. Padmanabhan, astro-ph/0004334 (2000).
[43] A. C. Croft *et al.*, astro-ph/0012324 (2000).
[44] A. J. S Hamilton, astro-ph/0006089 (2000).
[45] S. Burles, K. M. Nollett, and M. S. Turner, astro-ph/0010171 (2000).
[46] W. L. Freedman *et al.*, astro-ph/0012376 (2000). B. K. Gibson 2000, private communication
[47] K. Scholberg *et al.*, hep-ex/9905016 (1999).
[48] P. de Bernardis *et al.*, Nature **404**, 955 (2000).
[49] S. S. McGaugh, ApJL **541**, L33 (2000).
[50] J. R. Bond, G. Efstathiou, and M. Tegmark, MNRAS **291**, L33 (1997).
[51] D. J. Eisenstein, W. Hu, and Tegmark M, ApJ **518**, 2 (1999).
[52] J. Martin and D. Schwarz, Phys.Rev. D **62**, 103520 (2000).
[53] A. R. Liddle and D. H. Lyth, Phys. Lett. B **291**, 391 (1992).
[54] S. D'Odorico, M. Dessauges-Zavadsky, and P. Molaro, astro-ph/0102162 (2001).
[55] M. Fukugita, C. J. Hogan, and P. J. E Peebles, ApJ **503**, 518 (1998).
[56] D. Weinberg *et al.*, ApJ **490**, 564 (1997).
[57] M. Rauch *et al.*, ApJ **489**, 7 (1997).
[58] L. Hui, Z. Haiman, M. Zaldarriaga, and T. Alexander, astro-ph/0104442 (2001).
[59] J. J. Mohr, Z. Haiman, and G. P. Holder, astro-ph/0004244 (2000).
[60] J. Diego M *et al.*, astro-ph/0104217 (2001).
[61] N. A. Bahcall *et al.*, ApJ **541**, 1 (2000).
[62] R. G. Carlberg *et al.*, ApJ **516**, 552 (1999).
[63] L. Grego *et al.*, astro-ph/0012067 (2000).
[64] N. A. Bahcall and X. Fan, ApJ **504**, 1 (1998).
[65] V. R. Eke, S. Cole, C. S. Frenk, and J. P. Henry, MNRAS **298**, 1145 (1998).
[66] J. P. Henry, ApJ **534**, 565 (2000).
[67] A. R. Liddle and P. T. P Viana, MNRAS **303**, 535 (1999).
[68] I. Zehavi and A. Dekel, Nature **401**, 252 (1999).
[69] *et al.*. Peacock, Nature **410**, 169 (2001).
[70] Y. Fukuda *et al.*, Phys. Rev. Lett. **82**, 1810 (1999).
[71] J. W. F Valle, astro-ph/0104085 (2001).
[72] S. Perlmutter *et al.*, Nature **391**, 51 (1998).
[73] A. G. Riess *et al.*, Astron. J. **116**, 1009 (1998).
[74] R. A. C Croft, W. Hu, and R. Davé, Phys. Rev. Lett. **83**, 1092 (1999).