

Robust constraints on tensor perturbations from cosmological data: A comparative analysis from Bayesian and frequentist perspectives

Giacomo Galloni^{1,2,3,4,*} Sophie Henrot-Versillé⁵ and Matthieu Tristram⁵

¹*Dipartimento di Fisica, Università di Roma Tor Vergata,
Via della Ricerca Scientifica, 1, 00133 Roma, Italy*

²*Dipartimento di Fisica, Università La Sapienza, P.le A. Moro 2, 00185 Roma, Italy*

³*Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Roma 2,
Via della Ricerca Scientifica, 1, 00133 Roma, Italy*

⁴*Dipartimento di Fisica e Scienza della Terra, Università degli Studi di Ferrara,
via Giuseppe Saragat 1, 44122 Ferrara, Italy*

⁵*Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France*

 (Received 8 May 2024; accepted 12 August 2024; published 4 September 2024)

We analyze primordial tensor perturbations using the latest cosmic microwave background and gravitational waves data, focusing on the tensor-to-scalar ratio, r , and the tensor spectral tilt, n_t . Utilizing data from Planck PR4, BICEP/Keck, and LIGO-Virgo-KAGRA, we employ both Bayesian and frequentist methods to provide robust constraints on these parameters. Our results indicate more conservative upper limits for r with profile likelihoods compared to Bayesian credible intervals, highlighting the influence of prior selection and volume effects. The profile likelihood for n_t shows that the current data do not provide sufficient information to derive quantitative bounds, unless extra assumptions on r are used. Additionally, we conduct a 2D profile likelihood analysis of r and n_t , indicating a closer agreement between both statistical methods for the largest values of r . This study not only updates our understanding of the tensor perturbations but also highlights the importance of employing both statistical methods to explore less constrained parameters, crucial for future explorations in cosmology.

DOI: [10.1103/PhysRevD.110.063511](https://doi.org/10.1103/PhysRevD.110.063511)

I. INTRODUCTION

The cosmic microwave background (CMB) is a relic from the distant past, containing within its faint signals the secrets of the Universe's beginnings [1–3]. Of particular interest is B-mode polarization, which detection would reveal the presence of primordial gravitational waves (GWs), predicted by inflationary scenarios [4–6].

For this reason, observing B modes has been the goal of past, present, and future experiments such as the BICEP/Keck Array [7], the Simons Observatory [8], CMB-S4 [9], and the Light satellite for the study of B-mode polarization and inflation from cosmic microwave background radiation detection (LiteBIRD) [10,11]. On top of that, *Planck* represents the cornerstone of modern cosmology [12,13] and has recently provided the most precise constraints on the standard Λ CDM model with its fourth release (PR4) [14,15]. In this last iteration, PR4 also provides information on B modes, which, in combination with the measurements of BICEP/Keck, represents our most advanced knowledge on primordial GWs exploiting CMB data [16,17].

In parallel, GW interferometers are also proving to be a very insightful tool for providing complementary information on the primordial power spectrum of tensor modes [18]. Indeed, CMB experiments are notably not very effective in constraining its tilt (n_t). This is due to the lensing signal dominating the B-mode spectrum for $\ell > 300$. Although GW interferometers, such as LIGO-Virgo-KAGRA (LVK) [19–21], are designed to detect ripples in spacetime emanating from cataclysmic cosmic events, they also provide a unique synergy with CMB observations. Their ability to probe the GW energy density in a frequency range totally different from that of the CMB enhances our chances of constraining the primordial Universe scenarios [22,23].

The primordial scalar and tensor perturbations are customarily parametrized with power laws as

$$P_s(k) = A_s \left(\frac{k}{k_s} \right)^{n_s-1} \quad \text{and} \quad P_t(k) = A_t \left(\frac{k}{k_t} \right)^{n_t}, \quad (1)$$

where A_s (A_t), n_s (n_t), and k_s (k_t) are respectively the scalar (tensor) amplitude, spectral tilt, and pivot scale. On top of this, the tensor-to-scalar ratio at a generic scale is defined as

*Contact author: giacomo.galloni@unife.it

$$r_k = \frac{P_t(k)}{P_s(k)}, \quad (2)$$

which, together with n_t , represent the tensor sector of parameter space. The state-of-the-art knowledge on this sector can be divided into two cases: in the first, the spectral tilt is fixed to the so-called consistency relation of single-field slow-roll inflation $n_t = -r_{0.05}/8$ [4], and only $r_{0.05}$ is constrained through the data. In this context, the current lowest limit on $r_{0.05}$ is $r_{0.05} < 0.032$ at 95% confidence level (CL) based on the combination of Planck and BICEP/Keck data [16]. On the other hand, one can avoid the choice of single-field slow-roll inflation violating the relation just mentioned. In this case, the current bounds are $r_{0.01} < 0.028$ and $-1.37 < n_t < 0.42$ at 95% CL, accounting also for LVK data [17].

In a scenario where we do not have a detection of primordial GWs, the choice of the statistical analysis method used to extract information from the data is crucial. Bayesian methods, based on Markov chain Monte Carlo (MCMC), have been pivotal [24–28]. However, challenges arise, particularly in dealing with volume effects that can mislead our conclusions in multidimensional parameter spaces [29]. This is particularly the case for the tensor sector as we shall see in the rest of the paper. Indeed, for a given sensitivity, the constraints on n_t can be very broad if r is pushed to very low values.

This work provides an update on the current knowledge on the parameters of the primordial tensor power spectrum using the latest *Planck* release (PR4). After employing a MCMC approach to obtain information about the posterior probability, we gauge the impact of volume effects and prior choices by using the frequentist approach based on profile likelihoods (PLs), which offers independence from both [30–34]. Indeed, MCMC and PL answer different and complementary questions, delivering us with a complete picture of the tensor parameter space.

In Sec. II, we offer a concise overview of the datasets utilized in this study, encompassing contributions from the CMB derived from *Planck* and BICEP/Keck, as well as insights from a GW standpoint, particularly from the LVK collaboration. Section III discusses the methodologies underpinning both the MCMC and the PL analyses. Special attention is paid to the techniques employed to extract confidence intervals, ensuring their adherence to the correct statistical properties. Finally, Sec. IV presents our findings, revisiting the outcomes of [17] through the lens of updated datasets using the MCMC perspective. Simultaneously, we introduce a novel frequentist exploration of the tensor sector of parameter space using PL, shedding light on volume effects and prior-choice dependencies that impact Bayesian results.

II. DATASETS

A. Data

1. *Planck*

This study uses *Planck*'s PR4 maps.¹ They were produced by the NPIPE processing pipeline, which reconstructs temperature- and polarization-calibrated frequency maps from *Planck*'s LFI and HFI data. The NPIPE process, detailed in [36], incorporates data from previously neglected repointing periods, along with several enhancements that reduce noise and systematic errors across frequencies and component-separated maps, enhancing consistency between different frequencies.

We use several likelihoods that cover the multipole range from $\ell = 2$ to $\ell = 2500$. For large angular scales in temperature ($\ell \leq 30$), we consider the Commander *TT* likelihood (lowIT) based on a Bayesian posterior Gibbs sampling that combines the separation of astrophysical components and the estimation of likelihood [13,37]. For large angular scales in polarization, we use the low- ℓ likelihood polarized for *Planck* (LoLLiPoP) based on the Hamimeche-Lewis approximation for the *EE*, *BB*, and *EB* power spectra [16], which can cover multipoles $\ell \leq 150$. The cross-correlation spectrum *TE* is not used at large scales. At small angular scales ($\ell > 30$), we alternatively use the high- ℓ likelihood polarized for *Planck* (HiLLiPoP) or CamSpec which both combine the *TT*, *TE*, and *EE* CMB spectra over a large fraction of the sky (75% and 80%, respectively). HiLLiPoP is a multifrequency likelihood based on cross spectra of the 100, 143, and 217 GHz frequency maps, with astrophysical models for the residuals of foreground emissions [15]. CamSpec is based on cross spectra at 143 and 217 GHz which are preprocessed by a cleaning procedure using the 545 GHz maps as a template of Galactic dust emission [14]. Each likelihood comes together with its own nuisance parameters mostly related to instrumental calibration and residual foreground modeling.

2. BICEP/Keck Array

We use the BICEP/Keck likelihood (BK18), representing data collected by the BICEP2, Keck Array and BICEP3 CMB polarization experiments up to the 2018 observing season [38]. This likelihood is based on the Hamimeche-Lewis approximation [39] for the joint likelihood of the *BB* autospectra and cross spectra obtained across multiple frequency maps: BICEP/Keck (two at 95 GHz, one each at 150 and 220 GHz), WMAP (23 and 33 GHz), and *Planck* (PR4 at 30, 44, 143, 217, and 353 GHz). Covering an effective area of roughly 400 square degrees (equivalent to 1% of the sky), this dataset is centered on a region characterized by minimal foreground

¹This is available on the *Planck* Legacy Archive: [35].

emission. The data model encompasses Galactic dust and synchrotron emission, incorporating correlations between dust and synchrotron components.

3. LIGO-Virgo-KAGRA interferometers

We adopt the same approach as described in [17,18] taking advantage of the fact that GW interferometers probe scales ($k \sim 10^{16} \text{ Mpc}^{-1}$) nearly 18 orders of magnitude above those probed by the CMB ($k \sim 10^{-2} \text{ Mpc}^{-1}$). Under the assumption that the tensor tilt n_t remains constant across the large frequency range between the CMB and GW interferometers, we use the upper limit on the energy density of GWs (Ω_{GW}) provided by the LVK collaboration to obtain a constraint on n_t at small scales, with [18,40]

$$\Omega_{\text{GW}}(k) = \frac{r_{0.05} A_s}{24 z_{\text{eq}}} \left(\frac{k}{k_t} \right)^{n_t}, \quad (3)$$

where we considered the pivot scale to be $k_t = k_s = 0.05 \text{ Mpc}^{-1}$ and A_s is the scalar amplitude defined in Eq. (1). Then, we define a Gaussian likelihood on the energy density of GWs,

$$-2 \log(\mathcal{L}_{\text{LVK}}) = \frac{(\Omega_{\text{GW}} - \mu_{\text{LVK}})^2}{\sigma_{\text{LVK}}^2}, \quad (4)$$

predicted as centered on $\mu_{\text{LVK}} = 0$ and with a σ_{LVK} that equals half of the LVK's 95% bound [18].

In particular, we use the limit $\Omega_{\text{GW}}(25 \text{ Hz}) < 6.6 \times 10^{-9}$ (95% CL) provided in [23] using data from the third observing run of Advanced LIGO and Advanced Virgo (O3) combined with upper limits of the previous runs O1 and O2.

B. Combinations of data

Throughout the remainder of this paper, we refer to “PLK20(CamSpec/HiLLiPoP)” as the combination of *Planck* likelihoods:

- (i) *Planck* PR3 low- ℓ TT [13,37],
- (ii) *Planck* PR4 LoLLiPoP [16],
- (iii) *Planck* PR4 lensing [41],
- (iv) *Planck* PR4 high- ℓ (CamSpec [14] or HiLLiPoP [15]).

We then add the two other datasets:

- (i) BICEP/Keck array 2018 (BK18) [38],
- (ii) LIGO-Virgo-KAGRA 2021 (LVK21) [23],

to form the two main combinations we study in this work, PLK20(CamSpec) + BK18 + LVK21 and PLK20(HiLLiPoP) + BK18 + LVK21. Furthermore, we will consider the same combinations without LVK21 to emphasize the role of GW interferometers.

Note that we slightly modified the high- ℓ likelihoods (both HiLLiPoP and CamSpec) to avoid any correlation with the low- ℓ LoLLiPoP. Indeed, to maximize the

information coming from *BB*, we use LoLLiPoP up to $\ell = 150$ and consequently adjust the minimum multipole of high- ℓ likelihoods for *EE* at $\ell_{\text{min}} = 151$.

As previously done in [16], we neglect correlations between *Planck* and BICEP/Keck datasets and simply multiply the likelihood distributions. This is justified because the BK18 spectra are estimated on 1% of the sky, while the *Planck* analysis is derived from at least 50% of the sky.

III. METHODOLOGY

In this work, we explore both the tensor-to-scalar ratio r and the tensor spectral tilt n_t . This allows us to test the consistency relation imposed by the single-field slow-roll inflationary model ($n_t = -r/8$ at first order), leaving instead n_t free to vary. The only assumption we make is that the tensor power spectrum can still be described as a power law [see Eq. (1)]. In a context where n_t is fixed, the pivot is usually chosen to match the scalar pivot scale of $k_t = k_s = 0.05 \text{ Mpc}^{-1}$. Instead, in our case the pivot is taken to be $k_t = 0.01 \text{ Mpc}^{-1}$, since this scale is close to the decorrelation scale [18].

It is important to emphasize that we also fit for the six parameters of the Λ CDM cosmological model as well as all nuisance parameters associated to the likelihoods (see Sec. II B and the references therein for more details). For this reason, we refer to the model we study in this paper as $\Lambda\text{CDM} + r_{0.01} + n_t$. Then, taking into account the nuisance parameters, the size of the parameter space considered is of the order of ~ 30 dimensions.

To compute the angular power spectra of all the observables considered here, we employ CAMB [28,42,43].

Before discussing the details of the Bayesian and frequentist approaches to parameter estimation, it is essential to understand that these two methodologies address complementary questions. Bayesian credible intervals provide information about the probability that the true value of a parameter falls within a given interval. Conversely, frequentist confidence intervals inform us about the probability of obtaining the observed data. Thus, together they provide a complete picture on the parameter space under consideration, but they must be compared with care. In fact, for example, [31,44,45] shows the importance of comparing and confronting results from both methods in order to gain a comprehensive understanding of the data and validate the robustness of the conclusions.

A. MCMC analysis

MCMC, a versatile and indispensable tool of Bayesian statistics, has become fundamental in the navigation of intricate landscapes of probability distributions [24–28]. Its role in exploring large parameter spaces becomes particularly apparent when seeking information about individual parameters. Indeed, such a procedure provides us with the marginalized posterior of a parameter given a

set of data, which we can use to find the degree of belief of such a parameter.

If a physical boundary is present (as in our case with $r_{0.01}$), it is naturally encoded in a Bayesian framework by the prior distribution. Indeed, considering a generic parameter named θ and a dataset x , the posterior $P(\theta|x)$ is obtained through the Bayes theorem, which reads

$$P(\theta|x) = \frac{\mathcal{L}(x|\theta)\Pi(\theta)}{\mathcal{E}(x)}. \quad (5)$$

Here, $\mathcal{L}(x|\theta)$ is the likelihood, $\Pi(\theta)$ is the prior, and $\mathcal{E}(x)$ the evidence. A positivity bound on θ can be easily accounted for choosing

$$\Pi(\theta) = \begin{cases} 0 & \theta < 0 \\ 1 & \theta \geq 0. \end{cases} \quad (6)$$

All this comes without ambiguity in Bayesian statistics, as the random quantity evaluated here is θ and we are reconstructing its degree of belief.

Specifying the discussion to the tensor sector, [17] reports a comprehensive analysis on the two most used approaches for $r_{0.01}$ and n_r . Here, we follow the same procedure chosen in that work: Single-scale approach (SSA). This implies the introduction of a lower cutoff in the tensor-to-scalar ratio of $r_{0.01} > 10^{-5}$. The underlying idea is that this amplitude level is far from detectable with current experiments, thus we cut the prior to avoid pathological behaviors of the MCMC. From a physical point of view, we know that some minimal B-mode signal is present even in single-field slow-roll inflation with no primordial tensor modes. Indeed, we have measured scalar perturbations of our Universe, which source GWs at second order in the perturbations. In particular, those would produce a B-mode signal equivalent to $r \simeq 10^{-7}$ on large scales and $r \simeq 10^{-5}$ on small scales [46–50].

Despite all this, in a Bayesian framework, volume effects can have an important role. In a multidimensional problem, if a large part of the probability volume is in a certain area, the final posterior will be drawn towards that region just as a result of the marginalization procedure. On top of this, excluding *a priori* parts of the parameter space (as we do, following the SSA prescription) could also bring to differences in the final results, highlighting the prior dependence of Bayesian statistics (see [17]). In this context, another possible prior choice for a parameter whose order of magnitude is unknown is the log-uniform prior. This translates into equal weighting of the order of magnitude of this parameter [51].

These aspects can pose challenges in accurately gauging the significance of the obtained results and, possibly, mislead the derived conclusions. For this reason, in this work we confront and compare those results with the ones derived with the frequentist approach, i.e., the PL.

B. Profile likelihood

The procedure to get confidence intervals in a frequentist framework is slightly more involved due to the nature of frequentist intervals: in fact, the true values of parameters are not random variables, but rather fixed values that nature chooses. Thus, the boundaries obtained from a set of data are specific for the experiment considered and represent the random variable in this case. Repeating the experiment would cause these bounds to fluctuate. The “coverage probability” refers to the fraction of intervals that contain the true value of the parameters among the N different repetitions of an experiment. The confidence intervals are then determined to have a coverage probability greater than or equal to a certain confidence level (CL) [52].

To derive a confidence region that has the correct frequentist coverage properties, one can make use of likelihood ratio statistics. For multiparametric spaces, this amounts to constructing the PL: for fixed values of the parameter of interest (θ_i), we look for the maximum of the likelihood function in all the other dimensions (both for the physical parameters and the nuisance parameters). We then have access to the function $\chi_{\min}^2(\theta_i) = -2 \ln \mathcal{L}_{\max}(\theta_i)$, where \mathcal{L} is the likelihood considered. The best fit [or $\min(\chi_{\min}^2(\theta_i))$] gives the estimate of the parameter under consideration, which also corresponds to the maximum likelihood over all the other parameters.

This procedure also ensures that its determination is independent of any change of variable $f(\theta_i)$, making it parametrization invariant.

The error in the parameter θ can be deduced from the shape of the $\chi_{\min}^2(\theta_i)$ function. For Gaussian distribution, the function is parabolic and the 1σ error bounds are simply obtained by a cut at $\Delta\chi_{\min}^2 = \chi_{\min}^2(\theta_i) - \min(\chi_{\min}^2(\theta_i)) = 1$. When dealing with physical boundaries, one needs to use the Feldman-Cousins (FC) prescription and the Neyman construction [52,53]. Note that no integration in parameter space is performed to obtain the confidence intervals. Indeed, this is the manifestation of the core difference between Bayesian and frequentist approaches to parameter estimation mentioned above. The former requires to marginalize a posterior to extract bounds on a single parameter, while the latter does not.

One of the difficulties in building the PL is the precision with which we need to determine the values of $\chi_{\min}^2(\theta_i)$. We must rely both on a very accurate minimizer and a boost to the accuracy parameters of the Boltzmann code CAMB [36].

The PL procedure also generates a by-product, the “coprofiles.” In fact, we not only obtain a function $\chi_{\min}^2(\theta_i)$, but also $N - 1$ functions $\theta_j(\theta_i)$ where $j \neq i$. These coprofiles allow us to gauge the direction of degeneracy of other parameters on the profiled one and are very useful as a diagnostic tool of the PL procedure [31,34]. In other words, coprofiles allow us to explore the parameter space in the direction of the minimum $-2 \log \mathcal{L}$ valley.

C. Feldman-Cousins prescription

The first step of the FC prescription is to assume a true value of the parameter of interest θ . In fact, for each value of the parameter we can find an interval $[x_1(\theta, \alpha), x_2(\theta, \alpha)]$ such that

$$\alpha = \int_{x_1}^{x_2} \mathcal{L}(x|\theta) dx. \quad (7)$$

Repeating this process for a set of “true values” of θ and drawing each segment $[x_1, x_2]$ in a plot, we can find a “confidence belt.” Now, having a measured value of $x = x_0$, the confidence interval $[\theta_1(x), \theta_2(x)]$ is found by drawing a vertical line at x_0 and looking at the maximum and minimum values of θ (or its relative segment) that intersect this line.

Note that, actually, Eq. (7) does not define uniquely x_1 and x_2 , therefore, we need another equation to close the system. Here is where the FC prescription is defined. Consider the following test statistic:

$$\lambda(x, \theta) = \frac{\mathcal{L}(x|\theta)}{\mathcal{L}(x|\hat{\theta})}, \quad (8)$$

i.e., a likelihood ratio where $\hat{\theta}$ maximizes $\mathcal{L}(x|\theta)$. Solving Eq. (7) while asking $\lambda(x_1, \theta) = \lambda(x_2, \theta)$ represents the FC construction.

This prescription not only recovers the correct coverage probability, but it is also able to naturally account for any eventual physical boundaries of the parameters, shifting “automatically” from a two-sided interval to a one-sided one. Suppose that

$$\mathcal{L}(x|\theta) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(x-\theta)^2}{2}\right) \quad (9)$$

and that θ must be non-negative. Our physically allowed best estimate of θ is $\hat{\theta} = \max(0, x)$; then, we can write

$$\lambda(x, \theta) = \frac{\mathcal{L}(x|\theta)}{\mathcal{L}(x|\hat{\theta})} = \begin{cases} \exp\left(-\frac{(x-\theta)^2}{2}\right) & \text{if } x \geq 0 \\ \exp\left(x\theta - \frac{\theta^2}{2}\right) & \text{if } x < 0. \end{cases} \quad (10)$$

Now, if the measured value of x is too close to the physical limit or is negative, the correspondent vertical line would intercept only $x_1(\theta, \alpha)$, automatically defining an upper limit on θ . At some point, increasing the measured value, we would instead intercept also $x_2(\theta, \alpha)$, shifting to a two-sided confidence interval. For this reason, the FC intervals are also said to be *unified* (see [53] for more details).

IV. RESULTS

In this section, we fit for the parameters of the Λ CDM model extended to tensor perturbations with the dataset described in Sec. II and the methodology presented in Sec. III.

In particular, we focus on the tensor sector of the parameter space, i.e., $\{r_{0.01}, n_t\}$. As regards the Λ CDM parameters, we do not find any significant deviation from the current state-of-the-art results [14,15] (see Appendix A). We mention that we use a modified version of COBAYA [54] to obtain both Bayesian and frequentist results. For the latter, we added a new minimizer based on MINUIT² [55], which outperformed the two alternatives implemented in Cobaya (PY-BOBYQA [56] and SCIPY [57]).

Various configurations of CAMB and MINUIT have been tested in terms of reliability to obtain the absolute maximum of the likelihoods. More details can be found in Appendix B.

Note that very recently many new tools have been developed to perform PLs, which are not exploited in this work [58–60].

A. MCMC analysis

Figure 1 shows the 2D posterior distributions for $r_{0.01}$ and n_t from the MCMC analysis for the different datasets.

We start by comparing the results between different choices on the high- ℓ part of CMB, i.e., CamSpec or HiLLiPoP. In particular, the 95% CL marginalized intervals for PLK20(CamSpec) + BK18 + LVK21 are

$$r_{0.01} < 0.028 \quad \text{and} \quad -1.36 < n_t < 0.42, \quad (11)$$

while for PLK20(HiLLiPoP) + BK18 + LVK21

$$r_{0.01} < 0.029 \quad \text{and} \quad -1.39 < n_t < 0.41. \quad (12)$$

These intervals are almost identical, suggesting that this choice does not significantly impact the tensor sector, as expected. Furthermore, these are also very similar to the results presented in [17] which used a slightly different combination of datasets indicating the robustness of the derived limits regarding the choice of high- ℓ CMB likelihood.³ The bounds we obtain here for n_t are typically broader than the ones reported in [18]. We note that in [18], the tensor perturbations are characterized by fitting r at two different scales which can be affected by prior effects as discussed in [17].

As a test to gauge the importance of GWs interferometers, we also repeat the analysis removing LVK21. As expected, LVK21 severely constrains the tilt. For PLK20(CamSpec) + BK18 we obtain

²<https://github.com/CobayaSampler/cobaya/pull/332>.

³For completeness, [17] used a combination of datasets made up of LoLLiPoP of PR4 + *plik* and lensing from PR3 + BK18 + LVK21 (see Sec. II).

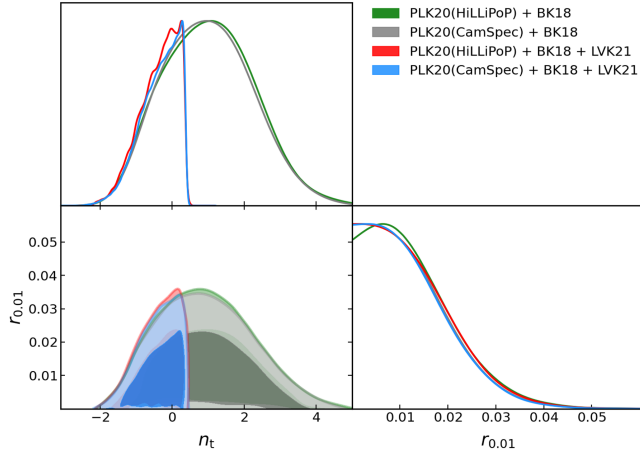


FIG. 1. 2D posterior of $r_{0.01} - n_t$ for *baseline* + CamSpec and *baseline* + HiLLiPoP.

$$r_{0.01} < 0.029 \quad \text{and} \quad -1.33 < n_t < 3.37, \quad (13)$$

while using PLK20(HiLLiPoP) + BK18 we obtain

$$r_{0.01} < 0.030 \quad \text{and} \quad -1.35 < n_t < 3.40. \quad (14)$$

Once again, these bounds are similar to the ones obtained in [17]. The only sizable difference lies in the fact that the tensor spectral tilt has slightly shifted to higher values. We have investigated the causes of such a shift but have been unable to identify any specific behavior of the likelihood function, or distinctive features within the parameter space.

In conclusion of this section, we reassessed the state-of-the-art results using the most updated datasets (see Sec. II). We found that the constraints on $r_{0.01}$ and n_t have remained essentially stable as compared to [17], demonstrating that the constraints are driven mainly by the low- ℓ part of CMB polarization.

B. Profile likelihood

We have previously noted that results from the Bayesian analysis may be affected by volume effects or by the prior choice. In this section, we perform a frequentist analysis on the tensor sector of parameter space to gauge such effects.

1. Tensor-to-scalar ratio

First, we calculate the PL on the tensor-to-scalar ratio. Our underlying model is Λ CDM + $r_{0.01} + n_t$, so at each fixed value of $r_{0.01}$, the likelihoods are maximized both with respect to the standard Λ CDM parameters and n_t (together with all the nuisance parameters in each likelihood).

The left panel of Fig. 2 reports the results of PLK20 \times (CamSpec) + BK18 + LVK21 and PLK20(HiLLiPoP) + BK18 + LVK21 in terms of both $\Delta\chi_{\min}^2(r_{0.01})$ and likelihood ratio $\mathcal{L}(r_{0.01})/\mathcal{L}_{\max}$ to emphasize the correspondence between these two quantities. On the other hand, the right panel shows the profile likelihood ratio for all the datasets defined above. All of them are very similar. However, one can note that HiLLiPoP results in a slightly higher upper limit with respect to CamSpec. This resonates with the MCMC results of Sec. IV A, where PLK20 \times (HiLLiPoP) + BK18 + LVK21 gives slightly broader bounds. Indeed, even though the two profiles show quite similar widths, the HiLLiPoP case peaks at $r_{0.01} = 0.01$ while CamSpec at $r_{0.01} = 0.009$.

Figure 2 shows also the comparison between the posteriors from Sec. IV A and PL distributions, suggesting that some volume effect is present in the Bayesian framework, pushing the posterior toward $r = 0$.

We emphasize that all PLs point to a best-fit value for the tensor-to-scalar ratio of $r_{0.01} \simeq 0.01$ while being completely consistent with a nondetection (i.e., $r_{0.01} = 0$). Instead, the maximum posterior tends to be around $r_{0.01} \simeq 0$ for every




FIG. 2. Left: $\Delta\chi_{\min}^2(r_{0.01})$ and $\mathcal{L}(r_{0.01})/\mathcal{L}_{\max}$ for PLK20(CamSpec) + BK18 + LVK21 and PLK20(HiLLiPoP) + BK18 + LVK21. Right: profile likelihood ratio (solid lines) and posterior distributions (dashed lines) for $r_{0.01}$.

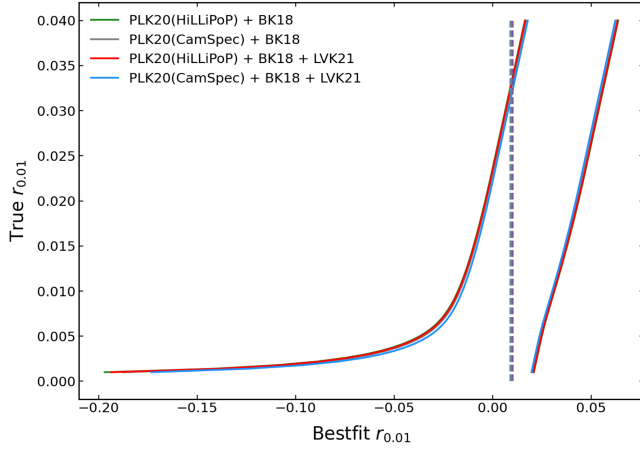


FIG. 3. FC belts obtained including the positivity condition on $r_{0.01}$. The vertical dashed lines are the best-fit values obtained from each dataset.

dataset (the only exception being PLK20(HiLLiPoP) + BK18, whose maximum posterior approaches the best fit).

2. Feldman-Cousins for $r_{0.01}$

We now apply the FC prescription to recover the confidence interval with the correct coverage probability and accounting for the physical limit $r_{0.01} \geq 0$.

For a Gaussian distribution of data given a parameter θ , we can include a physical bound in a frequentist framework by using the test statistic introduced by Eq. (10).

However, the PL ratios we obtain from both CamSpec and HiLLiPoP do not follow a Gaussian distribution far away from the best-fit value of $r_{0.01}$ (i.e., for $\Delta\chi^2 \gtrsim 2$). Indeed, we identify three distinct features in our PLs (see the right panel of Fig. 2): all of them present a relatively large probability tail, PLK20(HiLLiPoP) + BK18 + LVK21 shows a bump at low values of $r_{0.01}$, and the likelihood ratio for the combinations without LVK21 seems to drop faster than a Gaussian on the left of the best fit. We show in Appendix C that accounting for these features with a more complex approach with respect to the one presented in Sec. III B makes no significant difference.

Finally, in Fig. 3 we report the results for all the datasets considered. Intersecting these FC belts with the respective best-fit values, we obtain the upper bounds given in Table I.

TABLE I. 95% confidence intervals obtained with the FC prescription.

Dataset	95% confidence interval
PLK20(HiLLiPoP) + BK18	$r_{0.01} < 0.033$
PLK20(CamSpec) + BK18	$r_{0.01} < 0.032$
PLK20(HiLLiPoP) + BK18 + LVK21	$r_{0.01} < 0.033$
PLK20(CamSpec) + BK18 + LVK21	$r_{0.01} < 0.032$




FIG. 4. PL on the tensor spectral tilt. Both PLK(CamSpec) + BK18 + LVK21 and PLK(HiLLiPoP) + BK18 + LVK21 are normalized to the minimum of their counterpart without LVK21.

None of these limits are affected by the volume effect, or the prior choice, previously discussed in Sec. IV B 1. As a result, the upper limits of the PL are more conservative than those derived with the Bayesian analysis (see Sec. IV A).

3. Tensor spectral tilt

The profile likelihood on n_t is shown on Fig. 4 for various combinations of the datasets. Given the shape of the PLs, the derivation of quantitative upper bounds is beyond the scope of this work, and we shall nevertheless go to a qualitative description of the obtained results.

First, PLs without GW data show a mild preference for a best fit around $n_t \simeq 2$. However, the fact that $\Delta\chi^2$ is always $\lesssim 1$ means that current data are not sufficiently constraining to statistically disentangle the different values of n_t . In fact, the part at low n_t exhibits a constant- χ^2 region, which we will call “plateau” and will be discussed below.

The addition of LVK21 constrains the value of n_t to be lower than $n_t \lesssim 0.4$ due to the huge level arm in the frequency domain. Given the low level of statistical significance, we would not interpret this as a sign for a nonconstant spectral index. To emphasize this cutoff given by LVK21, we normalize the PLK20(CamSpec/HiLLiPoP) + BK18 + LVK21 profiles to the minimum of their counterpart without GW data. This also emphasizes that the corresponding plateaus are found at the same value of χ^2_{\min} . Furthermore, LVK21 also causes another plateau at the same χ^2_{\min} level with respect to the previously mentioned one, but this time for blue tilts.

Finally, we note that HiLLiPoP shows a “bump” at $n_t \sim -1$ and that the plateaus are slightly higher compared to the minimum χ^2 when considering HiLLiPoP rather than CamSpec, suggesting that HiLLiPoP is slightly more constraining n_t compared to CamSpec.

The plateaus observed in the PL result from the fact that the combination of likelihoods is no longer sensitive to the




FIG. 5. Coprofiles of $r_{0,01}$ as a function of the profiled value of n_t .

value of n_t in this region. This can be understood by looking at the coprofile $r_{0,01}(n_t)$ (Fig. 5). Indeed, in order to accommodate extreme values of n_t , the tensor-to-scalar ratio is suppressed to very small values. At some point, the likelihoods are no longer sensitive to whether tensors are there or not, and the χ^2 just sits on the same value for each tilt.

This is exactly the reason behind the analysis of [17]. Indeed, if we do not impose a lower cutoff in $r_{0,01}$, the datasets we consider are not sufficient to constrain n_t given that one can always find a value of $r_{0,01}$ low enough to accommodate any tilt. As a consequence, we do not have a clear upper or lower bound on the n_t profile (Fig. 4). On the contrary, introducing a cutoff allows to get constraints on n_t as illustrated in Fig. 6. First, note that as soon as the tensor-to-scalar ratio is restricted, $\Delta\chi_{\min}^2$ diverges, indicating that we could recover both an upper and a lower bound. In fact, with $r_{0,01} = 10^{-7}$ (respectively 10^{-5}), at some point the tilt

will be so red that the corresponding B modes should be observed by *Planck* or BICEP/Keck array. Despite this, once again we stop at a qualitative description of the profile in n_t , as a further study is necessary for an actual confidence interval.

The same reason explains the difference between the posterior and the profile on n_t (see Fig. 6). Indeed, the values of $r_{0,01}$ found by the PL minimization are so low ($r_{0,01} \sim 10^{-10}$ – 10^{-17}) that it would be impossible to explore that region with our MCMC using a flat prior. The region explored by the MCMC corresponds to the region of the coprofile with the highest values of $r_{0,01}$, which also corresponds to the region of the PL between the plateaus (see Fig. 6). Indeed, it is a well-known feature of uniform priors to poorly explore the region where the parameter of interest is extremely close to the boundaries. To showcase this, we show in the right panel of Fig. 6 the posterior obtained imposing a log-uniform prior on $r_{0,01}$. This allows us to better explore the region of very low tensor-to-scalar ratio and results in a broader distribution of n_t .

Once again, this comparison highlights a possible flaw of the Bayesian analysis: a tilt of $n_t = -2$ is almost excluded by our MCMC analysis, however, it is not statistically excluded by the PL analysis as it only corresponds to $\Delta\chi_{\min}^2 \simeq 0.7$. This is not the case if we look at the log-uniform results. In fact, values as low as $n_t \simeq -3$ are still accepted and correspond to $\Delta\chi_{\min}^2 \gtrsim 3$, due to a better representation of the low- r region. Despite this, such a choice is known to provide extremely underestimated bounds on $r_{0,01}$ [51].

C. 2D profile likelihood

In order to explore deeper the $(r_{0,01}, n_t)$ plane, we construct the 2D profile likelihood, fitting the best fit over

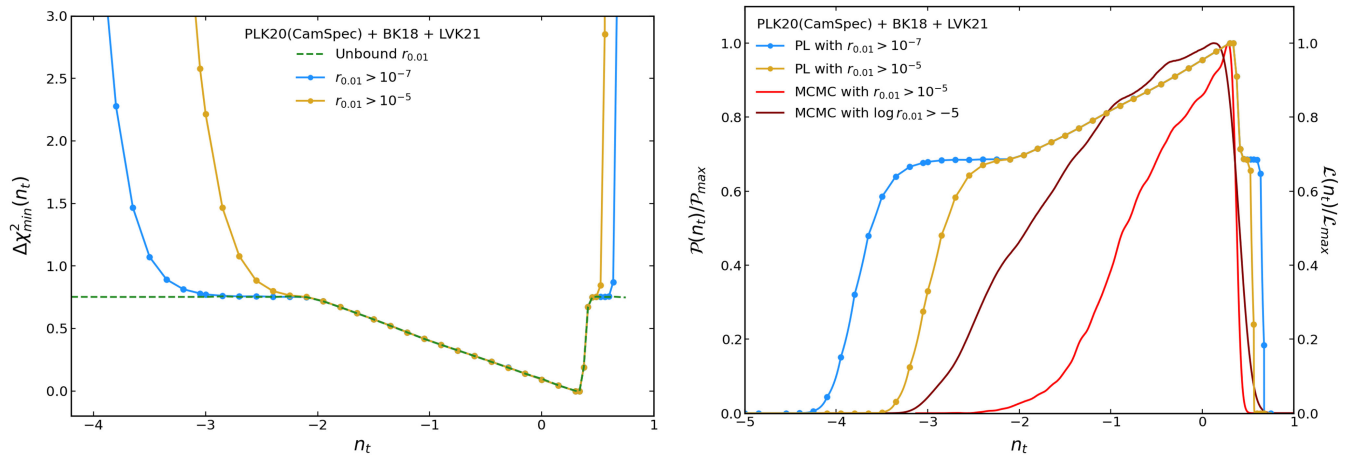


FIG. 6. Left: comparison between $\Delta\chi_{\min}^2$ as a function of n_t for PLK20(CamSpec) + BK18 + LVK21, while imposing $r_{0,01} > 10^{-7}$ (blue) and $r_{0,01} > 10^{-5}$ (gold). Here we also show the corresponding result from Fig. 4 (dashed green) to emphasize the effect of bounding $r_{0,01}$. Right: comparison between $\mathcal{L}/\mathcal{L}_{\max}$ of the cases mentioned above and $\mathcal{P}/\mathcal{P}_{\max}$ from the MCMC (red) of Fig. 1. Here, we plot the posterior obtained assuming a uniform prior on $\log r_{0,01} \in [-5, 0]$.

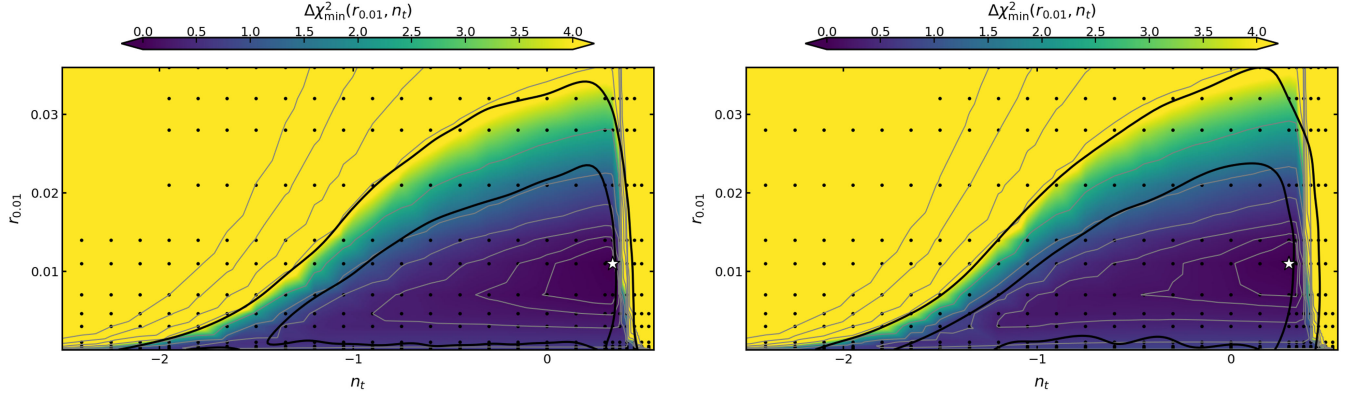


FIG. 7. 2D PL on $r_{0.01}$ and n_t using PLK20(CamSpec) + BK18 + LVK21 (left) and PLK20(HiLLiPoP) + BK18 + LVK21 (right). The star indicates the absolute minimum we find with this procedure, whereas the gray lines show some iso- χ^2 curves to emphasize the 2D shape of the profile. The black dots are the points in which a minimization is performed. We show in solid black the 68% and 95% contours of the 2D marginalized posterior from Sec. IV A.

the rest of the parameters for any fixed combination of $r_{0.01}$ and n_t .

Of course, such an analysis is very demanding in terms of CPU time, given that we have to iterate an already computationally heavy procedure. For this reason, we perform it only on PLK20(CamSpec/HiLLiPoP) + BK18 + LVK21 and not on their counterparts without LVK21. We assume a range of the tensor-to-scalar ratio from a maximum of $r_{0.01} = 0.036$ to a minimum of $r_{0.01} = 10^{-6}$, for a total of $N = 15$ steps. The range in n_t varies case by case, as Fig. 7 shows (black dots). Then, we interpolate these points to get a smooth surface that represents the result for $N \rightarrow \infty$.

We overplot the marginalized 2D posteriors from the Bayesian analysis discussed in Sec. IV A. The posteriors seem to follow the same behavior as the iso- χ^2 curves, at least for $r_{0.01} \gtrsim 5 \times 10^{-4}$. In particular, both 95% posterior contours follow approximately the iso- χ^2 curve at $\Delta\chi_{\min}^2 \simeq 4.6$. Similarly to the discussion at the end of

Sec. IV B 3, this depends on the choice of the prior for the MCMC exploration, thus this correspondence is not guaranteed if the SSA is not employed.

Here, if we imagine intersecting this surface with any horizontal (vertical) plane for a value of $r_{0.01}$ (n_t), we would obtain the PL of n_t ($r_{0.01}$) conditioned on that value. We see that for the values of $r_{0.01}$ considered here, the PL do not show any plateau (see Fig. 4), and instead χ^2 diverges for low and high values of n_t allowing in principle to derive some confidence interval.

To illustrate the results, we fix the tensor-to-scalar ratio to the value predicted by the Starobinsky inflation $r_{0.01} \sim 0.0046$ [61]. The resulting profiles on n_t are shown in Fig. 8 for PLK20(HiLLiPoP) + BK18 + LVK21 (solid blue) and PLK20(CamSpec) + BK18 + LVK21 (solid red). On top of this, we obtain the Bayesian equivalent of this by running an MCMC imposing again $r_{0.01} = 0.0046$. The resulting posterior distributions are shown in the same figure with dashed lines.




FIG. 8. $\Delta\chi_{\min}^2$ as a function of n_t , while imposing $r_{0.01} = 0.0046$ corresponding to Starobinsky inflation. In red we plot the results of PLK20(CamSpec) + BK18 + LVK21, while in blue the ones of PLK20(HiLLiPoP) + BK18 + LVK21.

Note that the frequentist and Bayesian results agree quite well in this case. As in the full 2D case, this correspondence indicates that the other parameters do not induce any volume effects on n_t . This also ensures that the credible and confidence intervals should be the same, allowing us to obtain frequentist confidence intervals.

Of course, we do not know *a priori* the $\Delta\chi^2_{\min}$ corresponding to the 95% CL, since the distribution is clearly non-Gaussian. Still, we can obtain it from the Bayesian intervals by finding the $\Delta\chi^2_{\min}$ that gives the same values. In particular, the Bayesian intervals read

$$\begin{aligned} -1.43 < n_t < 0.41 & \text{ PLK20(CamSpec) + BK18 + LVK21,} \\ -1.47 < n_t < 0.41 & \text{ PLK20(Hillipop) + BK18 + LVK21,} \end{aligned}$$

which correspond, respectively, to $\Delta\chi^2_{\min} = 1.5$ and $\Delta\chi^2_{\min} = 1.3$. As expected, these are lower than the corresponding Gaussian prescription for the 95% CL interval of $\Delta\chi^2_{\min} = 3.84$, since both of our distributions are much flatter than a Gaussian.

As a final remark, HiLLiPoP has a mild preference for $n_t \simeq -1$, also shown in Fig. 4, which is not shown by CamSpec.

V. CONCLUSIONS

In this work, we explore the primordial tensor perturbation constraints utilizing the latest available data from CMB and gravitational wave measurements. We focus more specifically on the tensor-to-scalar ratio $r_{0.01}$ and the tensor spectral tilt n_t parametrizing the amplitude and the tilt of the tensor perturbations power spectrum. Our dataset includes CMB measurements from the latest *Planck* PR4 dataset, the BICEP/Keck array, and LIGO-Virgo-KAGRA upper bound on the energy density of GWs to constraint the tensor spectrum at much smaller scales.

In order to provide a complete picture of the constraints, we use both Bayesian and frequentist approaches. While the former provides with a statistical statement about the probability distribution of the different parameters after marginalization, the latter can describe the probability distribution of data given the theoretical model insensitive to the choice of priors and free from volume effects. Thus, it is clear that comparing credible and confidence intervals is not trivial; however, it allows us to extract useful information on relatively unconstrained parameters. The results are summarized in Fig. 9.

For the tensor-to-scalar ratio, we find upper limits with profile likelihoods slightly more conservative than the Bayesian ones. This indicates that both volume effects and prior choices have a role in obtaining the credible intervals.

Concerning the tensor spectral tilt n_t , the recovered profile likelihoods exhibit a very non-Gaussian behavior with the absence of distinct bounds. In fact, $\Delta\chi^2_{\min}$ never

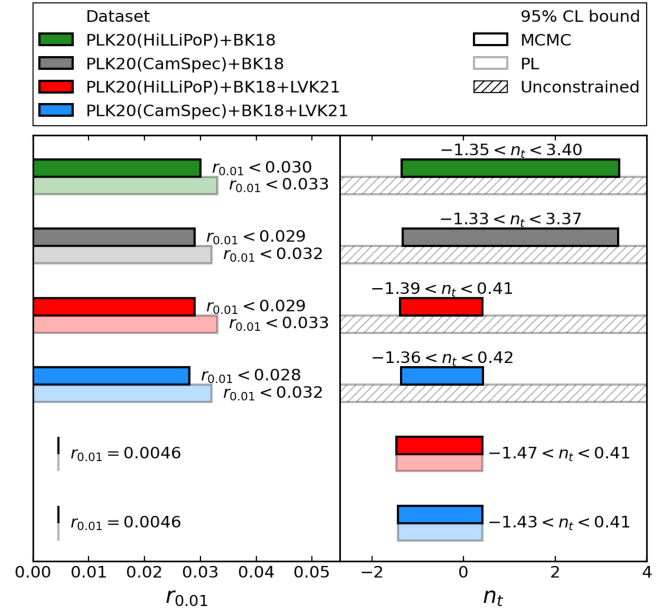


FIG. 9. Summary of the 95% CL intervals obtained in this work. The solid patches refer to the credible intervals, while the shaded ones to the confidence intervals. The last two rows refer to $r_{0.01} = 0.0046$ as predicted by Starobinsky inflation.

exceeds $\Delta\chi^2_{\min} \simeq 1$, which cannot be captured by Bayesian analysis. We identify that this behavior is the result of the potential accommodation of highly extreme spectral tilts at the expense of driving $r_{0.01}$ toward zero. Despite this, datasets accounting for GW interferometers show a mild preference for $n_t \simeq 0.3$. Instead, those without LIGO-Virgo-KAGRA prefer $n_t \simeq 2$. Some constraints on n_t can be found by introducing a cutoff on small values of $r_{0.01}$ as done in Bayesian analysis ($r_{0.01} > 10^{-5}$). However, it cannot solve the discrepancy with the marginalized posterior of the MCMC analysis when it is driven by a flat prior on $r_{0.01}$. For example, sampling with a log-uniform prior allows us to explore the region for very low values of $r_{0.01}$ and provide a marginalized posterior significantly wider. Indeed, in the case of n_t , given today's accuracy of the data, Bayesian analysis is dominated by the choice of priors and consequently credibility intervals are driven by this choice. In contrast, frequentist analysis, which is independent of any prior, leads to unconstrained n_t .

When profiling $r_{0.01}$ and n_t simultaneously in 2D, we recover constraints close to the 2D marginalized posterior for large values of r ($r_{0.01} \gtrsim 5 \times 10^{-4}$). This indicates that other parameters (physical or nuisance) do not induce volume effects in this 2D plane, proving that in this range the tensor sector is sufficiently decoupled from the Λ CDM parameters. For the largest value of $r_{0.01}$, including the prediction for the Starobinsky model ($r_{0.01} = 0.0046$), constraints on n_t can be derived (see Fig. 9). In this case, the correspondence between Bayesian and frequentist results is evident, ensuring that credible and confidence intervals are the same.

Projecting this comparison into the future, when more data will be available, we can expect a similar difference between these two methods as they answer different questions, especially in the case of an upper limit on r . Indeed, the prior dependence of the Bayesian analysis is an unavoidable part of the framework, so one must deal with the fact that a prior must be taken before performing any analysis. Instead, performing the same analysis with different priors is a way to showcase the consequence of such a choice [17]. On the other hand, the more the new data will be constraining on the tensor sector, the more frequentist and Bayesian results will converge to each other. This is because the likelihood function will dominate both methods in the limit of a full detection of r .

To summarize, this paper delivers a comprehensive and statistically robust analysis of the tensor sector of parameter space. Beyond offering an updated perspective using the latest datasets currently available, the analysis underscores the significance of probing relatively unconstrained parameters with frequentist approaches, complementing the widely used Bayesian methods. This insight proves valuable for forthcoming investigations into CMB polarization, such as with LiteBIRD, and for any endeavor exploring extensions to the standard Λ CDM model.

ACKNOWLEDGMENTS

The authors thank Gilles Weymann-Despres and Stéphane Ilić for useful comments and discussions. This work is based on observations obtained with *Planck*,⁴ an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada. The *Planck* PR4 data are publicly available on the *Planck* Legacy Archive [35]. Both likelihoods LoLLiPoP and HiLLiPoP based on PR4 are publicly available on GITHUB⁵ as external likelihoods for COBAYA. Some of the results in this paper have been derived using the following packages: CAMB [28,42,43], COBAYA [62], GETDIST [63], MATPLOTLIB⁶ [64], and NUMPY⁷ [65]. G. G. acknowledges support from ASI/LiteBIRD Grant No. 2020-9-HH.0, and by the InDark and LiteBIRD INFN projects.

APPENDIX A: CONSTRAINTS ON THE Λ CDM PARAMETERS

As mentioned in Sec. III, the focus of this work is the tensor sector of parameter space, which consists of $r_{0.01}$ and n_t . Despite this, to correctly capture the variability in these parameters, it is also important to sample all other Λ CDM and nuisance parameters. First, we define a set of priors in Table II, encoding our knowledge about physical

TABLE II. Priors on the parameters of Λ CDM + $r_{0.01}$ + n_t . Here, A_s is the scalar perturbations amplitude, n_s the scalar spectral tilt, Ω_b and Ω_{cdm} are the abundances of baryons and CDM, $h \equiv H_0/100$ is the Hubble constant divided by 100, τ_{reio} the optical depth, and θ_{MC} is an approximate quantity representing the sound horizon.

Parameter	Prior	Parameter	Prior
$\log(10^{10}A_s)$	[1.61, 3.91]	τ_{reio}	[0.01, 0.8]
n_s	[0.8, 1.2]	θ_{MC}	[0.5, 10]
$\Omega_b h^2$	[0.005, 0.1]	$r_{0.01}$	$[10^{-5}, 3]$
$\Omega_{\text{cdm}} h^2$	[0.001, 0.99]	n_t	$[-5, 5]$

parameters. For what regards the nuisance parameters of the various likelihoods, we stick to the default settings.

In addition, we fix the number of relativistic degrees of freedom to $N_{\text{eff}} = 3.046$, asking for a massive neutrino with $M_\nu = 0.06$ eV.

In Sec. IVA, the results on the tensor sectors are discussed, while here we show those on the Λ CDM parameters. In fact, Fig. 10 shows the triangle plot of what we may call the “scalar” sector of parameters. Note that here we substitute θ_{MC} with H_0 as it is a more physical parameter controlling the sound horizon. Here, we consider PLK(CamSpec) + BK18 + LVK21, PLK(CamSpec) + BK18, PLK(HiLLiPoP) + BK18 + LVK21, and PLK × (HiLLiPoP) + BK18 showing that they are all compatible with each other.

To be more quantitative, we also report in Table III the mean values and standard deviations of these parameters.

As expected, removing LVK21 does not produce any difference in Λ CDM parameters. The only noticeable difference is found by looking at n_s for PLK20 (CamSpec) and PLK20(HiLLiPoP), which is estimated to be slightly higher (approximately $0.5\sigma_{n_s}$ higher) by the latter. This feature is also remarked in [15]. For completeness, we also mention that HiLLiPoP also estimates H_0 to be higher of $\sim 0.2\sigma_{H_0}$.

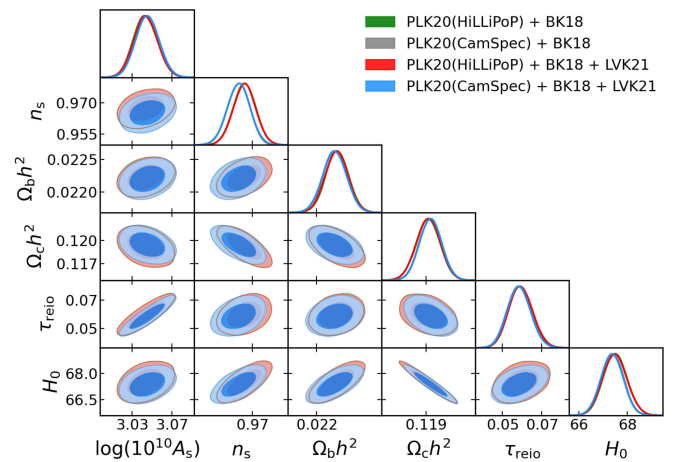


FIG. 10. Λ CDM contours.

⁴<http://www.esa.int/Planck>.

⁵<http://github.com/planck-npipe>.

⁶<https://github.com/matplotlib/matplotlib>.

⁷<https://github.com/numpy/numpy>.

TABLE III. Summary of the constraints on the Λ CDM + $r_{0.01}$ + n_t model. For the six Λ CDM parameters, we report the mean and the standard deviation; instead, we report the 95% CL intervals for the tensor sector, as in Sec. IV A.

Parameter	PLK20(CamSpec)+ BK18 + LVK21	PLK20(HiLLiPoP)+ BK18 + LVK21	PLK20(CamSpec)+ BK18	PLK20(HiLLiPoP)+ BK18
$\log(10^{10}A_s)$	3.046 ± 0.012	3.044 ± 0.012	3.045 ± 0.012	3.044 ± 0.012
n_s	0.9653 ± 0.0039	0.9673 ± 0.0039	0.9653 ± 0.0038	0.9673 ± 0.0039
$\Omega_b h^2$	0.02220 ± 0.00013	0.02222 ± 0.00013	0.02220 ± 0.00013	0.02222 ± 0.00013
$\Omega_{\text{cdm}} h^2$	0.1194 ± 0.0010	0.1192 ± 0.0011	0.1194 ± 0.0010	0.1192 ± 0.0011
τ_{reio}	0.0586 ± 0.0059	0.0594 ± 0.0060	0.0586 ± 0.0059	0.0594 ± 0.0060
H_0	67.37 ± 0.46	67.49 ± 0.50	67.40 ± 0.45	67.50 ± 0.49
$r_{0.01}$	<0.028	<0.029	<0.029	<0.030
n_t	$-1.36 < n_t < 0.42$	$-1.39 < n_t < 0.41$	$-1.33 < n_t < 3.37$	$-1.35 < n_t < 3.40$

APPENDIX B: TESTS ON THE ACCURACY

In Sec. III we mention that, to perform a precise PL, it is crucial to efficiently maximize the likelihood, reaching at each point of the profiled parameter the absolute minimum of the χ^2 (and not a local one). Thus, to verify the precision of our minimizing procedure, we performed different tests. For example, we fix the tensor-to-scalar ratio to $r = 0.02$ and minimize the χ^2 8 times, storing the result of each. We repeat this while changing the accuracy parameters of CAMB and of MINUIT. The dispersion of the results is an indication of the precision of minimization performed by MINUIT, since it gauges the reliability of recovering the absolute minimum. In addition, the value of these points tends to depend to some degree on the accuracy of the Boltzmann solver [36].

In particular, the parameters we considered for these tests are as follows:

- (i) `stra` of MINUIT, which allows to change the strategy of the minimization from *fast* (`stra = 0`), *balanced* (`stra = 1`), and *accurate* (`stra = 2`). The default value is `stra = 1`.
- (ii) `AccuracyBoost` of CAMB, which controls several other accuracy parameters of the Boltzmann solver. The default value is 1.
- (iii) `lAccuracyBoost` of CAMB, which is related to the resolution in ℓ space of the Boltzmann solver. The default value is 1.

Figure 11 shows the different sets of minima we obtain, where we define

- Standard Accuracy*: Default settings;
- High Accuracy iMinuit*: `stra = 2`;
- High Accuracy CAMB*: `AccuracyBoost = 2` and `lAccuracyBoost = 2`;
- All High Accuracy*: `stra = 2`, `AccuracyBoost = 2` and `lAccuracyBoost = 2`;
- Super High Accuracy CAMB*: `AccuracyBoost = 3` and `lAccuracyBoost = 3`;
- Super High CAMB + High iMinuit*: `stra = 2`, `AccuracyBoost = 3` and `lAccuracyBoost = 3`.

It is clear that these settings have an effect on the values we obtain for the minimum. Figure 11 shows that when we increase the accuracy of MINUIT the dispersion of the points gets severely reduced. Instead, the absolute minimum we obtain here is given by increasing the accuracy of CAMB to *high*, since going to *super high* actually increases the values of χ^2 . Despite this, all of these nondefault settings increase significantly the computation time necessary to maximize the likelihood; thus, we must ask ourselves if it is worth imposing some extra accuracy.

First, note that with *high* CAMB accuracy, the typical χ^2 decreases of ~ 0.125 . Still, if this difference is just a constant offset between the two configurations, it will essentially disappear when we normalize the likelihood values on the absolute minimum. Furthermore, note that although increasing the accuracy of *MINUIT* leads all points to converge to the absolute minimum, it is not worth it in terms of computation time. In fact, usually more than half of the points with the default MINUIT converge to the same point, so statistically we always get at least one point there.




FIG. 11. Different set of minima obtained with different settings for the accuracy of CAMB and MINUIT.

Despite these considerations, we build a parabolic fit for each accuracy configuration [with $\mathcal{O}(10)$ points] to verify that we can stick to the default settings. We did not find significant differences in the final PL.

APPENDIX C: NON-GAUSSIAN FELDMAN-COUSINS

In Sec. IV B 1 we discuss the fact that our PLs on the tensor-to-scalar ratio present some non-Gaussian features. Thus, we ask ourselves whether it is possible to generalize the FC prescription to a more complex curve that fits our PLs.

We consider two simple alternatives that have two characteristics: we require that one of the parameters of the function corresponds to our parameter of interest; in addition, we will consider only this parameter as free to vary in the FC procedure. In other words, the curve we choose will shift to the right or to the left with the parameter changing, while keeping its shape unchanged.

In particular, we consider what we call *double Gaussian*, defined as

$$P(x|\mu, \Delta\mu, \sigma_1, \sigma_2, A) \propto \exp\left[-\frac{(x-\mu)^2}{2\sigma_1^2}\right] + A \exp\left[-\frac{(x-\mu-\Delta\mu)^2}{2\sigma_2^2}\right], \quad (\text{C1})$$

and *piecewise Gaussian*, which reads

$$P(x|\mu, \Delta\mu, \sigma_1, \sigma_2, x_{\text{sep}}) \propto \begin{cases} \exp\left[-\frac{(x-\mu)^2}{2\sigma_1^2}\right] & \text{for } x \leq x_{\text{sep}} \\ \exp\left[-\frac{(x-\mu-\Delta\mu)^2}{2\sigma_2^2}\right] & \text{for } x > x_{\text{sep}} \end{cases} \quad (\text{C2})$$

Of course, these curves do not find any physical justification, but are born from the pragmatic attempt to fit our PL with analytical formulas. The parameter values are estimated by fitting Eqs. (C1) and (C2) to our PL. These curves, together with the standard Gaussian, allow us to fit the features mentioned in Sec. IV B 1.

As an example, the left panel of Fig. 12 shows the results of PLK20(HiLLiPoP) + BK18 + LVK21, which is our most complex PL in terms of extra features.

- (i) Using the standard Gaussian with all the PL points, we clearly get a poor fit of the maximum likelihood because of the relatively high-probability tail. This can be solved by fitting just the points near the minimum ($\Delta\chi^2 \lesssim 2$), as we do in Sec. IV B 1.
- (ii) If instead we try to use the double Gaussian on the full range of PL points, we can recover quite well both the maximum likelihood and the high-probability tail. On the other hand, using it in the region where $\Delta\chi^2 \lesssim 2$ allows us to fit the maximum and the bump at low $r_{0.01}$.
- (iii) Similarly to the double Gaussian, using the piecewise Gaussian on the full range of points allows us to fit the maximum likelihood and the tail.

At this point, we perform the FC computation for all the cases mentioned above. The right panel of Fig. 12 shows the intersections between the various FC belts (solid) and their relative best-fit values for $r_{0.01}$ (dashed). We find that the upper bound of $r_{0.01}$ does not change significantly from case to case, allowing the choice made in Sec. IV B 1. Of course, the only relevant exception is the Gaussian case on the full range, which clearly overestimates the upper bound.

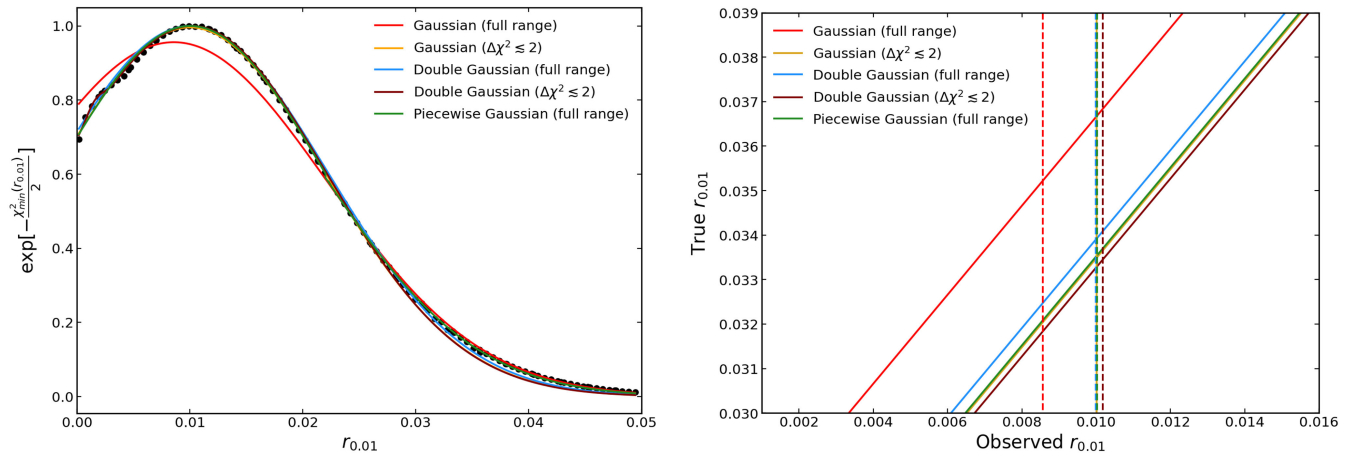


FIG. 12. Left: fits of different curves on the PLK20(HiLLiPoP) + BK18 + LVK21 PL. Right: intersection between the FC belts (solid) and their relative best-fit values of $r_{0.01}$ (dashed), defining its upper bound with different fitting curves.

- [1] U. Seljak and M. Zaldarriaga, *Phys. Rev. Lett.* **78**, 2054 (1997).
- [2] M. Kamionkowski, A. Kosowsky, and A. Stebbins, *Phys. Rev. D* **55**, 7368 (1997).
- [3] M. Kamionkowski and E. D. Kovetz, *Annu. Rev. Astron. Astrophys.* **54**, 227 (2016).
- [4] M. C. Guzzetti, N. Bartolo, M. Liguori, and S. Matarrese, *Riv. Nuovo Cimento* **39**, 399 (2016).
- [5] U. Seljak, *Astrophys. J.* **482**, 6 (1997).
- [6] W. Hu and M. J. White, *New Astron.* **2**, 323 (1997).
- [7] P. A. R. Ade, R. W. Aikin, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, J. A. Brevik, I. Buder, E. Bullock *et al.* (BICEP2 Collaboration), *Phys. Rev. Lett.* **112**, 241101 (2014).
- [8] P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton *et al.* (Simons Observatory Collaboration), *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [9] K. N. Abazajian, P. Adshead, Z. Ahmed, S. W. Allen, D. Alonso, K. S. Arnold, C. Baccigalupi, J. G. Bartlett, N. Battaglia, B. A. Benson *et al.*, [arXiv:1610.02743](https://arxiv.org/abs/1610.02743).
- [10] M. Hazumi *et al.*, *J. Low Temp. Phys.* **194**, 443 (2019).
- [11] E. Allys, K. Arnold, J. Aumont, R. Aurlien, S. Azzoni, C. Baccigalupi, A. J. Banday, R. Banerji, R. B. Barreiro *et al.* (LiteBIRD Collaboration), *Prog. Theor. Exp. Phys.* **2023**, 042F01 (2023).
- [12] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020).
- [13] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A5 (2020).
- [14] E. Rosenberg, S. Gratton, and G. Efstathiou, *Mon. Not. R. Astron. Soc.* **517**, 4620 (2022).
- [15] M. Tristram, A. J. Banday, M. Douspis, X. Garrido, K. M. Górski, S. Henrot-Versillé, L. T. Hergt, S. Ilić, R. Keskitalo, G. Lagache *et al.*, *Astron. Astrophys.* **682**, A37 (2024).
- [16] M. Tristram, A. J. Banday, K. M. Górski, R. Keskitalo, C. R. Lawrence, K. J. Andersen, R. B. Barreiro, J. Borrill, L. P. L. Colombo, H. K. Eriksen *et al.*, *Phys. Rev. D* **105**, 083524 (2022).
- [17] G. Galloni, N. Bartolo, S. Matarrese, M. Migliaccio, A. Ricciardone, and N. Vittorio, *J. Cosmol. Astropart. Phys.* **04** (2023) 062.
- [18] Y. Akrami, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A10 (2020).
- [19] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
- [20] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin *et al.*, *Classical Quantum Gravity* **32**, 024001 (2014).
- [21] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams *et al.* (LIGO Scientific, Virgo, KAGRA Collaborations), *Living Rev. Relativity* **23**, 3 (2020).
- [22] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **100**, 061101 (2019).
- [23] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams *et al.* (LIGO Scientific, Virgo, KAGRA Collaborations), *Phys. Rev. D* **104**, 022004 (2021).
- [24] N. Metropolis and S. Ulam, *J. Am. Stat. Assoc.* **44**, 335 (1949).
- [25] R. Eckhardt, *Los Alamos Sci.* **15**, 131 (1987), <https://www.bibsonomy.org/bibtex/20a9c00e0909f45c5d6567b6c49bed7c3/rspreeuw>.
- [26] D. P. Kroese, T. Brereton, T. Taimre, and Z. I. Botev, *WIREs Comput. Stat.* **6**, 386 (2014).
- [27] A. Lewis and S. Bridle, *Phys. Rev. D* **66**, 103511 (2002).
- [28] A. Lewis, *Phys. Rev. D* **87**, 103529 (2013).
- [29] G. Efstathiou, E. Rosenberg, and V. Poulin, *Phys. Rev. Lett.* **132**, 221002 (2024).
- [30] S. Henrot-Versillé, F. Robinet, N. Leroy, S. Plaszczynski, N. Arnaud, M.-A. Bizouard, F. Cavalier, N. Christensen, F. Couchot, S. Franco *et al.*, *Classical Quantum Gravity* **32**, 045003 (2015).
- [31] S. Henrot-Versillé, O. Perdureau, S. Plaszczynski, B. R. d'Orfeuill, M. Spinelli, and M. Tristram, [arXiv:1607.02964](https://arxiv.org/abs/1607.02964).
- [32] L. Herold, E. G. M. Ferreira, and E. Komatsu, *Astrophys. J. Lett.* **929**, L16 (2022).
- [33] R. Murgia, G. F. Abellán, and V. Poulin, *Phys. Rev. D* **103**, 063502 (2021).
- [34] E. McDonough, J. C. Hill, M. M. Ivanov, A. La Posta, and M. W. Toomey, [arXiv:2310.19899](https://arxiv.org/abs/2310.19899).
- [35] <http://pla.esac.esa.int>.
- [36] Planck Collaboration, Y. Akrami, K. J. Andersen, M. Ashdown, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak *et al.*, *Astron. Astrophys.* **643**, A42 (2020).
- [37] H. K. Eriksen, J. B. Jewell, C. Dickinson, A. J. Banday, K. M. Górski, and C. R. Lawrence, *Astrophys. J.* **676**, 10 (2008).
- [38] P. A. R. Ade, Z. Ahmed, M. Amiri, D. Barkats, R. B. Thakur, C. A. Bischoff, D. Beck, J. J. Bock, H. Boenish *et al.* (BICEP/Keck Collaboration), *Phys. Rev. Lett.* **127**, 151301 (2021).
- [39] S. Hamimeche and A. Lewis, *Phys. Rev. D* **77**, 103013 (2008).
- [40] G. Cabass, L. Pagano, L. Salvati, M. Gerbino, E. Giusarma, and A. Melchiorri, *Phys. Rev. D* **93**, 063508 (2016).
- [41] J. Carron, M. Mirmelstein, and A. Lewis, *J. Cosmol. Astropart. Phys.* **09** (2022) 039.
- [42] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000).
- [43] C. Howlett, A. Lewis, A. Hall, and A. Challinor, *J. Cosmol. Astropart. Phys.* **04** (2012) 027.
- [44] F. Couchot, S. Henrot-Versillé, O. Perdureau, S. Plaszczynski, B. Rouillé d'Orfeuill, M. Spinelli, and M. Tristram, *Astron. Astrophys.* **606**, A104 (2017).
- [45] P. Campeti and E. Komatsu, *Astrophys. J.* **941**, 110 (2022).

- [46] V. Acquaviva, N. Bartolo, S. Matarrese, and A. Riotto, *Nucl. Phys.* **B667**, 119 (2003).
- [47] S. Matarrese, S. Mollerach, and M. Bruni, *Phys. Rev. D* **58**, 043504 (1998).
- [48] D. Baumann, P. Steinhardt, K. Takahashi, and K. Ichiki, *Phys. Rev. D* **76**, 084019 (2007).
- [49] K. N. Ananda, C. Clarkson, and D. Wands, *Phys. Rev. D* **75**, 123518 (2007).
- [50] C. Fidler, G. W. Pettinari, M. Beneke, R. Crittenden, K. Koyama, and D. Wands, *J. Cosmol. Astropart. Phys.* **07** (2014) 011.
- [51] L. T. Hergt, W. J. Handley, M. P. Hobson, and A. N. Lasenby, *Phys. Rev. D* **103**, 123511 (2021).
- [52] R. L. Workman, V. D. Burkert, V. Crede, E. Klempt, U. Thoma, L. Tiator, K. Agashe, G. Aielli, B. C. Allanach *et al.* (Particle Data Group), *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).
- [53] G. J. Feldman and R. D. Cousins, *Phys. Rev. D* **57**, 3873 (1998).
- [54] J. Torrado and A. Lewis, [arXiv:2005.05290](https://arxiv.org/abs/2005.05290).
- [55] F. James and M. Roos, *Comput. Phys. Commun.* **10**, 343 (1975).
- [56] C. Cartis, J. Fiala, B. Marteau, and L. Roberts, [arXiv:1804.00154](https://arxiv.org/abs/1804.00154).
- [57] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, *Nat. Methods* **17**, 261 (2020).
- [58] A. Nygaard, E. B. Holm, S. Hannestad, and T. Tram, *J. Cosmol. Astropart. Phys.* **11** (2023) 064.
- [59] E. B. Holm, A. Nygaard, J. Dakin, S. Hannestad, and T. Tram, [arXiv:2312.02972](https://arxiv.org/abs/2312.02972).
- [60] T. Karwal, Y. Patel, A. Bartlett, V. Poulin, T. L. Smith, and D. N. Pfeffer, [arXiv:2401.14225](https://arxiv.org/abs/2401.14225).
- [61] A. Starobinsky, *Phys. Lett.* **91B**, 99 (1980).
- [62] J. Torrado and A. Lewis, *J. Cosmol. Astropart. Phys.* **05** (2021) 057.
- [63] A. Lewis, [arXiv:1910.13970](https://arxiv.org/abs/1910.13970).
- [64] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [65] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith *et al.*, *Nature (London)* **585**, 357 (2020).