

The information content of cosmic microwave background anisotropies

Douglas Scott,^a Dagoberto Contreras,^a Ali Narimani,^a Yin-Zhe Ma^b

^aDepartment of Physics and Astronomy,
University of British Columbia,
Vancouver, BC, Canada V6T1Z1

^bAstrophysics and Cosmology Research Unit,
School of Chemistry and Physics,
University of KwaZulu-Natal,
Durban 4041, South Africa

E-mail: dscott@phas.ubc.ca, dagocont@phas.ubc.ca, anariman@phas.ubc.ca,
ma@ukzn.ac.za

Abstract. The cosmic microwave background (CMB) contains perturbations that are close to Gaussian and isotropic. This means that its information content, in the sense of the ability to constrain cosmological models, is closely related to the number of modes probed in CMB power spectra. Rather than making forecasts for specific experimental setups, here we take a more pedagogical approach and ask how much information we can extract from the CMB if we are only limited by sample variance. We show that, compared with temperature measurements, the addition of E -mode polarization doubles the number of modes available out to a fixed maximum multipole, provided that all of the TT , TE , and EE power spectra are measured. However, the situation in terms of constraints on particular parameters is more complicated, as we explain and illustrate graphically. We also discuss the enhancements in information that can come from adding B -mode polarization and gravitational lensing. We show how well one could ever determine the basic cosmological parameters from CMB data compared with what has been achieved with *Planck*, which has already probed a substantial fraction of the TT information. Lastly, we look at constraints on neutrino mass as a specific example of how lensing information improves future prospects beyond the current 6-parameter model.

Keywords: CMB theory – cosmological parameters from CMB

¹Corresponding author.

1 Introduction

Planck [1], the *Wilkinson Microwave Anisotropy Probe* (WMAP [2]), the Atacama Cosmology Telescope (ACT [3]), the South Pole Experiment (SPT [4]) and other cosmic microwave background (CMB) experiments have measured the CMB with high sensitivity, covering angular scales from essentially the whole sky down to the arcminute regime [5]. It is well known that the precision with which these measurements are made (together with an understanding of the physics generating the anisotropies) allows us to place very tight constraints on cosmological parameters. A large number of studies have focused on predicting how well the parameters can be constrained using existing or future CMB data [6–15]. In this paper we want to take a rather more global view, and ask just how much constraining power there is to be mined, or in other words, how much cosmological information there is to extract from the CMB anisotropies. This will lead us to address questions like:

- How is the overall information content related to the number of CMB modes measured?
- What does polarization add to temperature information?
- How do power spectrum measurements relate to constraints on parameters?
- Is the information finite, and how far have we progressed towards the goal of measuring all that is available?

The basic 2015 *Planck* data set (including fairly conservative masking of the sky, as well as fitting of foreground signals) gives a temperature power spectrum that is measured to approximately 750σ and polarization-related power spectra that are measured to around 280σ [1]. We might therefore naively expect that there is about 800σ worth of constraints to be shared out among the cosmological parameters. However, a check of the *Planck*-derived constraints shows that this same data set yields a value for $100\theta_*$ (a parameterization of the ratio of the sound horizon to the last-scattering surface distance) that is 1.04103 ± 0.00046 [16], which corresponds to an almost 2300σ measurement of θ_* . At the same time, the constraints on the other five parameters in the usual set give a quadrature sum of about 120σ (which is fairly negligible compared to the θ_* constraint).

The original motivation for this paper was to ask “how is it that an 800σ measurement of anisotropy power leads to a roughly 2300σ combined constraint on cosmological parameters?” In attempting to answer this question, we hope to illuminate some issues concerning cosmological parameter constraints in general, and how they might relate to experimental design in the future.

In this paper we focus on the conventional cosmological-constant-dominated cold dark matter model, Λ CDM, and a standard 6-parameter set of cosmological parameters: A_s , the amplitude of the initial power spectrum; n_s , the power-law slope of the initial conditions; $\Omega_b h^2$, the baryonic density; $\Omega_c h^2$, the cold dark matter density; θ_* , which we have already defined; and τ , the optical depth to reionization. Here h is the Hubble parameter today, H_0 , in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Within this model we use the code `CAMB` [17] to calculate CMB power spectra.

2 CMB anisotropy information

The word “information” has many different meanings. Here, we use the word to mean the strength of our ability to constrain cosmologies. The total amount of information available

for cosmological surveys is related to the number of observable modes [18, 19]. This situation has been described in many papers related to measuring the 3-dimensional power spectrum in order to constrain cosmological parameters [20–24].

The situation for CMB temperature anisotropies is simpler, since it only involves assessing the information content of a purely 2-dimensional sky. On the other hand, as we shall see the relationship between the modes and the constraints on cosmological parameters is non-trivial.

Let us start by recalling that the temperature field $T(\theta, \phi)$ on the sky is usually expanded in terms of spherical harmonics, i.e.,

$$T(\theta, \phi) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{+\ell} a_{\ell m} Y_{\ell m}(\theta, \phi), \quad (2.1)$$

where we have removed the monopole (average CMB temperature) and dipole (which is dominated by our local velocity). Then, provided one goes to sufficiently high multipoles, and ignoring effects of beams and masking, one can use the set of $a_{\ell m}$ s as an alternative representation of the pixels in the map. The power spectrum C_ℓ is the expectation value of the variance of the $a_{\ell m}$ s as a function of ℓ , with each m being equivalent since there are no cosmologically preferred directions.

If the perturbations are Gaussian then each sky is a realization of this power spectrum. The scatter among these realizations is known as “cosmic variance.” The cosmic variance in estimates of the C_ℓ s on the full sky is

$$\Delta C_\ell = \sqrt{\frac{2}{(2\ell + 1)}} C_\ell \quad (2.2)$$

(e.g., Ref. [25]). The factor of 2 here is because this is effectively the “variance of the variance,” and for a Gaussian distribution that is twice the variance. The factor of $(2\ell + 1)$ is the number of m modes for each ℓ and if only a fraction f_{sky} is observed, then the approximate effect is to increase the uncertainty in C_ℓ so that the “sample variance” is f_{sky}^{-1} larger [26]. A more realistic expression can also be written that includes the instrumental noise and beam, as described in Ref. [27].

For the simple case of an all-sky, noise-free experiment, which measures multipoles perfectly up to ℓ_{max} , the total square of the signal-to-noise ratio (using Eq. 2.2) is

$$\begin{aligned} (S/N)^2 &\equiv \sum_{\ell=2}^{\ell_{\text{max}}} (C_\ell / \Delta C_\ell)^2 = \frac{1}{2} \sum_{\ell=2}^{\ell_{\text{max}}} (2\ell + 1) \\ &= \frac{1}{2} [\ell_{\text{max}}(\ell_{\text{max}} + 2) - 3]. \end{aligned} \quad (2.3)$$

Note that this calculation is *exactly half* the total number of modes, i.e., $\sum(2\ell + 1)$. This means that in terms of constraints on the power spectrum, each $a_{\ell m}$ mode contributes 1/2 to the square of the total signal-to-noise ratio (SNR). In other words, estimating the information in the power spectrum is effectively the same thing as counting modes.

To be clear, we are distinguishing here between trivial information that tells us about the particular realization of our Universe (which we can continue to measure as precisely as we wish) and the more valuable information we can extract from our Hubble patch, which

gives us constraints on the background cosmological model. The fact that the CMB sky is remarkably close to Gaussian allows us to reduce the information contained in the individual $a_{\ell m}$ s (or equivalently in the particular hot and cold spots on our sky) to estimates of the power at each ℓ (or equivalently the variance among pixels as a function of angular separation). For Gaussian skies, the amount of information in the power spectrum is directly proportional to the number of independent modes that can be measured.

Certainly one *could* regard all the data in a map as being “information,” i.e., the fact that there is a CMB hot spot in a particular direction is of some consequence, just as it matters that we live in the Milky Way galaxy, rather than M31. However, we discount these particulars about our realization, since they do not tell us about our overall cosmological model. So here, when we say “information,” we are referring to the constraining power for cosmological models, or more specifically the signal-to-noise ratio. Clearly there is a relationship between this “information” and the amount of computer memory required to store the related data, i.e., the number of bits needed. In a more formal information theoretical sense, the number of bits required corresponds to the base 2 logarithm of what we are defining as information (e.g., Ref. [28]) – but here we are talking about signal-to-noise ratio for power spectra, since that is where parameter constraints come from.

For the case of the CMB each $a_{\ell m}$ is a random number coming from a Gaussian distribution, with mean zero and variance C_ℓ . To estimate the C_ℓ from a set of observed $a_{\ell m}$ s, it is sufficient to have only a very few bits of information for each $a_{\ell m}$, since we only need each $a_{\ell m}$ to help us obtain an estimate of the variance. Hence the amount of cosmological information – i.e., the ability to eventually constrain parameters – is determined by the number of measurable modes times a (roughly) constant (but parameterization dependent) numerical factor. However, for a particular parameter it may be that some modes are more important than others; to understand how the information of power spectra maps onto the information of the cosmological model, we can perform a more rigorous calculation by considering the Fisher matrix, as we do in the following sections.

3 CMB Fisher information

The previous section considered CMB temperature anisotropies only, but since the CMB sky can be linearly polarized, there exists additional information in each pixel of a CMB map. If an experiment can measure the Q and U Stokes parameters (in addition to T), then in principle there are two additional pieces of information for each pixel. The most useful way of describing these additional degrees of freedom is through a geometrical approach, defining a divergence-like combination, usually called “ E ,” and a curl-like combination, usually called “ B ” [29–31]. T , E , and B fields can be used to determine auto- and cross-power spectra, and with parity considerations (at least for cosmological signals) making TB and EB uncorrelated, we are left with four CMB power spectra from which we can constrain parameters, namely C_ℓ^{TT} , C_ℓ^{TE} , C_ℓ^{EE} , and C_ℓ^{BB} .

The T - and E -mode maps have now been well measured [32–38], but estimates of B -modes are still in their infancy [39]. Moreover, even when primordial B -modes become detectable, we expect them to be small, and hence one will need an experiment with an entirely different sensitivity range to probe the information contained in those modes. For these reasons we will neglect B -modes in most of the discussion of this paper. However, the B -modes caused by the effects of gravitational lensing have now been detected by several experiments [38, 40–44], and we will discuss this later in Section 4.

Focusing only on T and E , an important fact is that the two fields on the sky are *not* independent, but contain correlations, which can be measured through the cross-power spectrum C_ℓ^{TE} . Because of this, it may not be entirely obvious how much additional information is provided by measurements of CMB polarization – does a measurement of T and E provide twice as much information as provided by each of them alone? We will answer this question in the following subsections.

3.1 The Fisher matrix

The Fisher matrix gives a powerful formalism for describing the information content coming from observables in terms of underlying parameters (see e.g., Ref. [45]). Under the assumption of Gaussian perturbations and with negligible instrumental noise, the Fisher information matrix for CMB temperature and polarization anisotropies is [46]

$$F_{ij} = \sum_\ell \sum_{XY} \frac{\partial C_\ell^X}{\partial p_i} (\mathbb{C}_\ell)^{-1}_{XY} \frac{\partial C_\ell^Y}{\partial p_j}, \quad (3.1)$$

where C_ℓ^X and C_ℓ^Y are the power in the ℓ th multipole for $X, Y = T, E$, or C (temperature, E -mode polarization, and TE correlation, respectively), and the p_i are cosmological parameters. Here we are ignoring the B -modes, as already explained; however, in principle one could easily extend Eq. (3.1) to include them.

We now define the vector \vec{x}_ℓ as

$$\vec{x}_\ell = \begin{pmatrix} C_\ell^{TT} \\ C_\ell^{EE} \\ C_\ell^{TE} \end{pmatrix}, \quad (3.2)$$

and \mathbb{C}_ℓ as a covariance matrix

$$\mathbb{C}_\ell = \begin{pmatrix} (\mathbb{C}_\ell)_{TT} & (\mathbb{C}_\ell)_{TE} & (\mathbb{C}_\ell)_{TC} \\ (\mathbb{C}_\ell)_{TE} & (\mathbb{C}_\ell)_{EE} & (\mathbb{C}_\ell)_{EC} \\ (\mathbb{C}_\ell)_{TC} & (\mathbb{C}_\ell)_{EC} & (\mathbb{C}_\ell)_{CC} \end{pmatrix}. \quad (3.3)$$

We can then formulate Eq. (3.1) into a matrix product as

$$F_{ij} = \sum_\ell \frac{\partial \vec{x}_\ell^\top}{\partial p_i} (\mathbb{C}_\ell)^{-1} \frac{\partial \vec{x}_\ell}{\partial p_j}, \quad (3.4)$$

where the $(\mathbb{C}_\ell)_{XY}$ entries for the noise-free case are

$$\begin{aligned} (\mathbb{C}_\ell)_{TT} &= \frac{2}{(2\ell+1)f_{\text{sky}}} (C_\ell^{TT})^2, \\ (\mathbb{C}_\ell)_{EE} &= \frac{2}{(2\ell+1)f_{\text{sky}}} (C_\ell^{EE})^2, \\ (\mathbb{C}_\ell)_{TE} &= \frac{2}{(2\ell+1)f_{\text{sky}}} (C_\ell^{TE})^2, \\ (\mathbb{C}_\ell)_{TC} &= \frac{2}{(2\ell+1)f_{\text{sky}}} (C_\ell^{TT} C_\ell^{TE}), \\ (\mathbb{C}_\ell)_{EC} &= \frac{2}{(2\ell+1)f_{\text{sky}}} (C_\ell^{EE} C_\ell^{TE}), \\ (\mathbb{C}_\ell)_{CC} &= \frac{1}{(2\ell+1)f_{\text{sky}}} \left[C_\ell^{EE} C_\ell^{TT} + (C_\ell^{TE})^2 \right], \end{aligned} \quad (3.5)$$

following Ref. [46]. For convenience we will define $\mathcal{N} \equiv f_{\text{sky}}(2\ell+1)/2$ and $r_\ell \equiv C_\ell^{TE}/\sqrt{C_\ell^{TT}C_\ell^{EE}}$, which is the correlation coefficient between T and E (see appendix A6 in Ref. [36]).

The Cramer-Rao bound states that we can assign the 1σ statistical uncertainties to be $\sigma_i = \sqrt{(F^{-1})_{ii}}$, this gives the smallest possible errors achievable. Now we will consider the simplest situation, where there is only one cosmological parameter to determine. This scenario, though simple, will explain the effects of including polarization along with temperature data, given that the *maps* are correlated. In Section 5 we will consider larger parameter sets which will explain the effects of correlations between *parameters*.




Figure 1. Square of the signal-to-noise ratio on an ℓ -by- ℓ basis for the usual six parameters of the standard Λ CDM model, using the best-fit *Planck* cosmology as the fiducial model. The curves are for a hypothetical noise-free, $f_{\text{sky}} = 0.5$ experiment. The results for the TT spectrum (see Eq. 3.7) are in red, EE (Eq. 3.10) in blue, and TE (Eq. 3.12) in green. To highlight the low- ℓ region, we have plotted the first 30 multipoles using a logarithmic x -axis (and drawn a vertical line to separate the log and linear parts). For reference, solid vertical lines mark the positions of peaks in the TT power spectrum (approximately the same positions as troughs in EE) and dashed vertical lines are troughs in TT (or peaks in EE); the numerical values of the peaks and troughs are taken from Ref. [36]. Note the different ranges plotted on the y -axis of each panel.

3.2 TT only

Let us first consider the case in which only the CMB temperature is mapped, i.e., we have not measured Q and U in this scenario. If we focus on a single parameter p , then the SNR can be written as

$$\left(\frac{p}{\Delta p}\right) = p F_{pp}^{1/2}, \quad (3.6)$$

where we sum over all the multipoles ℓ . For the temperature power spectrum alone, then the covariance matrix just has a single entry and we can write the squared SNR at each multipole as

$$\left(\frac{S^2}{N^2}\right)_\ell^{TT} = \mathcal{N} \left(\frac{d \ln C_\ell^{TT}}{d \ln p}\right)^2. \quad (3.7)$$

If we have a simple scaling parameter (similar to A_s) then $\partial C_\ell^X / \partial p = C_\ell^X / p$ and the logarithmic derivative is unity, so we just have

$$\begin{aligned} \left(\frac{p}{\Delta p}\right)^{TT} &= \left(\sum_\ell \mathcal{N}\right)^{1/2} \\ &= \left(f_{\text{sky}} \sum_\ell \left[\ell + \frac{1}{2}\right]\right)^{1/2} \end{aligned} \quad (3.8)$$

$$\simeq \frac{1}{\sqrt{2}} f_{\text{sky}}^{1/2} \ell_{\text{max}}. \quad (3.9)$$

Hence we see that in the simple case of a single parameter that is proportional to the amplitude of the power spectrum, the constraining power on this parameter from the TT power spectrum is the same as the mode counting presented in the previous section.

Things are more complicated when we consider general parameters. Figure 1 shows the squared SNR per multipole for the six parameters of the standard Λ CDM model following Eq. (3.7), i.e., with all other parameters held fixed.¹ The red curves in each panel show the TT -only case. We can see many details from this curve, illustrating how different multipole ranges affect constraints on each parameter.

Variations due to the parameter A_s are similar to those we have described for a ‘‘scaling parameter,’’ but not quite the same. If we looked at the TT spectrum without the effects of gravitational lensing, then we would find that the spectrum scales exactly like $A_s e^{-2\tau}$, and hence this is the scaling parameter we referred to in Eq. (3.7). Since for Fig. 1 all of the parameters are held fixed except each one individually, then both A_s and $e^{-2\tau}$ serve effectively as scaling parameters in the sense of Eq. (3.7). Because of this, the curve for τ looks like the one for A_s , multiplied by $4\tau^2$ (since $d \ln C_\ell^{TT} / d \ln \tau = -2\tau$), except for an extra variation at the lowest ℓ s. When we include lensing in the usual way (as we have done in Fig. 1) then there are also small wiggles in the A_s curve, which come from the smoothing effect of lensing on the peaks and troughs.

For the slope, n_s , we see the effect of the ‘‘pivot’’ point at $k = 0.05 \text{ Mpc}^{-1}$, which projects to $\ell \simeq 550$ for TT ; at multipoles around this point there is no constraint on n_s . The variations for most of the other parameters reflect the structure of the C_ℓ s themselves. For example, for

¹Gaussianity is not a valid assumption for the low multipoles, and hence the low- ℓ part of the plots should be taken as approximate estimates only.

θ_* the curve goes to zero near the positions of peaks and troughs (because the gradient of C_ℓ is zero there).

The A_s panel in Figure 1 can serve as a guideline for inferring the sensitivity of a C_ℓ to a particular parameter. One can approximately define the dependency of C_ℓ to a generic parameter p , by a polynomial as $C_\ell \propto p^\alpha$, with $\alpha = 1$ then representing linear dependence, and we might speak of “less than linear” and “non-linear” as corresponding to $\alpha < 1$ and $\alpha > 1$, respectively. Comparing each of the panels of Figure 1 with A_s , shows that the set $\{\Omega_b h^2, \Omega_c h^2, \tau\}$ all have a less than linear relation with the C_ℓ s at $\ell \lesssim 30$, while n_s exhibits a non-linear relationship. The parameter θ_* is non-linear over almost the entire multipole range (note the different y -axis range for this panel). In the high- ℓ region, $\ell > 30$, all of the parameters (except θ_*) show a mildly less than linear relation, with oscillations close to zero at some multipoles (and this is even true for θ_*), while n_s has a close to linear relation at $\ell \gtrsim 1500$.

We will discuss the blue and green (EE and TE) curves in the following subsections.

3.3 EE only

Now consider the situation where only the E -modes are mapped, and hence we only have access to the EE power spectrum for constraining cosmology. Here the situation is clearly exactly the same as it was for the TT -only case.

We have the squared SNR for a single parameter p being

$$\left(\frac{S^2}{N^2}\right)_\ell^{EE} = \mathcal{N} \left(\frac{d \ln C_\ell^{EE}}{d \ln p}\right)^2 \quad (3.10)$$

and the total for a scaling parameter is

$$\left(\frac{p}{\Delta p}\right)^{EE} \simeq \frac{1}{\sqrt{2}} f_{\text{sky}}^{1/2} \ell_{\text{max}}. \quad (3.11)$$

Again, the result is just what we expect from mode counting. If we have polarization data out to some ℓ_{max} , then it provides the same constraints on a scaling parameter as having temperature data out to the same ℓ_{max} .

The situation for more general parameters is presented by the blue curves in the panels of Figure 1. One can see several effects that are similar to the TT case. The situation for A_s and τ are essentially the same as for the TT case, with dramatic improvement for τ at low ℓ because of the sensitivity of large-scale polarization to reionization. The constraining power for n_s also has a zero in the EE case, but the pivot k_0 projects to a slightly different ℓ (reflecting the slightly different scales with which polarization is sourced compared to temperature). For the parameters θ_* , $\Omega_b h^2$, and $\Omega_c h^2$ we see that the constraining power from EE is generally higher than that for TT (as recently pointed out in Ref. [15] and explained in the next section). This illustrates the improved parameter constraints from polarization, essentially because of the sharper acoustic features in polarization.

3.4 TE correlation only

Now let us examine the case where we measure the TE power spectrum only. This is clearly not a realistic situation, but evaluating it will elucidate some interesting points. The Fisher

matrix again has a single entry, coming from the term $(\mathbb{C}_\ell)_{CC}$, and for a single parameter p , we obtain

$$\left(\frac{S^2}{N^2}\right)_\ell^{TE} = 2\mathcal{N} \left(\frac{d \ln C_\ell^{TE}}{d \ln p}\right)^2 \frac{r_\ell^2}{1+r_\ell^2}. \quad (3.12)$$

As before we have defined $r_\ell \equiv C_\ell^{TE}/\sqrt{C_\ell^{TT}C_\ell^{EE}}$, which is plotted in Figure 2. For a scaling parameter we can simply replace the derivative with unity.

For a general parameter p there can be a different situation than we saw for TT and EE . In the TE power spectrum the amount of correlation *and* (possibly surprisingly) how it changes under the influence of p is important. To see this we can re-write the square of the SNR of TE in terms of r , TT , and EE as

$$\begin{aligned} \left(\frac{S^2}{N^2}\right)_\ell^{TE} = 2\mathcal{N} \frac{r_\ell^2}{1+r_\ell^2} \left[\left(\frac{d \ln r_\ell}{d \ln p}\right)^2 + \frac{d \ln r_\ell}{d \ln p} \left(\frac{d \ln C_\ell^{TT}}{d \ln p} + \frac{d \ln C_\ell^{EE}}{d \ln p}\right) \right. \\ \left. + \frac{1}{4} \left(\frac{d \ln C_\ell^{TT}}{d \ln p} + \frac{d \ln C_\ell^{EE}}{d \ln p}\right)^2 \right]. \end{aligned} \quad (3.13)$$

We can then look at the two limiting cases $r_\ell \rightarrow \pm 1, 0$:

$$\lim_{r_\ell \rightarrow \pm 1} \left(\frac{S^2}{N^2}\right)_\ell^{TE} = \frac{\mathcal{N}}{4} \left(\frac{d \ln C_\ell^{TT}}{d \ln p} + \frac{d \ln C_\ell^{EE}}{d \ln p}\right)^2; \quad (3.14)$$

$$\lim_{r_\ell \rightarrow 0} \left(\frac{S^2}{N^2}\right)_\ell^{TE} = 2\mathcal{N} \left(\frac{dr_\ell}{d \ln p}\right)^2. \quad (3.15)$$

The first thing to notice is that in the limit of full correlation the information in TE is directly given by the information content of TT and EE , as expected. However, we also see that even in the case of a vanishing r_ℓ , there is still information to be obtained from measuring TE (this is because the SNR does *not* vanish as $r_\ell \rightarrow 0$). This will depend on the behaviour of the parameter p and specifically on how r_ℓ varies as the parameter changes. For a scaling parameter the information content of TE vanishes when r_ℓ vanishes, but this is not true in general.

Details of the parameter SNR² values per multipole are shown in Figure 1, with the green curves being for the TE case. Once again, the A_s panel is helpful in here because we can see that the dramatic drops in the green curve correspond to the points with $r_\ell = 0$. Checking the green curve value of these locations at other panels, such as θ_* , reveals that one gains information by measuring TE even if $r_\ell = 0$ at a specific angular scale (e.g., see Figure 1 just below $\ell = 1000$).

The TE curves are usually lower than the TT and EE curves, simply because the T - and E -modes are only partially correlated. This is manifested particularly strongly in the τ panel of Figure 1, where it is clear that TE is much less sensitive to the reionization bump than EE . However, we see that TE can be higher in some multipole ranges for some parameters, particularly for $\Omega_b h^2$ and $\Omega_c h^2$. For example, TE is more sensitive to $\Omega_b h^2$ than TT in the range $670 \lesssim \ell \lesssim 1920$ and more sensitive than EE in the range $810 \lesssim \ell \lesssim 1830$. In fact Ref. [15] already pointed out that TE can constrain $\Omega_c h^2$ *better* than TT by around 15% – this would be hard to determine directly from Figure 1, since the figure does not account for correlations among parameters (although this is something we do consider in Section 5).

3.5 TT and EE power spectra, no correlation

We would like to understand the basic way that polarization information combines with temperature information. So let us now consider the simple (although hypothetical) situation in which we have mapped out E -mode polarization, but when this polarization is *uncorrelated* with temperature anisotropies. Under these conditions the covariance matrix of Eq. (3.3) takes on a simple 2×2 form. In this case it is easy to verify that the information is exactly doubled compared with the temperature-only case:

$$\left(\frac{S^2}{N^2}\right) = \sum_{\ell} \left[\left(\frac{S^2}{N^2}\right)_{\ell}^{TT} + \left(\frac{S^2}{N^2}\right)_{\ell}^{EE} \right]. \quad (3.16)$$

Therefore for a scaling parameter the total SNR is

$$\begin{aligned} \left(\frac{p}{\Delta p}\right) &= pF_{pp}^{1/2} \\ &\simeq f_{\text{sky}}^{1/2} \ell_{\text{max}}, \end{aligned} \quad (3.17)$$

i.e., we obtain a factor of $\sqrt{2}$ improvement over the temperature-only case. This makes sense, because uncorrelated E -mode polarization is adding precisely one additional piece of information for every pixel on the sky, or equivalently, is adding an independent set of $a_{\ell m}^E$ modes to the $a_{\ell m}^T$ modes.

3.6 TT and EE correlated, but ignoring TE

The situation in the previous subsection is of course not realistic, because in reality there is a TE correlation in the CMB anisotropies, and hence there are *three* distinct power spectra to determine, C_{ℓ}^{TT} , C_{ℓ}^{EE} , and C_{ℓ}^{TE} . So how does this affect the total information content?

To see how this works, let us first of all imagine that although the temperature and polarization fields are correlated, we have *not* measured this correlation (we are imagining an impractical scenario here where the cosmologist has been careless and ignored TE). We will again treat the simple case of a scaling parameter. If we just use the 2×2 matrix in the upper left part of Eq. (3.3), we find

$$\begin{aligned} F_{pp} &= \frac{2}{p^2} f_{\text{sky}} \sum_{\ell=2}^{\ell_{\text{max}}} \left[\left(\ell + \frac{1}{2}\right) \frac{1}{1 + r_{\ell}^2} \right] \\ &\simeq \frac{1}{p^2} f_{\text{sky}} (\ell_{\text{max}}^2 M). \end{aligned} \quad (3.18)$$

Here we have defined

$$M \equiv \frac{\sum_{\ell=2}^{\ell_{\text{max}}} [(\ell + \frac{1}{2}) / (1 + r_{\ell}^2)]}{\frac{1}{2} \ell_{\text{max}}^2}, \quad (3.19)$$

giving the ratio between this case and the (previously considered) case where T and E are uncorrelated. Since $0 < r_{\ell}^2 < 1$, it is clear that $1/2 < M < 1$ (at least for sufficiently high ℓ). The limit $M = 1/2$ (equivalent to $r_{\ell}^2 = 1$) corresponds to a perfect correlation between temperature and polarization and is therefore the same as the temperature-only case. On the other hand, $M = 1$ (corresponding to $r_{\ell}^2 = 0$), is when T and E are uncorrelated.

The signal-to-noise ratio for our hypothetical scaling parameter p when C_ℓ^{TE} is unmeasured is just

$$\begin{aligned} \left(\frac{p}{\Delta p} \right) &= p F_{pp}^{1/2} \\ &\simeq f_{\text{sky}}^{1/2} \ell_{\text{max}} M^{1/2}. \end{aligned} \quad (3.20)$$




Figure 2. Various quantities related to the correlations between the T and E -modes in the CMB sky, for best-fit 6-parameter Λ CDM model found by *Planck* [16]. The red line shows the correlation coefficient $r_\ell \equiv C_\ell^{TE}/(C_\ell^{TT}C_\ell^{EE})^{1/2}$, while the green line shows the fractional polarization, defined as $s_\ell \equiv (C_\ell^{EE}/C_\ell^{TT})^{1/2}$. The blue line is the quantity M (defined in Eq. 3.19), which is effectively the fraction of the information lost by not observing C_ℓ^{TE} (along with C_ℓ^{TT} and C_ℓ^{EE}), plotted as a function of ℓ_{max} .

For the standard Λ CDM cosmology, we show the correlation coefficient r_ℓ in Figure 2, illustrating that for some multipoles the magnitude of the correlation between T and E can be as high as 60%. We also plot $1 - M$ as a function of ℓ_{max} ; this effectively shows the amount of information that would be lost by neglecting the TE cross-correlation power spectrum. The plot shows that this can be as much as 20%, but for high ℓ_{max} is a little under 10%.

3.7 Full TT , EE , and TE

Now let us consider the case when C_ℓ^{TT} , C_ℓ^{EE} , and C_ℓ^{TE} are *all* measured (for a simple scaling parameter first). We then need to consider the full 3×3 covariance matrix when calculating the Fisher information and we recover

$$\left(\frac{p}{\Delta p} \right) \simeq f_{\text{sky}}^{1/2} \ell_{\text{max}}, \quad (3.21)$$

i.e., the same as in the situation where T and E are assumed (unrealistically) to be uncorrelated.

This can be seen explicitly by inverting Eq. (3.3):

$$\begin{aligned}\mathbb{C}_\ell^{-1} &= \frac{2\mathcal{N}}{|C|^2} \times \begin{pmatrix} \frac{1}{2}C_\ell^{EE^2} & \frac{1}{2}C_\ell^{TE^2} & -C_\ell^{EE}C_\ell^{TE} \\ \frac{1}{2}C_\ell^{TE^2} & \frac{1}{2}C_\ell^{TT^2} & -C_\ell^{TT}C_\ell^{TE} \\ -C_\ell^{EE}C_\ell^{TE} & -C_\ell^{TT}C_\ell^{TE} & C_\ell^{TT}C_\ell^{EE} + C_\ell^{TE^2} \end{pmatrix} \\ &= \frac{\mathcal{N}}{|C|^2} C_\ell^{TT} C_\ell^{EE} \begin{pmatrix} s_\ell^2 & r_\ell^2 & -2r_\ell s_\ell \\ r_\ell^2 & s_\ell^{-2} & -2r_\ell/s_\ell \\ -2r_\ell s_\ell & -2r_\ell/s_\ell & 2(1+r_\ell^2) \end{pmatrix},\end{aligned}\quad (3.22)$$

where the quantity s_ℓ has been introduced as the ratio of the polarization to temperature anisotropy and we also define \mathcal{N} and $|C|$ to simplify the expression above, specifically with

$$\mathcal{N} \equiv \frac{(2\ell+1)f_{\text{sky}}}{2}, \quad (3.23)$$

$$|C| \equiv C_\ell^{TT}C_\ell^{EE} - C_\ell^{TE^2} \quad (3.24)$$

$$= C_\ell^{TT}C_\ell^{EE} (1 - r_\ell^2), \quad (3.25)$$

$$s_\ell^2 \equiv \frac{C_\ell^{EE}}{C_\ell^{TT}}. \quad (3.26)$$

The quantity s_ℓ is also plotted in Figure 2, showing that polarization anisotropies are only a few percent of temperature anisotropies at large angular scales, and asymptote to a value close to 20% at higher multipoles.

Following the algebra introduced above, the Fisher matrix becomes

$$F_{pp} = \frac{1}{p^2} \sum_\ell \begin{pmatrix} C_\ell^{TT} \\ C_\ell^{EE} \\ C_\ell^{TE} \end{pmatrix}^\top \mathbb{C}_\ell^{-1} \begin{pmatrix} C_\ell^{TT} \\ C_\ell^{EE} \\ C_\ell^{TE} \end{pmatrix} \quad (3.27)$$

$$= \frac{1}{p^2} \sum_\ell C_\ell^{TT} C_\ell^{EE} \begin{pmatrix} s_\ell^{-1} \\ s_\ell \\ r_\ell \end{pmatrix}^\top \mathbb{C}_\ell^{-1} \begin{pmatrix} s_\ell^{-1} \\ s_\ell \\ r_\ell \end{pmatrix} \quad (3.28)$$

$$= \frac{1}{p^2} \sum_\ell \frac{\mathcal{N}}{(1-r_\ell^2)^2} \begin{pmatrix} s_\ell^{-1} \\ s_\ell \\ r_\ell \end{pmatrix}^\top \begin{pmatrix} s_\ell(1-r_\ell^2) \\ s_\ell^{-1}(1-r_\ell^2) \\ -2r_\ell(1-r_\ell^2) \end{pmatrix} \quad (3.29)$$

$$= \frac{2}{p^2} \sum_\ell \mathcal{N} \frac{(1-r_\ell^2)^2}{(1-r_\ell^2)^2} \quad (3.30)$$

$$= \frac{2}{p^2} f_{\text{sky}} \sum_\ell \left(\ell + \frac{1}{2} \right), \quad (3.31)$$

which is similar to the result of the previous section, with $M = 1$. Therefore, measurements to some ℓ_{max} of the full TT , EE , and TE spectrum have *exactly* twice as much information as a temperature-only (or polarization-only) experiment. In appendix A we give a conceptually simpler derivation of this same result, by transforming to fields that are uncorrelated by construction.

What about more complicated parameters? For power spectra with arbitrary dependence on a single parameter p the total squared SNR is then

$$\begin{aligned} \left(\frac{S}{N}\right)^2 = \sum_{\ell} \frac{1}{(1-r_{\ell}^2)^2} \left\{ \left(\frac{S^2}{N^2}\right)_{\ell}^{TT} + \left(\frac{S^2}{N^2}\right)_{\ell}^{EE} + 2r_{\ell}^2 \left(\frac{S}{N}\right)_{\ell}^{TT} \left(\frac{S}{N}\right)_{\ell}^{EE} \right. \\ \left. - 4r_{\ell} \sqrt{2+2r_{\ell}^2} \left[\left(\frac{S}{N}\right)_{\ell}^{TT} + \left(\frac{S}{N}\right)_{\ell}^{EE} \right] \left(\frac{S}{N}\right)_{\ell}^{TE} + 4(1+r_{\ell}^2)^2 \left(\frac{S^2}{N^2}\right)_{\ell}^{TE} \right\}. \end{aligned} \quad (3.32)$$

In appendix A we present an interpretation of the contributions to the total SNR, which come from temperature (uncorrelated with polarization), polarization (uncorrelated with temperature), and the correlation itself, θ_{ℓ} (defined in Eq. A.3). Equation 3.32, however allows us to consider the situation at a scale ℓ for which $r_{\ell} = 0$. In this case we obtain

$$\left(\frac{S}{N}\right)_{\ell}^2 = \left(\frac{S^2}{N^2}\right)_{\ell}^{TT} + \left(\frac{S^2}{N^2}\right)_{\ell}^{EE} + 4 \left(\frac{S^2}{N^2}\right)_{\ell}^{TE}, \quad (3.33)$$

where the final term is given by 4 times Eq. (3.15). We see here (somewhat surprisingly) that the TE correlation will contribute information even when $r_{\ell} = 0$, provided that $dr_{\ell}/dp \neq 0$.

While Eqs. (3.32) and (A.12) are complete descriptions of the contributions to the information content of a single arbitrary parameter, they do not contain the details of how multiple parameters are constrained or how parameter correlations are involved, which we consider in Section 5.

3.8 Total SNR for a single parameter

To complete this section, let us be more explicit, and give a quantitative example. For a cosmic-variance-limited experiment up to $\ell_{\max} = 3000$ and covering the entire sky, the total signal-to-noise ratio in a single scaling parameter is $S/N \simeq 2100$. And if we add E -mode polarization information, also cosmic-variance-limited to the same ℓ , we obtain $S/N \simeq 3000$.

An experiment that makes ideal measurements of T and E out to some ℓ_{\max} has precisely twice as much information (i.e., constraining power for a scaling parameter) as a T -only experiment, provided that all of TT , EE , and TE are measured. A parameter with more complicated C_{ℓ} dependence, like θ_{*} for example, will have more information from EE due to its stronger contrast between peaks and troughs. Such a parameter will also have part of its constraint coming from how the correlation itself changes (an effect that cannot be seen with a simple scaling parameter).

In terms of constraints on specific parameters, we know that the situation is more complicated still. For example, polarization data are important for breaking particular degeneracies [47, 48] (especially for determining the reionization optical depth, τ), and so polarization may constrain some parameters much better than expected for simply twice as much information. We have already seen this presented in Figure 1 In Section 5 we will focus on several Λ CDM parameters and how their correlations affect parameter constraints.

4 Additional CMB information

Before investigating parameter dependence in detail, it is worth noting that some other information that could come from the CMB, in addition to the three power spectra we have been considering.

If one could measure BB to the same ℓ_{\max} then instead of the full $TT + TE + EE$ measurement giving twice as many modes as come from TT alone, we would now have *3 times* as many modes. In practice we expect primordial B -modes to be quite weak, and it is extremely unlikely that we could measure this power spectrum beyond the first few hundred multipoles (e.g., Ref. [49]). Hence B -modes are never going to add substantially to the mode count. On the other hand, *any* measurement of primordial B -modes would provide a direct constraint on the tensor-to-scalar ratio that would be better than the indirect constraints from other power spectra. Hence (as is well known) the constraints on this additional parameter are dramatically improved through better B -mode experiments.

We have been assuming that the CMB sky contains Gaussian perturbations, but we know that this cannot be exactly true. Certainly there is hope that we may one day detect non-Gaussianity from higher-order correlations in the CMB (e.g., Ref. [50]), and here polarization offers the promise of pushing the uncertainties down. However, we do not expect such a signal to give a very high SNR (at least compared to the power spectra), since the CMB is clearly very close to Gaussian.

One exception to this is that the 4-point function of the CMB sky contains correlations from the effects of gravitational lensing. This signature allows us to estimate an additional power spectrum, $C_L^{\phi\phi}$, which has already been done to $\text{SNR} \simeq 40$ by *Planck* [43]. If we could measure this power spectrum to the same maximum multipole as for the temperature and polarization power spectra, then it would add the same number of modes again. However, things are not so simple, because the index L for lensing comes from the coupling of modes at different scales, and hence making a noise-free temperature map out to ℓ_{\max} will not give a cosmic variance limited measurement of $C_L^{\phi\phi}$ to $L = \ell_{\max}$. But putting that aside, lensing will add effectively as many modes as temperature and polarization. Hence there is in principle $\sqrt{3}$ times as much information contained in a full CMB mapping experiment (plus some additional information from B -modes) as there is in a measurement purely of C_ℓ^{TT} . On the other hand, in terms of parameter constraints, the lensing power spectrum has little dependence on cosmological parameters other than amplitude [51, 52].

Since ϕ is a Gaussian random field, its power spectrum has the same statistical properties as the T or E maps. Therefore, adding lensing to the covariance matrix is similar to adding EE or TE to TT . One could also consider adding the information coming from lensed B -modes – but then one would have to account for the fact that these modes are *not* Gaussian (since they come from a convolution of E -modes with lensing ϕ -modes [53]). However, the addition of lensed B -modes provides little to no improvement on parameters if E -modes and lensing ϕ -modes are already accounted for. The reason is an intuitive one: the lensed B -modes come directly from *unlensed* E -modes and ϕ -modes. Thus adding B -modes simply double counts some combination of the E - and ϕ -modes (albeit at different scales). Thus adding the lensed B -modes only helps in the noise-dominated case. In Section 5.3 we consider including lensing ϕ -modes to our data vector (Eq. 3.1); for the reasons stated above we do not simultaneously consider lensed B -modes.

5 CMB parameter information

5.1 Relationship with overall SNR

We now discuss the connection between the information in the power spectra and the constraints on the 6-parameters of the standard Λ CDM model. One might expect the total SNR in ℓ -space to be very crudely of order the SNR in the parameter space; however, in detail we

do not expect them to be the same. This is due to the degeneracies between parameters and the sensitivity of subsets of the data to changes in specific parameters, as well as the important fact that the power spectra do not depend linearly on the parameters, as was discussed already (see Figure 1).




Figure 3. *Top:* ratio of the total SNR^2 in parameters to that in the power spectra for a cosmic-variance-limited experiment. Note that this quantity is independent of the sky coverage. The red line shows the results of a temperature-only experiment, the blue line is for a polarization-only experiment, and the green line shows the results for both temperature and polarization combined. The parameter set here consists of $\{\Omega_b h^2, \Omega_c h^2, \theta_*, \tau, A_s, n_s\}$. *Bottom:* same as the top panel, except that H_0 has been used as a parameter in place of θ_* , i.e., the parameter set here consists of $\{\Omega_b h^2, \Omega_c h^2, h, \tau, A_s, n_s\}$.

One can think of the N_p -dimensional parameter space as a data compression scheme. This compression reduces the ℓ measured power spectrum values and their uncertainties into N_p numbers and their corresponding uncertainties. The power spectra depend linearly on the parameter A_s (apart from small lensing effects), and hence if that was the only parameter, then its SNR would be the same as the SNR for the power spectra, and hence would be the same as the mode counting exercise discussed in the section 3. However, other parameters are not “linear” in this sense, and hence can be constrained better or worse than seen for the simple case of A_s .

As a simple example, if someone wanted to treat A_s^2 as a parameter instead of A_s , then the SNR would be better by a factor of 2. A more dramatic change will come from considering $\ln A_s$ rather than A_s . And in general the standard parameters affect the power spectra in ways that are fairly different from those of a scaling parameter. The most non-linear of the parameters is θ_* , as was already discussed, since fairly small changes in θ_* can result in large changes to the power spectra, because of the relative sharpness of the adiabatic peaks and troughs. The best way to understand the constraints on parameters in general is to use the Fisher matrix to investigate what happens for the standard 6-parameter cosmology. In general, however, we should not be surprised to find that the SNR values for some cosmological parameters only differ by factors of order unity from those of a linearly scaling parameter.

A comparison of the total SNR in parameters with the total SNR in the power spectra is shown in Figure 3. Here the error bars for each parameter are derived from a Fisher matrix calculation, assuming a noise-free experiment with sample variance only, with sky coverage of 50% (picked to approximately match that of *Planck*). The fiducial model used for the Fisher matrix calculation is Λ CDM with parameters equal to the best-fit values of the *Planck*-2015 “TT+lowTEB+lensing” combination [16]. The total SNR in parameter p_i is

$$\left(\frac{S}{N}\right)_i^2 = \sum_{\ell} \frac{p_i^2}{(F_{\ell}^{-1})_{ii}}, \quad (5.1)$$

and Eq. 2.3 (including an $f_{\text{sky}} = 0.5$ factor) is used to calculate the total SNR in power spectra.

Figure 3 (top panel) shows that the total SNR² in parameters can be a factor of as much as 100 larger than the total SNR² in the power spectra. However, this number is highly dependent on the set of parameters that are chosen and can vary dramatically, e.g., if θ_* is replaced by H_0 , as plotted in the bottom panel of Figure 3. The ratio of the total SNR in parameters to the SNR in the power spectra (temperature and polarization) is close to unity when H_0 is substituted for θ_* . This suggests that the combination of parameters in the bottom panel of Figure 3 is more “linearly” related to the power spectra than the set in the top panel. The behaviour seen in this figure explains the observation made in Section 1 that the SNR on θ_* from *Planck* exceeds the total SNR in the TT power spectrum.

One can learn more about parameter constraints by looking more closely at Figure 3. Firstly, the relatively poor constraints on the 6-parameter set at low multipoles is a result of parameter degeneracies, which are not broken until higher multipole data are included. Secondly, we see that adding polarization data makes a substantial improvement to the overall constraint on parameters. The more dramatic improvement in the total constraint (green lines in the figure) at low multipoles is a result of the well-known ability of CMB polarization data to break the $A_s e^{-2\tau}$ degeneracy. Thirdly, the polarization data on their own are more constraining (for the same ℓ_{max}) than the temperature data; this arises essentially because the polarization power spectra are “sharper” than for C_{ℓ}^{TT} , as has been stressed in other studies (e.g., Ref. [15]). We also see structure in these SNR curves that clearly reflects the shape of the C_{ℓ} s; we shall discuss this further in the next section.

5.2 Parameter constraints from power spectra

Although the SNR in the power spectra gets effectively shared out among the parameters, as we have seen this is only very crudely correct, and in practice the detailed constraints on the parameters will change based on many factors. In particular the parameter constraints

will depend on which power spectra are used, which multipole range is measured, and what set of parameters was chosen in the first place. We need to appreciate that some information is special for particular parameters, e.g., large-angle polarization for τ ; thus we should focus on the ℓ -range that is important for each parameter. We now describe this connection more comprehensively by showing some examples and comparing with the current uncertainties. It is important to realize that we are not intending here to make forecasts for specific experiments (with particular assumptions about beamsizes, noise, foreground contamination, etc.), since this has been done before. Instead we are asking the more general question of how good the parameter constraints could *ever* become, by comparing the ideal values with current constraints from *Planck*.

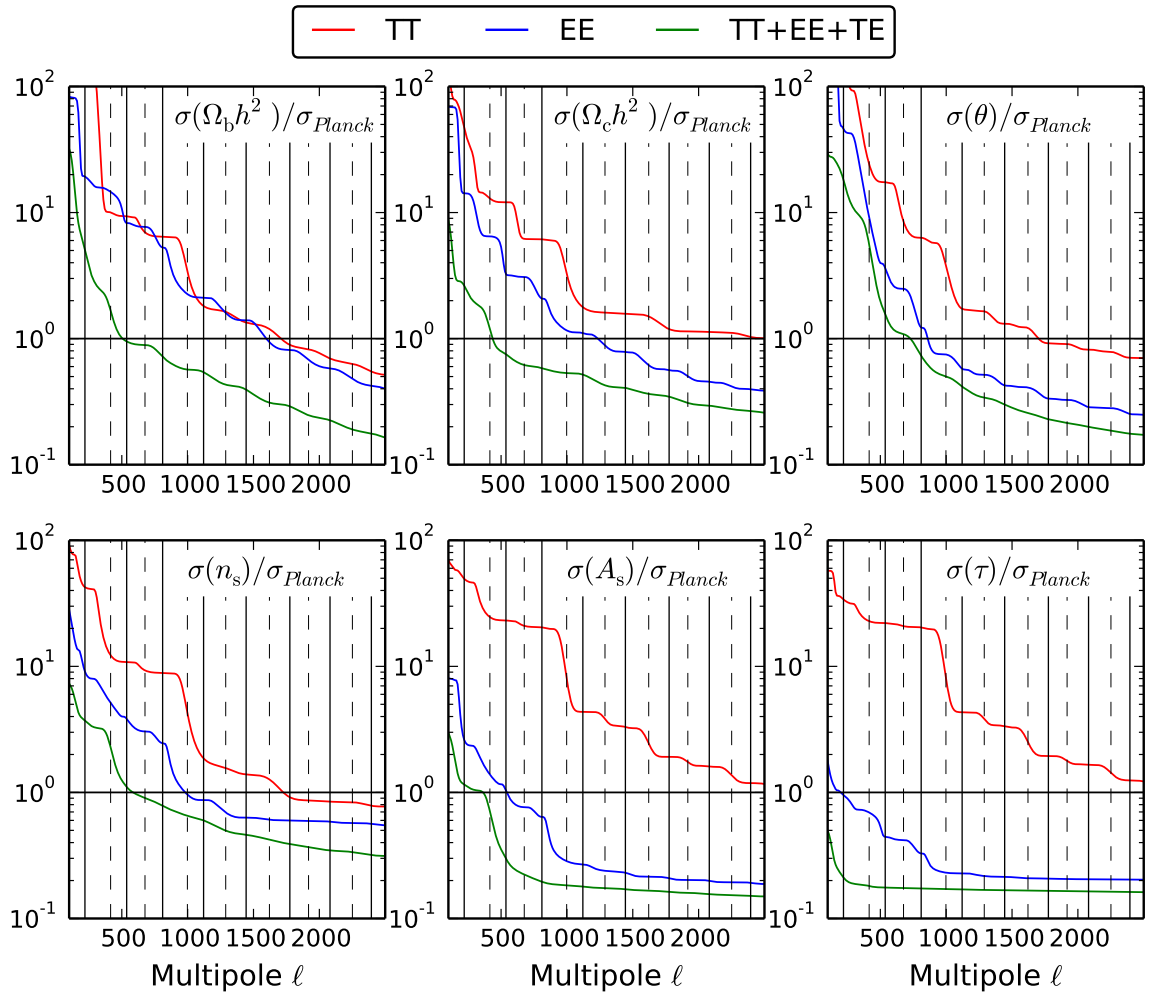


Figure 4. Cumulative Fisher forecast calculations for the six parameters of the standard Λ CDM model for a noise-free experiment with $f_{\text{sky}} = 0.5$. The red line is for a temperature-only experiment, the blue line is for pure polarization, and the green line is for the full power spectra (TT , TE , and EE). The error bars for each parameter are compared with the *Planck*-2015 “ TT +lowTEB+lensing” 68% confidence limits at different values of maximum multipole. The solid vertical lines indicate the temperature peak positions, while the dashed lines are for temperature troughs. Note that the peaks (troughs) in the TT power spectrum are in almost the same positions as the troughs (peaks) in the EE spectrum.

Figure 4 shows the results of a Fisher calculation for a cosmic-variance-limited, CMB-only experiment, where we have picked a sky coverage of 50%, which approximately matches the effective area used by *Planck* for the main parameter constraints (although this is frequency dependent, see Ref. [16]). Note that the errors would just scale as $f_{\text{sky}}^{-1/2}$ for other values. The x -axis here indicates the maximum ℓ used in the calculation. The red line is for a temperature-only experiment, while the blue and green lines are for polarization-only and the full set of three spectra, respectively. The error bars for each parameter are compared with the *Planck*-2015 “TT+lowTEB+lensing” 68% confidence limits. The vertical solid and dashed lines show the temperature peak and trough positions, respectively, which are almost the same as the E -polarization troughs and peaks [36].

The mode-counting argument (of Section 2) tells us that every $a_{\ell m}$ has equal weight for contributing to the SNR of the power spectrum, but, as we have seen, this is not true for the parameter SNR, since the C_ℓ s are not equal in delivering constraints on individual parameters. For example, while the multipoles around the third TT trough, $\ell \simeq 900$ –1100, are particularly useful for reducing the uncertainty on $\Omega_c h^2$, the information gain from the $\ell \simeq 1300$ –1650 range (from the fourth TT peak to the fifth trough) is almost negligible.

A careful examination of Figure 4 shows that in general the TT troughs are more important than the peaks for reducing the error bars (particularly clear when focusing on A_s and τ , for example). We also see (in the green lines) that when we add temperature to polarization the curves are much smoother. This is because the effects coming from the troughs and peaks of TT more or less cancel with the effects coming from the troughs and peaks of EE .

But there is still the issue to explain of why the TT troughs, which are obviously *lower* than the peaks, and therefore (one might expect) carry less information, should give stronger constraints. By comparing Figure 4 with Figure 5 – which shows the same predictions but with lensing effects turned off – one can see that the reason for the importance of the troughs is the effect of lensing on the TT spectrum. The explanation is that since lensing smooths the peaks and troughs, while preserving total power, then the *relative* change from lensing is larger around the troughs than around the peaks, and hence the troughs can give better constraints on parameters.

Comparing Figure 4 with Figure 5 also shows that while the lensing of the power spectra is useful for breaking the A_s – τ degeneracy, it makes the constraints on the set $\{\theta_*, \Omega_c h^2, \Omega_b h^2\}$ *weaker*. This is because the peaks and troughs are “sharper” in the unlensed spectra, and are therefore a better source of information for constraining θ_* . The uncertainties on the matter densities are also improved as a result of the better constraints on θ_* and the correlations among the parameters.

What we have highlighted here is the simple observation that constraints on cosmological parameters come from two basic factors: the first is having power spectra that are sensitive to changes in parameters (see Eq. 3.32 and Figure 1); the second is the ability for power spectra to break degeneracies between parameters. In the following subsection we will consider the effect of adding lensing ϕ -modes, which generally have weak dependence on cosmological parameters. Nevertheless, these modes can break parameter degeneracies, which is crucial for going beyond the 6-parameter model.

5.3 Extended parameters – neutrino mass

We will now choose a specific example to illustrate how we can think about the relationship between information and parameter constraints. Future CMB observations will target extensions to the 6-parameter Λ CDM model. In particular, the detection of the sum of the

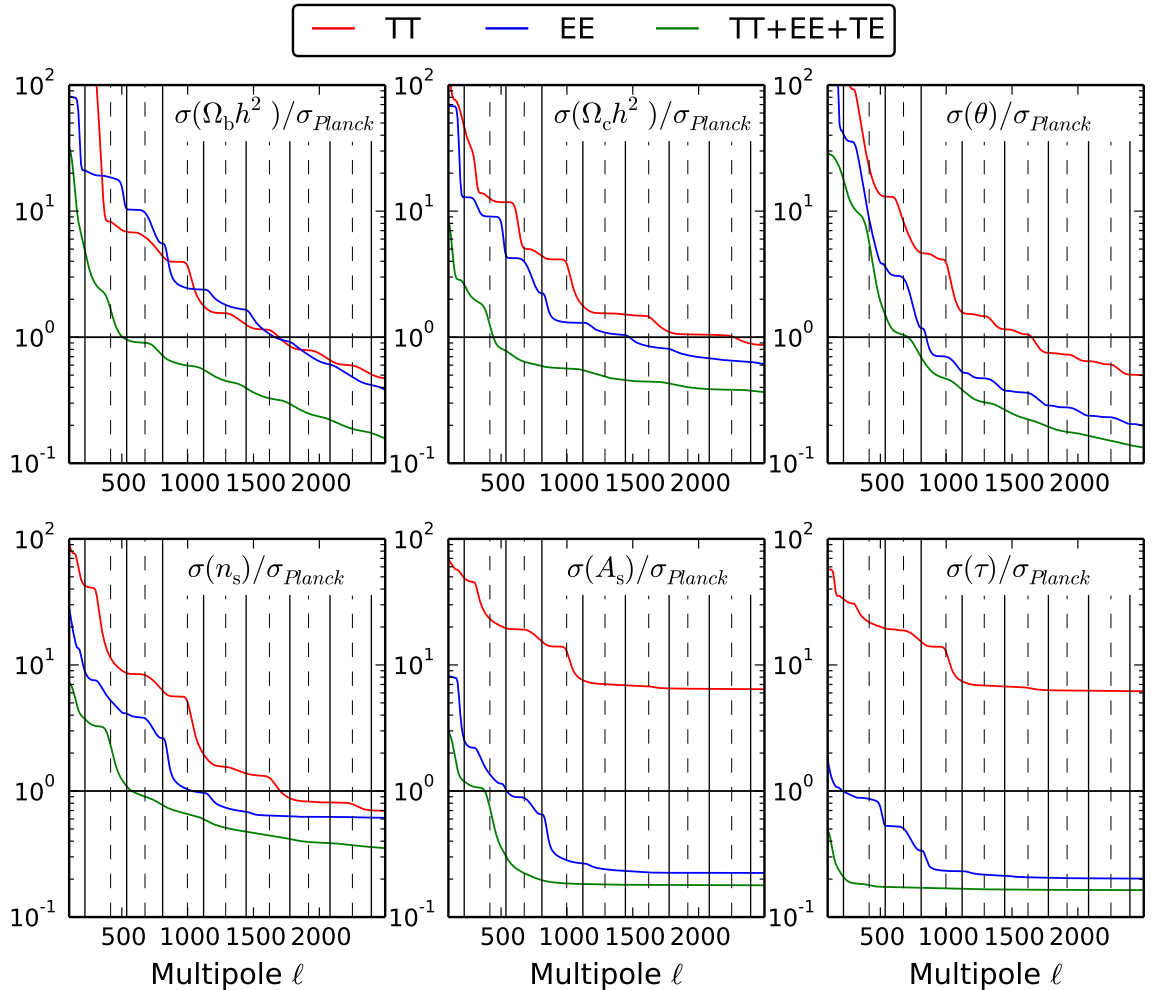


Figure 5. Cumulative Fisher forecast calculations for the six parameters of the standard Λ CDM model for a noise-free experiment with $f_{\text{sky}} = 0.5$, using *unlensed* spectra. The red line is for temperature only, the blue line for polarization only, and the green line for the full TT , TE , and EE spectra. The y -axis here is the ratio of the uncertainty on each parameter compared with the *Planck*-2015 “ TT +lowTEB+lensing” 68% confidence limits, at different maximum multipoles. The vertical lines are the positions of peaks (solid lines) and troughs (dashed lines) in the TT power spectrum (and opposite for the the EE spectrum).

masses of the neutrino species seems like a realistic (although challenging) possibility in the near future [54–56], since the current upper limit on m_ν^{tot} [16] is only a factor of a few higher than the 0.06 eV limit imposed by direct measurements of mass differences (e.g., Ref. [57]). However, the left panel of Figure 6 shows that such a measurement is not feasible with CMB temperature or polarization maps alone, because there is not enough information in the power spectra out to $\ell \simeq 3000$.

Besides the temperature and polarization fluctuations, one can also map the fluctuations of the gravitational potential or, equivalently, the lensing deflection angle. Inclusion of the lensing power spectrum in the Fisher formalism is discussed in Ref. [58], assuming Gaussianity and ignoring correlations between different multipoles. Making similar assumptions, we

find that CMB lensing can improve the constraints on the total neutrino mass dramatically, as is shown in the right panel of Figure 6. Here we have assumed that the CMB lensing power spectrum, $C_\ell^{\phi\phi}$, and its correlations with temperature and E -mode polarization can be measured with sample-variance accuracy over $f_{\text{sky}} = 0.5$ of the entire sky (see e.g., Ref. [58]). Given these assumptions, the full set of CMB power spectra reaches the fiducial sensitivity (corresponding to 60 meV in mass) at $\ell \simeq 1200$,² and can ultimately measure the mass at about the 2σ level. However, there is a very strong degeneracy between the neutrino mass and the dark matter density (with a correlation coefficient $\simeq 0.95$ at $\ell_{\text{max}} = 2500$), as well as a fairly strong degeneracy between neutrino mass and n_s (with a correlation coefficient of $\simeq 0.73$). Hence, any additional independent measurement of these parameters – e.g., via galaxy weak lensing, baryon acoustic oscillations, or redshift space distortions – that can break these degeneracies will lead to substantial improvement.




Figure 6. Fisher calculation results for the total neutrino mass using CMB data. The left panel compares the results of TT , EE , and the total ($TT+TE+EE$). The right panel compares the results for $TT+TE+EE$ with the case when CMB lensing (including cross-power spectra) is also added. It is assumed here that the CMB lensing spectra can be measured up to the sample variance limit for $f_{\text{sky}} = 0.5$. The fiducial neutrino density used here is $\Omega_\nu h^2 = 0.00064$, which corresponds to 60 meV in mass.

6 Discussion

In terms of mode-counting, the information contained in the CMB anisotropies is clearly finite, because it is limited by cosmic variance and the fact that the power spectra damp at the highest multipoles. For temperature information alone the total SNR in the power spectrum is $\sim \sqrt{f_{\text{sky}}/2} \ell_{\text{max}}$. *Planck* has measured most of what is available out to $\ell \simeq 2000$, with ACT and SPT continuing that out to higher ℓ , but over relatively small f_{sky} , and with foregrounds making it seem unrealistic to push beyond $\ell \simeq 3000$, say. This means that although we can continue to measure our CMB sky to ever more sensitive levels, we have already reached a point where the bulk of the useful information has already been extracted from the temperature anisotropies.

²Here we are using the same multipole symbol, ℓ , for lensing and temperature or polarization.

However, the situation is different for polarization information, since the foregrounds (from galaxies and clusters of galaxies) are very weakly polarized, and hence there is hope that we should be able to measure the primary polarization anisotropies to much higher multipoles, $\ell \simeq 5000$ and perhaps considerably higher [35, 59]. This means that there is at least an order of magnitude more polarization information to extract from the CMB sky. Moreover, as we have seen, the polarization data can place tighter constraints on parameters in general, and on some parameters in particular. The BB and $\phi\phi$ power spectra add additional information, but in practice this is probably a small fraction of what is available from the T - and E -modes.

Despite the dramatic improvement still expected from CMB polarization, the maximal SNR in the CMB power spectra is $\lesssim 10^4$ from primary CMB anisotropies. To improve cosmological parameter constraints we therefore need to go to 3-dimensional surveys (such as high- z 21-cm fluctuations) where there are considerably more modes [19]. As an illustration of what this CMB limitation means, let us consider the determination of the curvature of space, Ω_K . The cosmic variance limit is at the 10^{-5} level, since this is the amplitude of the curvature perturbation on the Hubble scale [60]. The current uncertainty from CMB data is at the 10^{-2} level, and given the above argument, we expect this to only decrease by about another order of magnitude. Hence, assuming that the Universe is sufficiently close to being spatially flat, then we will never be able to determine whether it is flat or curved using CMB data alone – there is simply not enough information for us to achieve the required SNR level on Ω_K from CMB anisotropies. Ambitious future experiments may probe $\gtrsim 10^{12}$ 3D modes, which would in fact allow us to reach below the cosmic-variance limit for the measurement of Ω_K within our observable volume.

The discussion here has focused on the Gaussian primary anisotropies, supplemented by CMB lensing. However, we should acknowledge that there is also some cosmological information content in the secondary anisotropies (i.e., cosmic IR background, integrated Sachs-Wolfe effect, Sunyaev-Zeldovich effects, etc.). These effects certainly enable further cosmological information to be measured, not just on the last-scattering surface, but also at other epochs along the light cone. Nevertheless the additional information seems limited in its scope for constraining background parameters, because either there is only a modest amount of information available at all (like in the ISW effect), or the additional information is still effectively on a 2D surface. The only way to obtain a dramatic improvement in constraining power will be to pursue methods that are fundamentally 3D.

7 Conclusions

We have taken a pedagogical approach to investigating the information content in CMB anisotropies, in the sense of constraining the cosmological model. It is clear that for temperature anisotropies, we have already mined a substantial part of what is available, and we are effectively running out of information. However, for CMB polarization we still have a way to go, and there may be an order of magnitude more constraining power still to extract from the CMB sky.

The CMB TT power spectrum has an SNR of approximately $\sqrt{f_{\text{sky}}/2} \ell_{\text{max}}$, which can be thought of as a simple mode-counting calculation. The SNR on a scaling parameter (like A_s) is the same, while some parameters (such as θ_*) have a dependence which is “non-linear,” hence allowing them to be constrained more tightly than the total SNR for the power spectrum as a whole.

We have shown that the mapping from information about the CMB power spectra to information about the cosmological parameters has two ingredients, namely the sensitivity of the power spectra to parameters (Eq. 3.32 and Figure 1), and the ability to break parameter degeneracies (Figures 4–6). The latter concept will likely become more important as we explore further data sets to constrain cosmology within and beyond the 6-parameter Λ CDM paradigm.

Temperature and polarization anisotropies are correlated, and hence TE contains information that enhances what is there from TT and EE alone. A full measurement of TT , EE , and TE yields one additional quantity per pixel in the map (compared with just measuring T), and hence a total SNR that is $\sqrt{2}$ times bigger than for TT alone. In addition to gaining back the information lost due to the TE correlation, when one measures all three power spectra one can also obtain additional information about how the correlation itself changes (most easily seen in Eq. A.12–A.13). In fact the information gained from the correlation can sometimes be *greater* than that from temperature (Ref. [15]) or polarization (Figure 1) alone.

Adding B -modes could in principle give one more quantity for each pixel, although in practice the primordial signal is expected to be weak. On the other hand CMB lensing provides an additional map of ϕ , which provides a whole other set of modes that can be used to constrain parameters. For the standard Λ CDM model lensing helps to break the A_s – τ degeneracy, but for extensions to the standard model (e.g., with neutrino mass included as an additional parameter) these data could be even more useful in future.

Constraints from the CMB will continue to improve as we measure more modes from polarization and from lensing. There is certainly a bright near-term future ahead as these measurements move towards being sample-variance-limited to small angular scales. In the longer term future there will be other secondary signals extractable from CMB measurements, but ultimately, to dramatically increase the number of modes probed, one will need to go to other observables (such as redshifted 21-cm maps), which can provide 3D surveys of our past light cone.

A Decorrelating T and E

As an alternative to the derivation of total CMB information content in Section 3, one can define new variables that are uncorrelated, so that the covariance matrix becomes diagonal. Doing this makes it clear that the improvement in SNR from including E -mode polarization information is exactly $\sqrt{2}$, regardless of how T and E are correlated. The change in adding B -modes is then trivial (since they are uncorrelated with either T or E). The approach we describe here is for any two correlated data sets, the specific example will be temperature and polarization data. Whether they are lensed or not does not change the arguments, although the addition of lensing data themselves would require a generalization of the method to deal with the associated correlations.

We can decorrelate temperature and polarization simply by rotating the data $\{a_{\ell m}^T, a_{\ell m}^E\}$ into a new basis, designated as $\{t_{\ell m}, e_{\ell m}\}$, through an appropriate angle θ_ℓ :

$$\begin{pmatrix} t_{\ell m} \\ e_{\ell m} \end{pmatrix} = \begin{pmatrix} \cos \theta_\ell & \sin \theta_\ell \\ -\sin \theta_\ell & \cos \theta_\ell \end{pmatrix} \begin{pmatrix} a_{\ell m}^T \\ a_{\ell m}^E \end{pmatrix}. \quad (\text{A.1})$$

We denote the power spectra derived from this new set of variables as $\lambda_\ell^{TT}, \lambda_\ell^{EE}, \lambda_\ell^{TE}$ and they are given by

$$\begin{pmatrix} \lambda_\ell^{TT} \\ \lambda_\ell^{EE} \\ \lambda_\ell^{TE} \end{pmatrix} = \begin{pmatrix} \cos^2 \theta_\ell & \sin^2 \theta_\ell & \sin 2\theta_\ell \\ \sin^2 \theta_\ell & \cos^2 \theta_\ell & -\sin 2\theta_\ell \\ -\frac{1}{2} \sin 2\theta_\ell & \frac{1}{2} \sin 2\theta_\ell & \cos 2\theta_\ell \end{pmatrix} \begin{pmatrix} C_\ell^{TT} \\ C_\ell^{EE} \\ C_\ell^{TE} \end{pmatrix}, \quad (\text{A.2})$$

or $\vec{\lambda}_\ell = R_\ell \vec{x}_\ell$, with R_ℓ defined as the transformation above. By demanding that $t_{\ell m}$ and $e_{\ell m}$ be uncorrelated (equivalently that $\lambda_\ell^{TE} = 0$) we fix the angle to be

$$\theta_\ell = \frac{1}{2} \tan^{-1} \left(\frac{2C_\ell^{TE}}{C_\ell^{TT} - C_\ell^{EE}} \right). \quad (\text{A.3})$$

Note that there are alternative approaches to decorrelating T and E , e.g., by leaving T unaltered, while removing the correlated part from E [61]; the approach we describe here is easy to picture as a rotation. The covariance matrix for these new power spectra can simply be derived by computing the 4-point functions of t and e . However, a much simpler method is to take the previous covariance matrix (i.e., \mathbb{C}) and make the following replacements: $C^{TT} \rightarrow \lambda^{TT}$; $C^{EE} \rightarrow \lambda^{EE}$; and $C^{TE} \rightarrow 0$. The covariance and inverse covariance matrices then become

$$\mathbb{L}_\ell = \frac{2}{(2\ell + 1)f_{\text{sky}}} \begin{pmatrix} \lambda_\ell^{TT^2} & 0 & 0 \\ 0 & \lambda_\ell^{EE^2} & 0 \\ 0 & 0 & \frac{1}{2} \lambda_\ell^{TT} \lambda_\ell^{EE} \end{pmatrix}, \quad (\text{A.4})$$

$$\mathbb{L}_\ell^{-1} = \left(\ell + \frac{1}{2} \right) f_{\text{sky}} \begin{pmatrix} 1/\lambda_\ell^{TT^2} & 0 & 0 \\ 0 & 1/\lambda_\ell^{EE^2} & 0 \\ 0 & 0 & 2/\lambda_\ell^{TT} \lambda_\ell^{EE} \end{pmatrix}, \quad (\text{A.5})$$

respectively (or equivalently $\mathbb{L}_\ell = R_\ell \mathbb{C}_\ell R_\ell^\top$ and $\mathbb{L}_\ell^{-1} = (R_\ell^\top)^{-1} \mathbb{C}_\ell^{-1} R_\ell^{-1}$). The data vector takes the simple form

$$\vec{\lambda}_\ell = \begin{pmatrix} \lambda_\ell^{TT} \\ \lambda_\ell^{EE} \\ \lambda_\ell^{TE} \end{pmatrix}. \quad (\text{A.6})$$

Note that even although λ_ℓ^{TE} is zero, this does not imply that for a general parameter p , $\partial \lambda_\ell^{TE} / \partial p$ will vanish.

The transformation performed here leaves the Fisher matrix unchanged. This is because

$$F'_{ij} = \sum_\ell \frac{\partial \vec{\lambda}_\ell^\top}{\partial p_i} \mathbb{L}_\ell^{-1} \frac{\partial \vec{\lambda}_\ell}{\partial p_j} \quad (\text{A.7})$$

$$= \sum_\ell \frac{\partial \vec{x}_\ell^\top}{\partial p_i} R_\ell^\top (R_\ell^\top)^{-1} \mathbb{C}_\ell^{-1} R_\ell^{-1} R_\ell \frac{\partial \vec{x}_\ell}{\partial p_j} \quad (\text{A.8})$$

$$= \sum_\ell \frac{\partial \vec{x}_\ell^\top}{\partial p_i} \mathbb{C}_\ell^{-1} \frac{\partial \vec{x}_\ell}{\partial p_j} \quad (\text{A.9})$$

$$= F_{ij}. \quad (\text{A.10})$$

Hence, for a single scaling parameter p , such that $\partial\vec{\lambda}_\ell/\partial p = \vec{\lambda}_\ell/p$, we trivially find

$$F_{pp} = \frac{2}{p^2} f_{\text{sky}} \sum_{\ell} \left(\ell + \frac{1}{2} \right), \quad (\text{A.11})$$

which is the same as the result of Eq. (3.31).

We can also consider a general parameter p like in Section 3.7, in which case we will have a fixed rotation matrix R and hence

$$\left(\frac{S}{N} \right)^2 = \sum_{\ell} \mathcal{N} \left[\left(\frac{d \ln \lambda_{\ell}^{TT}}{d \ln p} \right)^2 + \left(\frac{d \ln \lambda_{\ell}^{EE}}{d \ln p} \right)^2 + \frac{2}{\lambda_{\ell}^{TT} \lambda_{\ell}^{EE}} \left(\frac{d \lambda_{\ell}^{TE}}{d \ln p} \right)^2 \right]. \quad (\text{A.12})$$

The interpretation of the terms on the right-hand side is quite clear: the first two are the information coming from temperature and polarization anisotropies, respectively, after removing their correlation; and the final term is the information coming from the correlation itself. This can be seen by re-writing this in terms of C^{TE} and θ_{ℓ} as

$$\begin{aligned} \left(\frac{S}{N} \right)^2 = \sum_{\ell} \mathcal{N} \left[\left(\frac{d \ln \lambda_{\ell}^{TT}}{d \ln p} \right)^2 + \left(\frac{d \ln \lambda_{\ell}^{EE}}{d \ln p} \right)^2 \right. \\ \left. + \frac{(C_{\ell}^{TT} - C_{\ell}^{EE})^2}{2\lambda_{\ell}^{TT} \lambda_{\ell}^{EE}} \left(\cos 2\theta_{\ell} \frac{d \tan 2\theta_{\ell}}{d \ln p} \right)^2 \right]. \end{aligned} \quad (\text{A.13})$$

So we see, as we found in Section 3.7, that depending on the parameter p , there is extra information to be found from the correlation itself, even if $C_{\ell}^{TE} = 0$.

B The fully correlated case

Note that the Fisher formalism we have employed in this paper relies on the covariance matrix being invertible. In a realistic experiment this is not an issue, because, regardless of the true correlation between modes, noise will regularize the problem, such that the covariance will always be invertible. Additionally, on physical grounds, it is clear that we need not worry even in the case of no noise, because we will *never* be in the situation where $r_{\ell}^2 = 1$. We can see the problem with $r_{\ell}^2 \rightarrow 1$ in the following way – adding E should double the number of independent modes, irrespective of how correlated T and E are, *except* for the case $r_{\ell}^2 = 1$, when there is no new information from E (and hence the mode count appears to be discontinuous as $r_{\ell}^2 \rightarrow 1$). Clearly this situation is hypothetical and hence should not unduly concern us.

However, there is a similar case where we *would* have a singular covariance and that is the case of E -modes, lensing, and lensed B -modes. It is apparent that the lensed B -modes are completely determined by E -modes and lensing modes, ϕ . Therefore we can simply ask what there is to be gained from actually measuring the B -modes. The answer is that (in principle) there is nothing gained in this measurement! This is true if we assume that a perfect measurement of E , B , and ϕ can be performed. From an experimental point of view this will never really be true, due to the presence of noise (i.e., one can always hope to beat down the noise in lensed B -modes by actually measuring them). Of course the Fisher-formalism assumes the model to be correct, so from a theoretical point of view (and looking beyond standard Λ CDM), measuring lensed B -modes will of course always have merit.

Acknowledgments

This research was supported by the Natural Sciences and Engineering Research Council of Canada and by the Canadian Space Agency. We thank Jim Zibin for several enjoyable discussions on this topic and for providing useful comments on the paper.

References

- [1] Planck Collaboration, ArXiv e-prints (2015), [arXiv:1502.01582](#) .
- [2] C. L. Bennett *et al.*, *ApJS* **208**, 20 (2013), [arXiv:1212.5225](#) .
- [3] S. Das *et al.*, *JCAP* **4**, 014 (2014), [arXiv:1301.1037](#) .
- [4] K. T. Story *et al.*, *ApJ* **779**, 86 (2013), [arXiv:1210.7231](#) .
- [5] D. Scott and G. F. Smoot, in *The Review of Particle Physics*, Chin. Phys. C, Vol. 38, edited by K. Olive (2014) Chap. 27, p. 09001.
- [6] G. Jungman, M. Kamionkowski, A. Kosowsky, and D. N. Spergel, *PRD* **54**, 1332 (1996), [astro-ph/9512139](#) .
- [7] J. R. Bond, G. Efstathiou, and M. Tegmark, *MNRAS* **291**, L33 (1997), [astro-ph/9702100](#) .
- [8] M. Zaldarriaga, D. N. Spergel, and U. Seljak, *ApJ* **488**, 1 (1997), [astro-ph/9702157](#) .
- [9] G. Efstathiou and J. R. Bond, *MNRAS* **304**, 75 (1999), [astro-ph/9807103](#) .
- [10] G. Rocha, R. Trotta, C. J. A. P. Martins, A. Melchiorri, P. P. Avelino, R. Bean, and P. T. P. Viana, *MNRAS* **352**, 20 (2004), [astro-ph/0309211](#) .
- [11] L. Perotto, J. Lesgourgues, S. Hannestad, H. Tu, and Y. Y Y Wong, *JCAP* **10**, 013 (2006), [astro-ph/0606227](#) .
- [12] C. Burigana, C. Destri, H. J. de Vega, A. Gruppuso, N. Mandolesi, P. Natoli, and N. G. Sanchez, *ApJ* **724**, 588 (2010), [arXiv:1003.6108](#) .
- [13] S. Foreman, A. Moss, and D. Scott, *PRD* **84**, 043522 (2011), [arXiv:1106.4018](#) .
- [14] W. L. K. Wu, J. Errard, C. Dvorkin, C. L. Kuo, A. T. Lee, P. McDonald, A. Slosar, and O. Zahn, *ApJ* **788**, 138 (2014), [arXiv:1402.4108](#) .
- [15] S. Galli, K. Benabed, F. Bouchet, J.-F. Cardoso, F. Elsner, E. Hivon, A. Mangilli, S. Prunet, and B. Wandelt, *PRD* **90**, 063504 (2014), [arXiv:1403.5271](#) .
- [16] Planck Collaboration, ArXiv e-prints (2015), [arXiv:1502.01589](#) .
- [17] A. Lewis, A. Challinor, and A. Lasenby, *ApJ* **538**, 473 (2000), [astro-ph/9911177](#) .
- [18] A. Loeb, *JCAP* **5**, 028 (2012), [arXiv:1203.2622](#) .
- [19] Y.-Z. Ma and D. Scott, ArXiv e-prints (2015), [arXiv:1510.08309](#) .
- [20] U. Seljak, *ApJ* **503**, 492 (1998), [astro-ph/9710269](#) .
- [21] N. Padmanabhan, M. Tegmark, and A. J. S. Hamilton, *ApJ* **550**, 52 (2001), [astro-ph/9911421](#) .
- [22] C. D. Rimes and A. J. S. Hamilton, *MNRAS* **360**, L82 (2005), [astro-ph/0502081](#) .
- [23] P. McDonald and U. Seljak, *JCAP* **10**, 007 (2009), [arXiv:0810.0323](#) .
- [24] J. Carron, M. Wolk, and I. Szapudi, *MNRAS* **453**, 450 (2015), [arXiv:1412.5511](#) .
- [25] L. F. Abbott and M. B. Wise, *ApJL* **282**, L47 (1984).
- [26] D. Scott, M. Srednicki, and M. White, *ApJL* **421**, L5 (1994), [astro-ph/9305030](#) .
- [27] L. Knox, *PRD* **52**, 4307 (1995), [astro-ph/9504054](#) .

- [28] S. Grandis, S. Seehars, A. Refregier, A. Amara, and A. Nicola, ArXiv e-prints (2015), [arXiv:1510.06422](#) .
- [29] M. Zaldarriaga and U. Seljak, *PRD* **55**, 1830 (1997), [astro-ph/9609170](#) .
- [30] M. Kamionkowski, A. Kosowsky, and A. Stebbins, *PRD* **55**, 7368 (1997), [astro-ph/9611125](#) .
- [31] W. Hu and M. White, *PRD* **56**, 596 (1997), [astro-ph/9702170](#) .
- [32] G. Hinshaw *et al.*, *ApJS* **180**, 225 (2009), [arXiv:0803.0732](#) .
- [33] K. T. Story *et al.*, *ApJ* **779**, 86 (2013), [arXiv:1210.7231](#) .
- [34] S. Das *et al.*, *JCAP* **4**, 014 (2014), [arXiv:1301.1037](#) .
- [35] S. Naess *et al.*, *JCAP* **10**, 007 (2014), [arXiv:1405.5524](#) .
- [36] Planck Collaboration, ArXiv e-prints (2015), [arXiv:1507.02704](#) .
- [37] A. T. Crites *et al.*, *ApJ* **805**, 36 (2015), [arXiv:1411.1042](#) .
- [38] BICEP2 and Keck Array Collaborations, *ApJ* **811**, 126 (2015), [arXiv:1502.00643](#) .
- [39] BICEP2/Keck and Planck Collaborations, *Physical Review Letters* **114**, 101301 (2015), [arXiv:1502.00612](#) .
- [40] D. Hanson *et al.*, *Physical Review Letters* **111**, 141301 (2013), [arXiv:1307.5830](#) .
- [41] The Polarbear Collaboration, *ApJ* **794**, 171 (2014), [arXiv:1403.2369](#) .
- [42] R. Keisler *et al.*, *ApJ* **807**, 151 (2015), [arXiv:1503.02315](#) .
- [43] Planck Collaboration, ArXiv e-prints (2015), [arXiv:1502.01591](#) .
- [44] A. van Engelen *et al.*, *ApJ* **808**, 7 (2015), [arXiv:1412.0626](#) .
- [45] M. Tegmark, A. N. Taylor, and A. F. Heavens, *ApJ* **480**, 22 (1997), [astro-ph/9603021](#) .
- [46] D. J. Eisenstein, W. Hu, and M. Tegmark, *ApJ* **518**, 2 (1999), [astro-ph/9807130](#) .
- [47] M. Zaldarriaga, *PRD* **55**, 1822 (1997), [astro-ph/9608050](#) .
- [48] U. Seljak, *ApJ* **482**, 6 (1997), [astro-ph/9608131](#) .
- [49] G. Simard, D. Hanson, and G. Holder, *ApJ* **807**, 166 (2015), [arXiv:1410.0691](#) .
- [50] Planck Collaboration, ArXiv e-prints (2015), [arXiv:1502.01592](#) .
- [51] K. M. Smith, W. Hu, and M. Kaplinghat, *PRD* **74**, 123002 (2006), [astro-ph/0607315](#) .
- [52] A. Challinor, A. Lewis, G. Rocha, and S. Smith, *New Astron. Rev.* **51**, 421 (2007).
- [53] A. Lewis and A. Challinor, *PhysRep* **429**, 1 (2006), [astro-ph/0601594](#) .
- [54] J. Lesgourgues, L. Perotto, S. Pastor, and M. Piat, *PRD* **73**, 045021 (2006), [astro-ph/0511735](#) .
- [55] Z. Pan and L. Knox, *MNRAS* **454**, 3200 (2015), [arXiv:1506.07493](#) .
- [56] R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis, *PRD* **92**, 123535 (2015), [arXiv:1509.07471](#) .
- [57] G. J. Feldman, J. Hartnell, and T. Kobayashi, *Adv. High Energy Phys.* **2013**, 475749 (2013), [arXiv:1210.1778 \[hep-ex\]](#) .
- [58] L. Santos, P. Cabella, A. Balbi, and N. Vittorio, *PRD* **88**, 043505 (2013), [arXiv:1307.2919](#) .
- [59] B. A. Benson *et al.*, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, ProcSpie, Vol. 9153 (2014) p. 91531P, [arXiv:1407.2973 \[astro-ph.IM\]](#) .
- [60] T. P. Waterhouse and J. P. Zibin, ArXiv e-prints (2008), [arXiv:0804.1771](#) .

[61] M. Frommert and T. A. Enßlin, [MNRAS 395, 1837 \(2009\)](#).