CrossMark

New Constraint on the Tensor-to-scalar Ratio from the Planck and BICEP/Keck Array Data Using the Profile Likelihood

Paolo Campeti[1,2] and Eiichiro Komatsu[1,3]
[1] Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str.1, D-85741 Garching, Germany
[2] Excellence Cluster ORIGINS, Boltzmannstr. 2, D-85748 Garching, Germany
[3] Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU, WPI), UTIAS, The University of Tokyo, Chiba, 277-8583, Japan

Abstract

Motivated by the discrepancy between Bayesian and frequentist upper limits on the tensor-to-scalar ratio parameter r found by the SPIDER collaboration, we investigate whether a similar trend is also present in the latest Planck and BICEP/Keck Array data. We derive a new upper bound on r using the frequentist profile likelihood method. We vary all the relevant cosmological parameters of the ΛCDM model, as well as the nuisance parameters. Unlike the Bayesian analysis using Markov Chain Monte Carlo (MCMC), our analysis is independent of the choice of priors. Using Planck Public Release 4, BICEP/Keck Array 2018, Planck cosmic microwave background lensing, and baryon acoustic oscillation data, we find an upper limit of $r < 0.037$ at 95% Confidence Level (C.L.), similar to the Bayesian MCMC result of $r < 0.038$ for a flat prior on r and a conditioned Planck lowlEB covariance matrix.

1. Introduction

Detecting the stochastic background of primordial gravitational waves predicted within the inflationary paradigm (Grishchuk 1974; Starobinsky 1979) represents one of the principal objectives of the current cosmological research, as it would provide the definitive evidence for cosmic inflation (Guth 1981; Sato 1981; Albrecht & Steinhardt 1982; Linde 1982).

Whereas inflation produces gravitational waves (i.e., tensor modes) over a wide range in frequency measurable by several different probes (see, e.g., Campeti et al. 2021, for a review), the most promising route to detection is the B-mode polarization of the cosmic microwave background (CMB; Kamionkowski et al. 1997; Seljak & Zaldarriaga 1997).

The current data sets only provide upper bounds on the tensor-to-scalar ratio r (i.e., the ratio of the amplitudes of the tensor and scalar modes power spectra). To date, the tightest limit on r (customarily measured at the pivot scale $k_0 = 0.05$ Mpc^{-1}) is $r < 0.032$ at 95% Confidence Level (C.L.; Tristram et al. 2022), coming from the Planck latest CMB temperature and E and B-mode polarization data (Tristram et al. 2021), the BICEP/Keck Array B-mode data (BICEP/Keck Collaboration 2021, hereafter BK18), the baryon acoustic oscillations (BAO) of the large-scale structure (Alam et al. 2021), and the CMB lensing data (Planck Collaboration VIII 2018). This upper limit is derived using a standard Bayesian Monte Carlo Markov Chain (MCMC) procedure, varying the relevant cosmological parameters of a flat Λ cold dark matter (ΛCDM) model and adopting the Sellentin & Heavens (2016, hereafter SH) correction to the Hamimeche & Lewis (2008, hereafter HL) likelihood for the Planck large-scale EE, BB, and EB power spectra (the "lowlEB" likelihood). The SH correction is needed to account for the increased uncertainty in parameter estimation due to the limited number of simulations used to estimate the covariance matrix. This

is obtained by analytically marginalizing over the unknown true covariance matrix.

Most of the constraining power on r at the pivot scale comes from BK18's B-mode data. An upper limit of $r < 0.036$ at 95% C.L. (BICEP/Keck Collaboration 2021) is obtained just from the BK18 data, provided that we fix the ΛCDM parameters to their best-fitting values given in Planck Collaboration VI (2020). The Planck satellite provides, on the other hand, the tightest constraints to date on the B modes at the largest angular scales, which are not accessible from the ground. Exploiting the latest NPIPE-processed Public Release 4 (PR4) of temperature and polarization maps (Planck Collaboration Int. LVII 2020), the Planck collaboration reported a limit of $r < 0.056$ at 95% C. L. (Tristram et al. 2021), which is relaxed to $r < 0.075$ when properly accounting for the SH correction in the lowlEB likelihood (Beck et al. 2022).

While the SH correction accounts for the Monte Carlo noise in the estimated covariance matrix, it does not correct for the additional scatter in the best-fitting maximum a posteriori parameter (MAP) estimate, which can lead to a misestimation of confidence limits (Beck et al. 2022). This effect is especially relevant near the physical boundary of a given parameter (i.e., $r \geqslant 0$ in our case of interest) and can produce a significant underestimation of the upper limit. The issue can be corrected by increasing the number of (computationally expensive) time-ordered data simulations used in the covariance matrix estimation or by properly conditioning the covariance matrix. The latter method has been applied in Beck et al. (2022) to the lowlEB Planck likelihood (which we will refer to as "conditioned HL" in the following), resulting in a much weaker Planck-only upper limit of $r < 0.13$ at 95% C.L., associated to a large shift of the peak of the marginalized distribution to larger r values than in the SH case. Similarly, for the Planck + BK18 + BAO + lensing combination, the conditioning results in a more conservative upper limit of $r < 0.038$.

In this paper, we present constraints on r using the frequentist profile likelihood method, and compare them to the standard Bayesian MCMC procedure adopted throughout

the literature.[4] As shown by results from the SPIDER collaboration (Ade et al. 2022), the two approaches can give quite different answers, with the Bayesian upper limit being almost a factor of 2 larger than its frequentist counterpart. While the profile likelihood is a standard data analysis tool in particle physics (see, e.g., ATLAS Collaboration 2013; Particle Data Group 2020), it has been seldom used in cosmology, notable cases of use being the application to ΛCDM parameters estimation from the Planck data (Planck Collaboration Int. XVI 2014), to the Early Dark Energy fraction (Herold et al. 2022), to coupled dark energy and Brans-Dicke models (Gómez-Valent 2022) and, as anticipated above, to the estimation of r from the SPIDER data (Ade et al. 2022). Nonetheless, this approach bears several potentially interesting differences with Bayesian methods (Cousins 1995). First, the profile likelihood does not require priors, which may have an impact on the final constraints. Second, while in Bayesian methods the choice of a specific set of parameters to sample might represent an implicit prior choice, the maximum likelihood estimate (MLE) is invariant under model reparameterization. Third, the parameter estimates obtained from the profile likelihood are not affected by "volume effects" that can arise during marginalization in the MCMC approach (Hamann et al. 2007). Moreover, the profile likelihood formalism allows us to conveniently include the effect of the parameter's physical boundary in the confidence intervals via the Feldman–Cousins prescription (Feldman & Cousins 1998).

Our work aims to deconstruct the current constraints on r and scrutinize their robustness. Similarly to the profile likelihood analysis performed on the ΛCDM parameters (Planck Collaboration Int. XVI 2014), we study the effect of priors and marginalization on the inference of r from the Planck and BK18 data. We also explore the effect of conditioning the Planck lowlEB covariance matrix (Beck et al. 2022) on the profile likelihood.

The structure of the paper is the following. We describe the data and likelihood used in our analysis in Section 2. We review the profile likelihood formalism and the Feldman–Cousins prescription in Section 3. We discuss the new constraints on r from our frequentist analysis and compare them to the Bayesian credible intervals in Section 4. We conclude in Section 5.

2. Data and Likelihoods

We use the latest Planck NPIPE-processed PR4 maps (Planck Collaboration Int. LVII 2020) and the BK18 data set (BICEP/Keck Collaboration 2021). We use the data and likelihoods publicly available for the Cobaya[5] (Torrado & Lewis 2021) MCMC framework, as done in Tristram et al. (2022). We also use Cobaya as an interface with the CAMB Boltzmann solver (Lewis et al. 2000).

2.1. Planck Likelihoods

The Planck likelihood consists of three parts: the low-ℓ TT Commander likelihood (Planck Collaboration V 2020) for $\ell = 2$–30, the high-ℓ TT + TE + EE HiLLiPoP likelihood[6] (Planck Collaboration XV 2013; Planck Collaboration XI 2016;

Couchot et al. 2017) for $\ell = 30$–2500, and the low-ℓ EE + BB + EB LoLLiPoP or the lowlEB likelihood[7] (Tristram et al. 2021) for $\ell = 2$–150.

The low-ℓ TT likelihood is the same as in PR3, since no improvement is expected with the PR4 update for the high signal-to-noise temperature data. The HiLLiPoP likelihood is instead a Gaussian likelihood for cross-power spectra of the Planck 100, 143, and 217 GHz data.

The LoLLiPoP likelihood for large-scale EE, BB, and EB power spectra implements the HL approximation for a non-Gaussian likelihood (Hamimeche & Lewis 2008), adapted specifically for cross-power spectra (Mangilli et al. 2015). In this case, an offset term is needed to make the distribution of cross-power spectra similar to that of autopower spectra, as required by the HL approximation. The covariance matrix for this likelihood is estimated from 400 Monte Carlo simulations of PR4, which include Planck noise, systematic effects, and foreground residuals. The Planck lowlEB likelihood implements the SH correction to the HL likelihood to account for the Monte Carlo noise in the covariance matrix estimate. This is not sufficient to amend the additional scatter in the MAP estimate: a possible solution indicated in Beck et al. (2022) involves using the HL likelihood (without the SH correction) with a conditioned covariance matrix. The conditioning strategy removes all off-diagonal elements beyond the next-to-nearest neighbor for unbinned multipoles ($\ell \leqslant 35$) and all off-diagonal elements beyond the nearest neighbor for binned multipoles ($\ell > 35$). We will refer to this specific choice as "cond. HL" in the following.

2.2. BICEP/Keck Array 2018 Likelihood

The BK18 likelihood, which includes only B modes at $\ell \simeq 30$–300, also applies the HL approximation to auto- and cross-power spectra in conjunction with the WMAP data at 23 and 33 GHz and Planck NPIPE-processed data at 30, 44, 143, 217, and 353 GHZ. The bandpower covariance matrix is estimated from 499 simulations. The default BK18 likelihood already incorporates conditioning to reduce the Monte Carlo noise.

2.3. Likelihood Combination and Priors in the Default Analysis

We combine the Planck and BK18 likelihoods, neglecting correlations between them. This is a good approximation because the current B-mode data are noise-dominated, the two CMB surveys have uncorrelated noises, and they observe very different fractions of the sky (i.e., 50% for Planck and 1% for BK18; see Tristram et al. 2021, 2022). In the following, whenever we use the Planck likelihood, we will also include the BAO data (Alam et al. 2021) and the Planck CMB lensing data (Planck Collaboration VIII 2018).

There are in total 33 free parameters in the default Planck + BK18 analysis, including r, 6 parameters of a flat ΛCDM model $\{\Omega_b h^2, \Omega_c h^2, \tau, A_s, n_s, \theta_{MC}\}$, and the nuisance parameters. The tensor spectral index n_t is fixed via the inflationary consistency relation $n_t = -r/8$, similarly to previous analyses (Tristram et al. 2021, 2022). We also checked that fixing $n_t = 0$ as in the BICEP/Keck Collaboration (2021) analysis does not impact our results.

The Planck likelihoods introduce 19 nuisance parameters, accounting for map and absolute calibration and foreground

---

modeling (for a description see Appendix B in Tristram et al. 2021). Of these, eight parameters have a Gaussian prior in the default MCMC analysis, whereas the others have uniform priors. The BK18 likelihood has seven nuisance parameters accounting for Galactic dust and synchrotron foreground modeling. Of these, six parameters have uniform priors in the default BK18 analysis (BICEP/Keck Collaboration 2021), whereas the synchrotron spectral index β_s has a Gaussian prior $\beta_s = -3.1 \pm 0.3$ (motivated by the WMAP 23 and 33 GHz data; Fuskeland et al. 2014). As shown in BICEP/Keck Collaboration (2021), the constraint on β_s from the BK18 data is prior-dominated; therefore, for a more direct comparison with the Bayesian results in the literature, we also explore the possibility of fixing $\beta_s = -3.1$ in the profile likelihood, since frequentist analyses do not incorporate priors. We indicate such choice as "fixed β_s" in the following.

3. Profile Likelihood

We use the profile likelihood to investigate the effects of priors and marginalization on the current Bayesian constraints on r. The profile likelihood is a staple in the frequentist's toolbox. As it does not incorporate priors, explicitly or implicitly via the model parameterization, it is immune to volume effects which may appear during marginalization in MCMC.

The profile likelihood for a parameter of interest μ (in our case $\mu = r$) is obtained by fixing μ to multiple values within the range of interest and minimizing the $\chi^2(\mu) = -2 \log \mathcal{L}(\mu)$ with respect to all the remaining cosmological and nuisance parameters for each fixed value of μ. Here, $\mathcal{L}$ is the likelihood. By construction the minimum χ^2_{min} coincides with the global MLE (also called "best-fit").

We use $\Delta\chi^2(\mu) = \chi^2(\mu) - \chi^2_{min}$ to construct frequentist confidence intervals on μ. If μ is far away from its physical boundary, a confidence interval at α C.L. can be obtained by cutting $\Delta\chi^2(\mu)$ at a fixed threshold $\Delta\chi^2_{th}$ such that the cumulative distribution function of the χ^2 distribution with one degree of freedom is equal to α (e.g., cutting at $\Delta\chi^2_{th} = 1$ and $\Delta\chi^2_{th} = 3.84$ for 68% or 95% C.L., respectively; see, e.g., Trotta 2017). We can use this procedure for both parabolic (associated to a Gaussian-distributed parameter) and nonparabolic $\Delta\chi^2(\mu)$ thanks to invariance of the MLE under reparameterization.

3.1. The Feldman–Cousins Prescription

If the parameter estimate is instead close to its physical boundary, as in our case of interest, the classical Neyman's construction of frequentist confidence intervals is unsatisfactory. It can lead to empty intervals and to failure of the frequentist coverage property[8] if the choice of reporting an upper limit or a two-sided interval is made by looking at the data.

These issues can be solved by adopting the Feldman & Cousins (1998, hereafter FC) prescription. For each value μ of the parameter of interest (with unknown true value) and each observable x, we compute the *likelihood ratio*

$$R(x, \mu) = \frac{\mathcal{L}(x|\mu)}{\mathcal{L}(x|\mu_{best})}, \qquad (1)$$

where μ can take only physically allowed values and μ_{best} is the value of μ which maximizes the likelihood $\mathcal{L}(x|\mu)$. The so-called *confidence belt* at the desired α C.L. is then built by selecting for each μ an *acceptance interval* $[x_1, x_2]$ such that

$$\begin{cases} R(x_1, \mu) = R(x_2, \mu), \\ \int_{x_1}^{x_2} P(x|\mu)dx = \alpha, \end{cases} \qquad (2)$$

where $P(x|\mu)$ is the probability density function for x given μ and the values x are added to the acceptance interval in order of decreasing likelihood ratio. The confidence belt is then given by the union of all acceptance intervals $[x_1(\mu), x_2(\mu)]$: intercepting it with a line at $x = x_0$, with x_0 being the value of x minimizing χ^2 (i.e., the value measured in the experiment), we obtain the confidence interval $[\mu_1, \mu_2]$ for the parameter μ.

The FC prescription provides, therefore, an additional criterion to fix the extrema of the confidence intervals in Neyman's construction and to transition between an upper limit and a two-sided interval, giving exact frequentist coverage for a Gaussian parameter even in proximity of a physical boundary. This is in contrast with the conservatism (i.e., overcoverage) inherent to Bayesian limits in the same context (Cousins 1995; Feldman & Cousins 1998). While conservatism might not be as a severe issue as undercoverage, it certainly degrades our ability to discriminate against false hypotheses, making it worthwhile to examine frequentist intervals.

As we will see in the next section, the profile likelihood for r gives a parabolic $\Delta\chi^2$ near its minimum, and presents a physical boundary at $r = 0$. In this case, $\mu_{best} = \max(0, x)$ and the likelihood ratio in Equation (1) becomes (Feldman & Cousins 1998)

$$R(x, \mu) = \begin{cases} \exp(-(x - \mu)^2/2), & \text{for } x \geqslant 0, \\ \exp(x\mu - \mu^2/2), & \text{for } x < 0, \end{cases} \qquad (3)$$

where x and μ are expressed in units of σ, that is, the width of the parabolic fit to $\Delta\chi^2(\mu)$. The confidence interval is obtained by solving the system given in Equation (2) with $P(x|\mu)$ being a Gaussian with mean μ and unit variance.

3.2. Minimization Algorithm

We minimize $\chi^2(r)$ with the MIGRAD algorithm implemented in the iMinuit[9] package, a python interface for the popular Minuit multidimensional minimizer.[10] We scan the parameter space, fixing r to values over a wide range and minimizing $\chi^2(r)$ with respect to the remaining 32 free parameters for each fixed r. Each point in the profile likelihood typically requires around $\mathcal{O}(10^4)$ evaluations of the likelihood, with each evaluation taking $\mathcal{O}(1)$ s (using 10 logical central processing units on a computer cluster node), almost exclusively absorbed by the evaluation of the CAMB Boltzmann code.[11] To increase the chance of the minimizer reaching the

---

[8] The frequentist coverage property is realized at the level α if a fraction α of the confidence intervals obtained from Neyman's construction contains the fixed and unknown true value of the parameter of interest. (see, e.g., Cousins 2018).

[9] iminuit.readthedocs.io

[10] We found significantly better performances using iMinuit compared to other common minimizing algorithms (e.g., the scipy minimize module (Virtanen et al. 2020) and Py-BOBYQA (numericalalgorithmsgroup.github.io), which are already implemented in the Cobaya sampler).

[11] The profile likelihood is highly competitive with the more traditional MCMC approach, which requires $\mathcal{O}(10^6)$ points to reach convergence, due to inefficient sampling of the Metropolis–Hastings algorithm near the boundary of a parameter with a uniform positive prior.
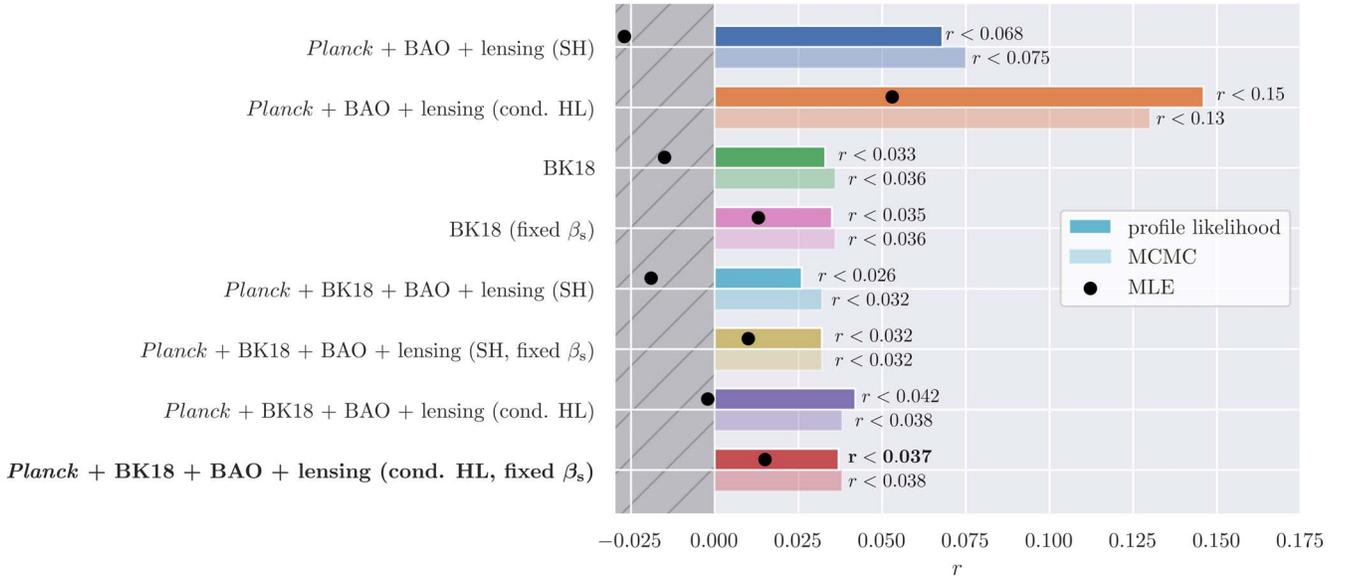
Figure 1. Summary of 95% C.L. upper limits on r for data sets considered in this work. The darker shaded bars indicate the upper limit from the profile likelihood, whereas the lighter shaded bars the MCMC one. The best-fitting r values from the profile likelihood analysis are shown as the black dots. Negative (unphysical) values of r are indicated by the hatched dark gray region. The baseline result of this work is highlighted in bold.

Table 1
Upper Limits on the Tensor-to-scalar Ratio Parameter r (95% C.L.) from the Profile Likelihood Method with the FC Prescription and the MCMC

Data	Likelihood	Fixed β_s	Profile (95% C.L.)	MCMC (95% C.L.)	r_{MLE}	Color
Planck + BAO + lensing	SH	⋯	$r < 0.068$	$r < 0.075$	-0.027	■
	cond. HL	⋯	$r < 0.15$	$r < 0.13$	0.053	■
BK18	fix ΛCDM params.	✗	$r < 0.033$	$r < 0.036$	-0.015	■
		✓	$r < 0.035$	$r < 0.036$	0.013	■
Planck + BK18 + BAO + lensing	SH	✗	$r < 0.026$	$r < 0.032$	-0.019	■
		✓	$r < 0.032$	$r < 0.032$	0.01	■
	cond. HL	✗	$r < 0.042$	$r < 0.038$	-0.0021	■
		✓	$\mathbf{r < 0.037}$	$r < 0.038$	0.015	■

Note. We also report the MLE for r obtained from the profile likelihood method. For "cond. HL" we adopt the conditioning prescription defined in Beck et al. (2022) for the Planck lowlEB likelihood. For "SH" we marginalize the likelihood over the covariance matrix (Sellentin & Heavens 2016). Note that for BK18-only data we fix all six ΛCDM parameters to the best-fitting values given in Planck Collaboration VI (2020). For each case involving the BK18 likelihood, we indicate whether we are fixing the synchrotron spectral index to $\beta_s = -3.1$ (see Section 2 for details). The baseline result of this work is highlighted in bold. The colors shown in the rightmost column match those in Figures 1 and 2.

global minimum, each minimization is started from 10 different random initial parameter sets. We take the point with the lowest χ^2 as the final result. We checked that increasing the accuracy settings of the CAMB code does not change our results.

## 4. Results and Comparison to the Bayesian Analysis

In Figure 1 and Table 1 we report 95% C.L. upper limits on r obtained from the profile likelihood (darker shaded bars) and compare them to their Bayesian MCMC counterparts (lighter shaded bars). We also show the best-fitting r, that is, the global MLE values found in the profile likelihood analysis as the black dots, and indicate negative (unphysical) values of r by the hatched dark gray region.

We consider three data set combinations: Planck + BAO + lensing, BK18-only, and Planck + BK18 + BAO + lensing. For each combination involving the Planck data, we show the results obtained marginalizing over the lowlEB covariance matrix ("SH") and the ones conditioning it ("cond. HL"), as discussed in Section 2. For each combination involving

the BK18 data, we show the results fitting for the synchrotron spectral index β_s and the ones fixing it to $\beta_s = -3.1$ (i.e., to the mean of the Gaussian prior imposed in the MCMC default analysis; see Section 2.3), labeled as "fixed β_s."

In Figure 2 we show $\Delta\chi^2(r)$ and the respective parabolic fits (solid lines), the upper limits from the FC prescription (vertical dashed lines), and the values of $\Delta\chi^2$ corresponding to each upper limit (horizontal dashed lines). The colors indicated in Table 1 match those in Figures 1 and 2.

We start by discussing the results of Planck + BAO + lensing (blue and orange bars in Figure 1). If the SH correction is used, the MLE from the profile likelihood lies in the unphysical region of the parameter space ($r_{\mathrm{MLE}} = -0.027$), whereas conditioning the Planck lowlEB covariance matrix shifts it to a large positive value ($r_{\mathrm{MLE}} = 0.053$). This results in a significantly larger upper limit ($r < 0.15$ instead of $r < 0.068$) in the latter case, despite the width of $\Delta\chi^2(r)$ being the same in both cases. This confirms the findings of Beck et al. (2022) in a prior-independent manner.
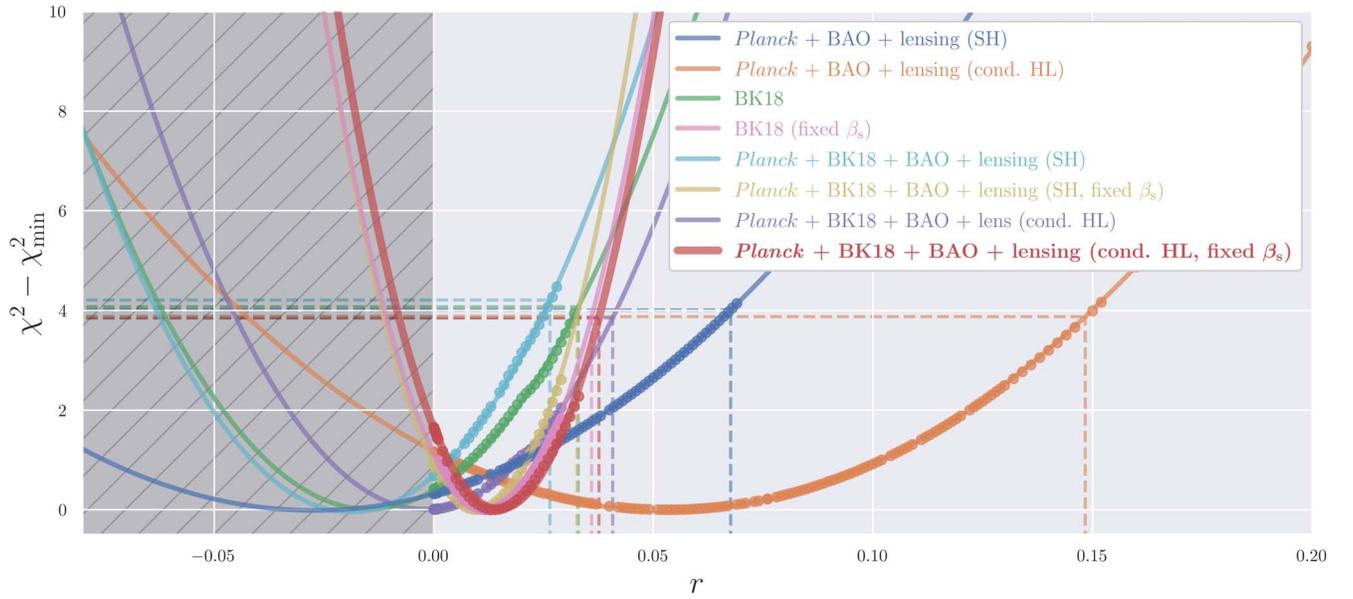
Figure 2. Profile likelihoods for r from the data sets combinations considered in this work. The points are the $\chi^2 - \chi^2_{\min}$ values obtained from the likelihood maximization, whereas the parabolic fits are shown as the solid lines. The dashed lines indicate the upper limits at 95% C.L. according to the FC prescription. Unphysical (negative) values of r are shown as the dark gray hatched area. The baseline result is shown in the thick red line.

Comparing the MCMC and profile likelihood limits, we observe that the limit from the profile likelihood is tighter in the SH-corrected case, whereas the opposite is true in the conditioned HL case. In the SH case, the MLE lies deep in the negative region, where Bayesian intervals notoriously over-cover in the presence of a boundary (Feldman & Cousins 1998). In the conditioned HL case, instead, the weaker profile likelihood limit is partly due to the different coverage properties and definitions of Bayesian limits compared to FC far from the boundary and partly to the effect of the Gaussian priors in the Planck likelihood.

The BK18-only constraints (green bars in Figure 1) are obtained fixing the ΛCDM parameters to their Planck Collaboration VI (2020) best-fitting values, as done in BICEP/Keck Collaboration (2021). The upper limit from the profile likelihood is slightly tighter than the corresponding MCMC case ($r < 0.033$ versus $r < 0.036$), with an MLE lying deep in the negative region. We notice, however, that the best-fitting model prefers a value $\beta_s \simeq -2$ for the synchrotron spectral index nuisance parameter in the BK18 likelihood (see the Appendix). On the other hand, constraints on this parameter in the default MCMC analysis are prior-driven (see Section 2.3) and prefer a value $\beta_s \simeq -3$. Therefore, a more straightforward comparison with the frequentist approach can be drawn after fixing β_s to the central value of the Gaussian prior in the profile likelihood (pink bars in Figure 1). We then recover $r < 0.035$, very close to the Bayesian result given in BICEP/Keck Collaboration (2021).

The difference between the FC and Bayesian limits obtained from the Planck+BK18+BAO+lensing data, both in the SH-corrected and the conditioned HL cases (the light blue and purple bars in Figure 1), is also due to the prior-dominated constraint on β_s in the MCMC analysis, and to its consequent effect on the position of the MLE and the different width of $\Delta\chi^2$ (compare, e.g., the purple and red solid lines). Fixing β_s in the profile likelihood leads to equal or slightly tighter limits than the MCMC ones (see the pink, yellow, and red bars). Specifically, these small differences can be fully ascribed to the overcoverage of the

Bayesian limit near the boundary (Section 3), since we found that fixing β_s in the MCMC analysis produces the same upper limit as imposing the Gaussian prior on it.

We also checked the effect of fixing the Planck likelihood nuisance parameters to the mean values of their Gaussian prior (Tristram et al. 2021, 2022) in the profile likelihood analysis. For the Planck+BK18 combinations (with fixed β_s and conditioned HL covariance) this leads to the same upper limit and MLE as when fitting those nuisance parameters. In other words, the constraint is not prior-dominated and the nuisance parameters are constrained by the data.

We address the relevance of volume effects due to marginalization in the context of Bayesian inference for r. As evident from Figure 1 and Table 1, no substantial difference exists between MCMC marginalized limits and the prior-independent FC ones, as long as prior-dominated nuisance parameters such as β_s are fixed in the profile likelihood analysis. This suggests that volume effects do not play a prominent role in the Bayesian constraints. We note also that, because of the inefficiency of the Metropolis–Hastings algorithm in sampling near the boundary when a uniform positive prior is imposed on r, an apparent lower limit $r > 0$, which is entirely caused by the prior-dominated posterior, appears (Hergt et al. 2021). In fact, very small values of r are too sparsely sampled when such a prior is used. Achieving convergence (conventionally assessed via the Gelman–Rubin diagnostic) can thus be very time-consuming and require numerous evaluations of the likelihood (see Section 3.2 and footnote 11). In the available literature, MCMC is used to explore also negative values of r in order to try to circumvent this efficiency issue; however, this is unphysical as the model and the angular power spectra cannot be defined for negative r. The problem is exacerbated when one explores also the tensor spectral index n_T in the MCMC fit. These issues can be addressed, for instance, with the adoption of a logarithmic prior on r, introducing, however, a dependence of the constraints on the choice of the prior lower edge, as well as with the profile likelihood approach we adopt in this paper.

5. Conclusions

In this paper, we derived confidence intervals on r from the state-of-the-art CMB data sets Planck and BK18 via the frequentist profile likelihood method, and compared with the Bayesian MCMC procedure typically adopted in the literature. This is a useful robustness test for a potential future detection of r or for putting robust upper limits on this parameter, checking simultaneously for the dependence on priors and the volume effect upon marginalization in the Bayesian constraints. The profile likelihood is not affected by the inefficiency of the MCMC sampling near the boundary when a uniform prior is imposed on $r \geqslant 0$.

The profile likelihood and MCMC results did not agree for all the data combinations explored in the paper. Specifically, we reported an upper limit of $r < 0.042$ at 95% C.L. for the combination of Planck, BK18, BAO, and lensing with a conditioned Planck lowlEB covariance matrix as suggested in Beck et al. (2022). This limit is more conservative than the corresponding MCMC limit of $r < 0.038$. We find that the Bayesian constraint is driven by the Gaussian prior adopted for the synchrotron spectral index β_s in the BK18 likelihood. Fixing this nuisance parameter to the central value of the prior, $\beta_s = -3.1$, we eventually obtained an upper limit of $r < 0.037$ from the profile likelihood, slightly tighter than the MCMC limit because of the well-known overcoverage of Bayesian intervals near the parameter boundary.

We also confirmed the findings in Beck et al. (2022) regarding the conditioning of the Planck lowlEB covariance matrix: the additional scatter due to the limited number of simulations used in the covariance matrix construction moves the MLE of the lowlEB likelihood toward lower values, deceptively tightening the resulting upper limit.

The profile likelihood method is computationally more efficient than the MCMC, providing a useful alternative for a fast and robust evaluation of confidence limits near the physical boundary of a parameter.

We anticipate that the profile likelihood will represent a useful sanity check for prior effects in future and increasingly sensitive surveys, such as the BICEP array (Moncelsi et al. 2020), the Simons Observatory (Ade et al. 2019), the LiteBIRD satellite (LiteBIRD Collaboration 2022), and the CMB-S4 (Abazajian et al. 2016) experiments.

Appendix
Best-fitting Parameters

In Table 2 we compare the best-fitting parameters for the Planck+BK18+BAO+lensing combination and conditioned lowlEB covariance matrix with ("Fixed β_s" Column) and without ("Free β_s" Column) fixing the synchrotron spectral index β_s in the BK18 likelihood. See Section 2 and references therein for details on likelihoods and parameters used here.

Table 2
Best-fit Parameters for Planck + BK18 + BAO + Lensing with Conditioned HL Covariance, Fitting or Fixing β_s in the BK18 Likelihood

Parameter	Type	Best-fit (Free β_s)	Best-fit (Fixed β_s)
Cosmological parameters			
r	profile MLE	-0.0021	0.015
θ_{MC}	free	0.104062	0.0104062
$\log(10^{10}A_s)$	free	3.054	3.054
n_s	free	0.966	0.967
$\Omega_b h^2$	free	0.0223	0.0223
$\Omega_c h^2$	free	0.119	0.119
τ	free	0.0603	0.0604
A_s	derived	2.12×10^{-9}	2.12×10^{-9}
H_0	derived	67.43	67.43
σ_8	derived	0.813	0.813
BK18 nuisance parameters			
β_s	free or fixed	-2.0 (free)	-3.1 (fixed)
A_d	free	4.433	4.397
A_{sync}	free	0.17	0.517
α_d	free	-0.641	-0.657
α_s	free	-1.9×10^{-7}	-4.3×10^{-6}
β_d	free	1.500	1.484
ϵ	free	0.03	-0.131
Planck nuisance parameters			
A_{pl}	free (HiLLiPoP, lowlTT, lensing)	1.00189	1.00192
$c_0(100A)$	free (HiLLiPoP)	3.85×10^{-3}	3.86×10^{-3}
$c_1(100B)$	free (HiLLiPoP)	-1.0053×10^{-2}	-1.0053×10^{-2}
$c_3(143B)$	free (HiLLiPoP)	-1.0031×10^{-2}	-1.0026×10^{-2}
$c_4(217A)$	free (HiLLiPoP)	-1.0053×10^{-2}	-1.0053×10^{-2}
$c_5(217B)$	free (HiLLiPoP)	-4.419×10^{-3}	-4.417×10^{-3}
$A^{PS}(100 \times 100)$	free (HiLLiPoP)	2.620×10^2	2.619×10^2
$A^{PS}(100 \times 143)$	free (HiLLiPoP)	1.245×10^{-2}	1.244×10^2
$A^{PS}(100 \times 217)$	free (HiLLiPoP)	84.71	84.65
$A^{PS}(143 \times 143)$	free (HiLLiPoP)	53.09	53.05
$A^{PS}(143 \times 217)$	free (HiLLiPoP)	37.70	37.67
$A^{PS}(217 \times 217)$	free (HiLLiPoP)	74.39	74.41
A_{dust}^{100}	free (HiLLiPoP)	1.694×10^{-2}	1.688×10^{-2}
A_{dust}^{143}	free (HiLLiPoP)	3.966×10^{-2}	3.963×10^{-2}
A_{dust}^{217}	free (HiLLiPoP)	0.1322	0.1322
A_{SZ}	free (HiLLiPoP)	1.050	1.048
A_{CIB}	free (HiLLiPoP)	1.056	1.055
A_{kSZ}	free (HiLLiPoP)	7.722×10^{-5}	2.723×10^{-5}
$A_{SZ \times CIB}$	free (HiLLiPoP)	3.961×10^{-5}	3.582×10^{-6}
$c_2(143A)$	fixed (HiLLiPoP)	0.0	0.0
A_{radio}^{PS}	fixed (HiLLiPoP)	0.0	0.0
A_{dust}^{PS}	fixed (HiLLiPoP)	0.0	0.0
A_{dust}^{100T}	derived (HiLLiPoP)	1.694×10^{-2}	1.688×10^{-2}
A_{dust}^{143T}	derived (HiLLiPoP)	3.966×10^{-2}	3.963×10^{-2}
A_{dust}^{217T}	derived (HiLLiPoP)	0.1322	0.1322
A_{dust}^{100P}	derived (HiLLiPoP)	1.694×10^{-2}	1.688×10^{-2}
A_{dust}^{143P}	derived (HiLLiPoP)	3.966×10^{-2}	3.963×10^{-2}
A_{dust}^{217P}	derived (HiLLiPoP)	0.1322	0.1322
χ^2 values			
χ_{BAO}^2	...	17.87	17.84
χ_{BK18}^2	...	534.70	536.20
χ_{lowlEB}^2	...	156.50	155.65
$\chi_{hillipop}^2$	...	30346.11	30345.86
χ_{lowlTT}^2	...	22.90	23.36
$\chi_{lensing}^2$	...	8.71	8.72
χ_{tot}^2	...	31086.786	31087.63

ORCID iDs

Paolo Campeti ⓘ https://orcid.org/0000-0002-5637-519X
Eiichiro Komatsu ⓘ https://orcid.org/0000-0002-0136-2404

References

Abazajian, K. N., Adshead, P., Ahmed, Z., et al. 2016, arXiv:1610.02743
Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, JCAP, 02, 056
Ade, P. A. R., Amiri, M., Benton, S.J., et al. 2022, ApJ, 927, 174
Alam, S., Aubert, M., Avila, S., et al. 2021, PhRvD, 103, 083533
Albrecht, A., & Steinhardt, P. J. 1982, PhRvL, 48, 1220
ATLAS Collaboration 2013, Combined Measurements of the Mass and Signal Strength of the Higgs-like Boson with the ATLAS Detector using up to 25 fb^{-1} of Proton-proton Collision Data, Report, ATLAS-CONF-2013-014, CERN
Beck, D., Cukierman, A., & Wu, W. L. K. 2022, MNRAS, 515, 229
BICEP/Keck Collaboration 2021, PhRvL, 127, 151301
Campeti, P., Komatsu, E., Poletti, D., & Baccigalupi, C. 2021, JCAP, 01, 012
Couchot, F., Henrot-Versillé, S., Perdereau, O., et al. 2017, A&A, 602, A41
Cousins, R. D. 1995, AmJPh, 63, 398
Cousins, R. D. 2018, arXiv:1807.05996
Feldman, G. J., & Cousins, R. D. 1998, PhRvD, 57, 3873
Fuskeland, U., Wehus, I. K., Eriksen, H. K., & Næss, S. K. 2014, ApJ, 790, 104
Gómez-Valent, A. 2022, PhRvD, 106, 063506
Grishchuk, L. P. 1974, ZhETF, 67, 825
Guth, A. H. 1981, PhRvD, 23, 347
Hamann, J., Hannestad, S., Raffelt, G. G., & Wong, Y. Y. Y. 2007, JCAP, 08, 021
Hamimeche, S., & Lewis, A. 2008, PhRvD, 77, 103013
Hergt, L. T., Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2021, PhRvD, 103, 123511
Herold, L., Ferreira, E. G. M., & Komatsu, E. 2022, ApJL, 929, L16
Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, PhRvL, 78, 2058
Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473
Linde, A. D. 1982, PhLB, 108, 389
LiteBIRD Collaboration 2022, PTEP, 11, ptac150
Mangilli, A., Plaszczynski, S., & Tristram, M. 2015, MNRAS, 453, 3174
Moncelsi, L., Ade, P.A.R., Ahmed, Z., et al. 2020, Proc. SPIE, 11453, 1145314
Particle Data Group 2020, PTEP, 2020, 083C01
Planck Collaboration Int. LVII 2020, A&A, 643, A42
Planck Collaboration Int. XVI 2014, A&A, 566, A54
Planck Collaboration V 2020, A&A, 641, A5
Planck Collaboration VI 2020, A&A, 641, A6
Planck Collaboration VIII 2018, A&A, 641, A8
Planck Collaboration XI 2016, A&A, 594, A11
Planck Collaboration XV 2013, A&A, 571, A15
Sato, K. 1981, MNRAS, 195, 467
Seljak, U., & Zaldarriaga, M. 1997, PhRvL, 78, 2054
Sellentin, E., & Heavens, A. F. 2016, MNRAS, 456, L132
Starobinsky, A. A. 1979, JETPL, 30, 682
Torrado, J., & Lewis, A. 2021, JCAP, 05, 057
Tristram, M., Banday, A.J., Górski, K. M., et al. 2021, A&A, 647, A128
Tristram, M., Banday, A.J., Górski, K. M., et al. 2022, PhRvD, 105, 083524
Trotta, R. 2017, arXiv:1701.01467
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261