

PAPER • OPEN ACCESS

Accurate and efficient likelihood modeling for large-scale CMB data







To cite this article: Giacomo Galloni *et al* JCAP12(2025)052

View the [article online](#) for updates and enhancements.

You may also like

- [First release of *LiteBIRD* simulations from an end-to-end pipeline](#)
M. Bortolami, N. Raffuzzi, L. Pagano et al.
- [Particle identification capability of a homogeneous calorimeter composed of oriented crystals](#)
Pietro Monti-Guarnieri, Laura Bandiera, Nicola Canale et al.
- [Investigation of the burst phenomenon in SiPMs in liquid nitrogen](#)
M. Guarise, M. Andreotti, A. Balboni et al.

Accurate and efficient likelihood modeling for large-scale CMB data

Giacomo Galloni ^{a,b,*} Paolo Campeti ^{b,c} Luca Pagano ^{a,b,d} Martina Gerbino ^b
Massimiliano Lattanzi ^b and Paolo Natoli ^{a,b}

^a*Dipartimento di Fisica e Scienze della Terra, Università degli Studi di Ferrara,
via Saragat 1, I-44122 Ferrara, Italy*

^b*Istituto Nazionale di Fisica Nucleare, Sezione di Ferrara,
via Saragat 1, I-44122 Ferrara, Italy*

^c*ICSC, Centro Nazionale “High Performance Computing, Big Data and Quantum Computing”,
Casalecchio di Reno, Bologna, Italy*

^d*Institut d’Astrophysique Spatiale, CNRS, Université Paris-Sud, Université Paris-Saclay,
Bât. 121, 91405 Orsay cedex, France*

E-mail: giacomo.galloni@unife.it, paolo.campeti@fe.infn.it,
pgnlcu@unife.it, gerbinom@fe.infn.it, lattanzi@fe.infn.it,
ntlpla@unife.it

ABSTRACT: Accurate parameter estimation from cosmic microwave background data requires reliable likelihood modeling, particularly at large angular scales where angular power spectrum estimators exhibit non-Gaussian statistics. We present a novel approach, based on the Hamimeche-Lewis formalism, that marginalizes over auto-spectra, thus reducing residual biases from noise misestimation and partial sky coverage. We validate our approach by simulating three independent CMB channels, or data splits, in a multi-field setting, comparing to the pixel-based likelihood ground truth estimates for the optical depth τ and the tensor-to-scalar ratio r . We benchmark our method against the main power spectrum based alternatives available in the literature, showing that it outperforms all of them in terms of accuracy, while remaining fast and computationally efficient <https://github.com/ggalloni/CHARM-Like>.

KEYWORDS: Bayesian reasoning, CMBR polarisation, CMBR theory, cosmological parameters from CMBR

ARXIV EPRINT: [2505.24829](https://arxiv.org/abs/2505.24829)

*Corresponding author.

Contents

1	Introduction	1
2	Modeling large-scale CMB data	3
2.1	The Hamimeche & Lewis likelihood	3
2.2	The new auto-spectra marginalized Hamimeche & Lewis	6
2.3	Pixel-based likelihood	7
3	Methodology	8
3.1	Our setup	8
3.2	Simulations	9
3.3	Posterior evaluation	10
4	Results	11
4.1	Full-sky	11
4.2	Cut-sky	12
4.3	Using only two channels	15
5	Conclusions	16
A	Assuming a wrong fiducial model	19
B	The role of the offset	20
B.1	Changing offset	22
C	Approximating the likelihood with a Gaussian	22

1 Introduction

The Cosmic Microwave Background (CMB) is a cornerstone of modern cosmology, offering deep insights into the early Universe. Its anisotropies are thought to originate from quantum fluctuations during inflation, which seeded small inhomogeneities in the primordial plasma [1–3].

As CMB photons travel through the Universe, they acquire additional imprints from various processes [4–9]. In particular, polarization — typically decomposed into curl-free “E-mode” and divergence-free “B-mode” components — serves as a key observable. While scalar perturbations only generate E-modes, inflation-driven tensor perturbations would also produce a primordial B-mode signal [4, 10, 11] that remains unobserved [12–14].

A wide range of experiments has been deployed to study the CMB, or will be in the near future. Full-sky observations, unique to satellites [15–20], are especially important when studying large-scale signals, such as the potential imprint of primordial gravitational waves (GWs), or the reionization epoch [13, 21–26].

In a full-sky scenario, assuming that the fields (pure CMB and noise) observed are Gaussian and statistically isotropic, the likelihood of the observed data given a theoretical

model is exactly described by a Wishart distribution of the angular power spectra, the C_ℓ [27, 28]. However, even satellites which can access the full sky must contend with bright foreground signals that prevent direct complete CMB imaging (see e.g. [20, 29, 30]). Multi-frequency observations allow for foreground separation, but this process is not perfect [31–38]. Thus, it is reasonable to assume a partial sky coverage — a cut-sky — for future CMB experiments.

This creates a statistical challenge in how to model the probability of the data at large angular scales. In principle, one could work at the map level, assuming that pixel values follow a multivariate Gaussian distribution [28]. Various examples of this approach can be found in the literature, e.g. the analysis of COBE satellite [39, 40], where multiple frequency channels were combined at map-level, filtering non-stochastic real-space components [41–44].

However, this approach quickly becomes computationally expensive and is highly sensitive to systematics, as it depends on the ability to estimate potentially large pixel covariance matrices [27]. Other approaches to overcome this issue are based on sampling the full joint CMB posterior, e.g. through Gibbs sampling [45–47]. Instead, since the angular power spectrum provides a sufficient summary statistic of cosmological information in the map, harmonic-space likelihoods are commonly used [48, 49].

One such approach in the cut-sky is the Hamimeche & Lewis (HL) likelihood, which accounts for the non-Gaussianity of C_ℓ in full-sky and for couplings between multipoles introduced by masking [27]. Another interesting strength of this approximation is its native capability of working in a “multi-field” scenario. In fact, it can be applied in joint analysis across correlated fields, which could be CMB temperature and polarization, or equivalently maps observed at different frequencies, or data splits from a mission dataset [27]. Despite being a good approximation in the cut-sky regime and widely used in the literature, the HL likelihood does not fully capture the distortion of the C_ℓ distribution, which deviates from a Wishart form when masking is applied. In addition, semi-analytical cross-spectrum-based likelihood approximation schemes were explored in [26, 50] providing interesting alternatives.

In this work, we revisit the HL likelihood in the context of multi-field CMB polarization analyses, focusing on its practical application in realistic settings. In particular, we address two common challenges: partial sky coverage and imperfect knowledge of the noise bias [51, 52], both of which affect the performance of this approximation. Indeed, imperfect knowledge of the systematics of an experiment, or of the residual foreground signal mentioned above, could lead to a mismatch between the assumed noise bias and the true one. To this end, we study two variants of the HL prescription. The first is what we call cross-spectra HL (cHL), which tries to extend the offset HL (oHL) already adopted in the literature [12, 53, 54] to the multi-field case [55, 56]. The cHL is designed to use only cross-spectra, i.e. cross-correlations between two different fields, to by-pass the misestimation of the noise bias. The second, which we introduce here for the first time, is a new *marginalized HL* (mHL) approach, in which auto-spectra are marginalized over analytically to effectively remove dependence on uncertain noise levels.

We carry out a systematic comparison of the HL, cHL, and our proposed mHL using synthetic maps, both in full-sky and cut-sky regimes. As noted above, we focus on large angular scales with $\ell \leq 32$, targeting the reionization peak of the E- and B-mode spectra. Although the likelihood on smaller angular scales is not perfectly Gaussian — hence the

original motivation for the HL approach [27], which is validated on such scales — the central limit theorem ensures that the distribution becomes increasingly Gaussian as higher multipoles are considered. In this regime, residual foregrounds or systematic effects typically represent a more significant limitation [20, 29, 30]. To explore the limits of the application of each method, we simulate mismatches in the noise bias and assess their impact on parameter inference. To our knowledge, this is the first comprehensive study that jointly investigates the behavior of HL-based likelihoods under simultaneous noise mismatches and sky cuts, and the first to introduce and validate the mHL as a robust alternative. These results provide new insights into the practical use of HL approximations and offer concrete recommendations for future CMB polarization analyses. For example, we study whether a single-field approach is preferable to multi-field HL-based likelihoods.

2 Modeling large-scale CMB data

2.1 The Hamimeche & Lewis likelihood

On the full sky, the exact likelihood of the angular power spectra of two correlated CMB fields, for example temperature T and E-modes, is [28]

$$-2 \ln \mathcal{L}(\hat{\mathbf{C}}_\ell | \mathbf{C}_\ell) = \sum_\ell (2\ell + 1) \left(\text{Tr} \{ \hat{\mathbf{C}}_\ell \mathbf{C}_\ell^{-1} \} - \ln | \hat{\mathbf{C}}_\ell \mathbf{C}_\ell^{-1} | - n \right), \quad (2.1)$$

where

$$\mathbf{C}_\ell = \begin{pmatrix} C_\ell^{TT} & C_\ell^{TE} \\ C_\ell^{TE} & C_\ell^{EE} \end{pmatrix}, \quad (2.2)$$

C_ℓ indicates the true power spectra, the hat indicates their estimates, and $n = 2$ is the number of fields used to normalize the likelihood.

From this starting point, we can define the HL likelihood. Equation (2.1) can be rewritten as [27]

$$-2 \ln \mathcal{L}(\hat{\mathbf{C}}_\ell | \mathbf{C}_\ell) = \sum_{\ell\ell'} X_\ell^T M_{\ell\ell'}^{-1} \Big|_{\text{fid.}} X_\ell, \quad (2.3)$$

where $M_{\ell\ell'}^{-1}$ is the inverse covariance matrix of the $n(n+1)/2$ independent C_ℓ components, evaluated for a fiducial model. Then,

$$X_\ell = \left(X_\ell^{\text{T}\times\text{T}}, X_\ell^{\text{T}\times\text{E}}, X_\ell^{\text{E}\times\text{E}} \right) \quad (2.4)$$

is the vector formed from the independent elements of the matrix $\mathbf{C}_{g\ell} = \mathbf{C}_{f\ell}^{1/2} \mathbf{U}_\ell \mathbf{g}(\mathbf{D}_\ell) \mathbf{U}_\ell^T \mathbf{C}_{f\ell}^{1/2}$. Here, $\mathbf{C}_{f\ell}$ is the fiducial version of the matrix in equation (2.2), \mathbf{U}_ℓ is the diagonalization matrix of $\mathbf{C}_\ell^{-1/2} \hat{\mathbf{C}}_\ell \mathbf{C}_\ell^{-1/2}$, and \mathbf{D}_ℓ is the diagonal matrix of eigenvalues. The function acting on it is defined as [27]

$$\mathbf{g}(\mathbf{D}_\ell) = [\mathbf{g}(\mathbf{D}_\ell)]_{ij} = g(D_{\ell,ii}) \delta_{ij} = \text{sign}(D_{\ell,ii} - 1) \sqrt{2(D_{\ell,ii} - \ln D_{\ell,ii} - 1)} \delta_{ij}. \quad (2.5)$$

Up to now, introducing the HL likelihood, we used the familiar example of CMB temperature and E-mode polarization [28]. However, this formalism is defined in general for n Gaussian and statistically isotropic correlated fields [27]. For example, we could consider three generic fields $\{A, B, C\}$, which in principle could be different frequency channels, data-splits, or actual CMB observables such as temperature, E-modes, and B-modes. The only difference between these cases is the degree to which these fields are correlated.

In practice, CMB analyses are often done in such a multi-field framework, either because one wants to include multiple observables, or because one would benefit from the fact that noise is often uncorrelated between splits. An example is the likelihood constructed for the BICEP/Keck analysis [22], where eleven maps are combined in a single HL likelihood.

In this work, we consider only B-(E-)modes from n_{ch} “channels” to estimate the tensor-to-scalar ratio r (optical depth τ), corresponding to n_{ch} frequency bands observing the same CMB. Each channel is affected by isotropic white noise N_{ℓ}^{ch} , which does not correlate between channels. For $n_{\text{ch}} = 3$ channels $\{A, B, C\}$, the matrix of equation (2.2) reads

$$\mathbf{C}_{\ell} = \begin{pmatrix} C_{\ell}^{\text{A}\times\text{A}} + N_{\ell}^{\text{A}} & C_{\ell}^{\text{A}\times\text{B}} & C_{\ell}^{\text{A}\times\text{C}} \\ C_{\ell}^{\text{A}\times\text{B}} & C_{\ell}^{\text{B}\times\text{B}} + N_{\ell}^{\text{B}} & C_{\ell}^{\text{B}\times\text{C}} \\ C_{\ell}^{\text{A}\times\text{C}} & C_{\ell}^{\text{B}\times\text{C}} & C_{\ell}^{\text{C}\times\text{C}} + N_{\ell}^{\text{C}} \end{pmatrix}, \quad (2.6)$$

and the data vector reads

$$X_{\ell} = \left(X_{\ell}^{\text{A}\times\text{A}}, X_{\ell}^{\text{A}\times\text{B}}, X_{\ell}^{\text{A}\times\text{C}}, X_{\ell}^{\text{B}\times\text{B}}, X_{\ell}^{\text{B}\times\text{C}}, X_{\ell}^{\text{C}\times\text{C}} \right). \quad (2.7)$$

The HL likelihood requires an estimate of the noise bias and assumes a fiducial cosmological model to compute both X_{ℓ} and $M_{\ell\ell'}^{-1}$. In practice, both of these can be misestimated, either because the noise is not perfectly known, or because the fiducial spectrum differs from the true one. Although the second issue can be addressed by iterating the analysis over multiple fiducial models, the noise bias is trickier to correct since the true level may not be directly accessible. For this reason, we show the results of the wrong fiducial assumption in appendix A, as this would not be an issue in a real-life scenario. There we show that the HL is extremely stable against this effect, as expected [27]. mHL yields similar robustness, while the cHL appears more sensitive.

To handle this, a variation called the offset HL (oHL) likelihood was proposed [53] (section 2.1.1). Assuming uncorrelated noise between fields, one can build a likelihood based only on cross-spectra, avoiding the need to subtract noise bias explicitly. The name “offset” will be explained in the next section, where we will try to extend this method to multiple fields.

In this work, we propose a new alternative: the *marginalized HL* (mHL) likelihood, where we marginalize the HL likelihood over the auto-spectra to mitigate this problem (section 2.2).

We perform an extensive validation of these multi-field HL-based likelihoods, providing concrete recommendations for CMB data analysis. We assess their robustness when the noise levels and fiducial spectra used in the likelihood analysis are mismatched. This is done both in the full-sky, where the HL likelihood is exact, and in the cut-sky, where all of them are approximations. In addition, the usage of a multi-field approach can answer a simple, but very important dilemma. As mentioned before, the n_{ch} maps used here can be seen as data-splits from a “full-mission” dataset. Thus, with this full-mission map, we can ask ourselves what would be the corresponding single-field result. What we call from now the “single-field” HL is very useful to understand whether a future CMB analysis should rely on a single-field approach, or prefer a multi-field HL-based likelihood at large angular scales.

2.1.1 Offset Hamimeche & Lewis

Ref. [53] proposed a method to handle misestimates of the noise bias using only cross-spectra. In particular, the oHL likelihood was introduced and validated using pairs of frequency

channels (e.g., 100 GHz and 143 GHz from the *Planck* satellite [17, 18]), resulting in a single cross-spectrum. For example, $C_\ell^{100 \times 143}$ is transformed — following the same steps described earlier — into $X_\ell^{100 \times 143}$, which is then used in the likelihood computation. However, it is evident from the logarithm and square-root of equation (2.5) that cross-spectra can be problematic, as they are not necessarily positive-definite. To address this, ref. [53] introduced an offset O_ℓ , designed to mimic the presence of a noise bias and enable the cross-spectrum to behave like an auto-spectrum (see the appendix of [53] for further details).

The proposed offset is defined such that more than 99% of their Monte Carlo (MC) simulations yield a positive-definite matrix $\mathbf{C}_\ell^{-1/2} \hat{\mathbf{C}}_\ell \mathbf{C}_\ell^{-1/2}$. Thus, it depends on the noise levels in the simulated maps, the fiducial spectra $\mathbf{C}_{f\ell}$, and the mask used [53].

This offset, however, could be implemented in different ways, for example based on the expected width of the distribution of C_ℓ [12], bringing about different outcomes. Although this was discussed in ref. [57] and we give some extra detail in appendix B, we leave a detailed investigation of this dependence to future work. In the following sections, we always adopt the procedure used in [53].

To further avoid numerical issues and regularize the transformation in equation (2.5), ref. [53] proposed modifying the function g as

$$g(x) \longrightarrow \text{sign}(x)g(|x|). \quad (2.8)$$

Although this modification has only a minor impact once the offset is applied, it is important to stress that it is introduced solely for numerical stability — not for statistical reasons. As such, if the offset is not large enough to compensate for negative elements, this adjustment could bias the resulting distribution.

2.1.2 Extending oHL to multiple fields

One crucial aspect is still missing: the oHL is proposed for a single cross-spectrum, so it provides no way of combining the information from multiple cross-spectra. In our case, however, we would like to combine three independent cross-spectra: $C_\ell^{A \times B}$, $C_\ell^{A \times C}$, and $C_\ell^{B \times C}$. Following the logic of ref. [53], we try to extrapolate a multi-field generalization, which we will name cross-spectra HL (cHL). To define the cHL, one is easily tempted to build the data vector by concatenating these three spectra, ensuring that no auto-spectra are included. Thus, we can define the new data vector as

$$C_\ell^{\text{cHL}} = \left(C_\ell^{A \times B} + O_\ell^{A \times B}, C_\ell^{A \times C} + O_\ell^{A \times C}, C_\ell^{B \times C} + O_\ell^{B \times C} \right), \quad (2.9)$$

which is three times as long as $C_\ell^{A \times B}$, for example, and is then transformed using the HL prescription into X_ℓ^{cHL} .¹ Then, the procedure is the same as in section 2.1. Consider as

¹Since each cross-spectrum is treated independently when transforming C_ℓ^{cHL} , this approach effectively corresponds to using matrices like

$$\mathbf{C}_\ell^{\text{cHL}} = \begin{pmatrix} C_\ell^{A \times B} + O_\ell^{A \times B} & 0 & 0 \\ 0 & C_\ell^{A \times C} + O_\ell^{A \times C} & 0 \\ 0 & 0 & C_\ell^{B \times C} + O_\ell^{B \times C} \end{pmatrix}$$

in the HL procedure (compare this with equation (2.6)).

an example the $X_\ell^{A \times B}$ element of the HL from equation (2.7). Given that the data vector X_ℓ^{cHL} is obtained through a single-field HL transformation of equation (2.9), we note that $X_\ell^{A \times B, \text{cHL}} \neq X_\ell^{A \times B, \text{full HL}}$.

To our knowledge, such an approach was attempted in ref. [55], where cross-spectra from *Planck* and WMAP were combined into a single concatenated spectrum (see equation (2.9)). Something similar was also applied in ref. [56], where the *Planck* single-frequency spectra are included as cross-spectra of two data-split maps, treating them as auto-correlations on the diagonal of equation (2.2).

Of course, when $n_{\text{ch}} = 2$ the cHL is identical to the original oHL, given that it yields only one cross-spectrum. We study this case in section 4.3.

We conclude this section by emphasizing that cross-spectra are not the only reason C_ℓ can lose positive definiteness. When estimated on a partial sky, even auto-spectra can be negative [27]. The more aggressive the mask, the more likely the estimated auto-spectra will exhibit negative elements. While the offset in the oHL (and in the cHL) helps mitigate this, no such safeguard is included in the HL likelihood defined at the beginning of the section. Although the partial-sky case is, in principle, distinct from having two anti-correlated fields — since negative spectra in the former arise from the estimation process itself — introducing an offset in the diagonal of equation (2.6) may still be beneficial, despite the presence of noise bias. We investigate this possibility in the following.

2.2 The new auto-spectra marginalized Hamimeche & Lewis

We now propose a variant of the HL likelihood, which handles noise bias uncertainty by marginalizing over the auto-spectra, named for brevity the *marginalized HL* (mHL). The procedure is very similar to the one described in section 2.1, where, in our case, we start from matrices as equation (2.6). However, when computing the likelihood of equation (2.3), we employ only the “cross-elements” of the X_ℓ data vector defined in equation (2.7):

$$X_\ell^{\text{cross}} = \left(X_\ell^{A \times B}, X_\ell^{A \times C}, X_\ell^{B \times C} \right). \quad (2.10)$$

Correspondingly, the elements of the matrix $M_{\ell\ell'}$ associated with the “auto-elements” (e.g. $X_\ell^{A \times A}$) are removed. This approach relies on a simple but often overlooked observation: the HL likelihood has the structure of a Gaussian distribution for the X_ℓ , with covariance $M_{\ell\ell'}$. Indeed, it can be shown empirically that $X_\ell \sim \mathcal{N}(\mu, M_{\ell\ell'})$. Therefore, we can apply the standard marginalization rule for a multivariate Gaussian to obtain a likelihood defined only on X_ℓ^{cross} , without the need to specify μ .² Note that, if $X_\ell \sim \mathcal{N}(\mu, M_{\ell\ell'})$ this procedure is fully analogous to performing the integral of the HL likelihood over the auto-elements of X_ℓ , similarly to what is done in ref. [48] to obtain the probability of $\widehat{C}_\ell^{\text{TE}}$ marginalized over $\widehat{C}_\ell^{\text{TT}}$ and $\widehat{C}_\ell^{\text{EE}}$. The difference is what is being marginalized: in the mHL we integrate out

²For a random variable $Y = (Y_1, Y_2, Y_3) \sim \mathcal{N}(\mu, \Sigma)$ with $\mu = (\mu_1, \mu_2, \mu_3)$ and Σ a 3×3 covariance matrix, the marginal distribution on Y_2 is:

$$\begin{aligned} \tilde{Y} &\sim \mathcal{N}(\tilde{\mu}, \tilde{\Sigma}) \\ \tilde{\mu} &= (\mu_1, \mu_3) \\ \tilde{\Sigma} &= \begin{pmatrix} \Sigma_{11} & \Sigma_{13} \\ \Sigma_{13} & \Sigma_{33} \end{pmatrix}. \end{aligned}$$

the X_ℓ^{auto} , treated as random variables, exploiting the fact that the result of that integral is known for a multivariate Gaussian; instead, ref. [48] marginalizes over the individual $\hat{C}_\ell^{\text{A}\times\text{A}}$.

Also, we emphasize that the mHL and the HL likelihood share the same assumption of a Wishart-distributed C_ℓ . In turn, this translates into assuming a Gaussian and statistically isotropic field in pixel-space; thus, any violation of these properties makes both the mHL and the HL likelihood approximations of the true unknown distribution. The only difference is the information retained for parameter estimation: the HL exploits the full information encoded in X_ℓ (see equation (2.7)), while the mHL just the one in X_ℓ^{cross} (see equation (2.10)).

Hence, a key advantage of this method is its flexibility: it allows one to selectively retain any well-estimated auto-spectrum simply by adjusting the data vector and covariance accordingly. For instance, if the auto-spectrum for the C-channel $C_\ell^{\text{C}\times\text{C}}$ is deemed reliable, one could define $X_\ell = (X_\ell^{\text{C}\times\text{C}}, X_\ell^{\text{A}\times\text{B}}, X_\ell^{\text{A}\times\text{C}}, X_\ell^{\text{B}\times\text{C}})$. This is not as straightforward to do in the cHL framework. Moreover, since the algebra used to construct X_ℓ^{cross} is identical to that of the HL likelihood, each element remains unchanged; for instance, $X_\ell^{\text{A}\times\text{B},\text{mHL}} = X_\ell^{\text{A}\times\text{B},\text{HL}}$. This is not the case for the cHL (see section 2.1.1).

In principle, this approach avoids defining an offset, or modifying equation (2.5), since it uses the original HL formulation. However, as already mentioned, when dealing with partial-sky data, negative eigenvalues can arise in the HL transformation from the matrix $\mathbf{C}_\ell^{-1/2} \hat{\mathbf{C}}_\ell \mathbf{C}_\ell^{-1/2}$. To mitigate this cut-sky effect and stabilize the eigenvalues, we add an offset only to the diagonal elements of the various matrices involved, as mentioned above for the HL. Specifically, we assume the same offset definition used in the oHL analysis of [53]. Differently from the oHL (and cHL), we compute this offset for the auto-spectra, e.g. $O_\ell^{\text{A}\times\text{A}}$, which is typically smaller than its cross counterparts in the cHL, say $O_\ell^{\text{A}\times\text{B}}$. Moreover, as f_{sky} goes to 100%, $O_\ell^{\text{A}\times\text{A}}$ goes to zero since no auto-spectrum is negative on the full sky, where the power spectrum estimation does not introduce this artifact. This does not happen for $O_\ell^{\text{A}\times\text{B}}$: the noise realizations are independent; thus, if the noise is large enough with respect to the signal, the final C_ℓ can be negative when the noise anti-correlates between channels. Furthermore, to ensure complete numerical stability, we use the same change of the g function as for the oHL (see equation (2.8)).

In appendix B, we show the performance of both the HL and mHL likelihoods when this regularization step is omitted and when we use two possible definitions of the offset.

2.3 Pixel-based likelihood

Finally, we emphasize once again that none of the likelihoods discussed so far are exact on a cut sky. On the partial sky, C_ℓ does not distribute as a Wishart; thus, we must assume otherwise to apply the HL prescription. For this reason, we also compute the Pixel-Based (PB) likelihood [28], which is exact under the assumption of Gaussian-distributed pixel power. Since we consider only map simulations obtained by summing CMB and white-noise realizations, we can safely adopt this assumption and compute analytically the pixel covariance matrix. The PB likelihood serves therefore as a reference against which the various approximations are evaluated, HL included. The PB is computed with the publicly available `parameter_estimation` algorithm.³

³https://baltig.infn.it/cosmology_ferrara/lowell-likelihood-analysis/-/tree/master/parameter_estimation

Parameter	Fiducial value
$A_s 10^{-9}$	2.12
n_s	0.9651
$\Omega_b h^2$	0.02237
$\Omega_{\text{cdm}} h^2$	0.1201
H_0	67.32 km/s/Mpc
τ	0.06
r	0 (0.01)
$\sum m_\nu$	0.06

Table 1. Fiducial values for the cosmological parameters. In parenthesis we report the value of the tensor-to-scalar ratio in the tensor-detection case.

As mentioned above, the three channels we study in this work can also be seen as three data-splits of a hypothetical mission. Thus, we compute the PB on the full-mission dataset, after combining the three maps from the channels. In this sense, the PB will contain the full information, which we try to recover with the multi-field HL-based likelihoods.

3 Methodology

3.1 Our setup

In this work, we assess the robustness of HL-based likelihoods against misestimation of the noise bias. We consider either C_ℓ^{EE} or C_ℓ^{BB} with $n_{\text{ch}} = 3$ channels, within the context of both the Λ CDM and Λ CDM + r models, to estimate τ and r , respectively (see section 4.3 for the $n_{\text{ch}} = 2$ case). Specifically, we define three test cases:

- Signal-dominated case: using E-modes to estimate the optical depth τ in a regime where the signal dominates over the noise. In particular, we assume a white-noise level of $N^E = 80 \mu\text{K-arcmin}$ for each of the three E-mode channels (A , B , and C).
- Noise-dominated case: using B-modes under the null-signal-hypothesis of $r = 0$. Here, we instead assume $N^B = 10 \mu\text{K-arcmin}$ for all channels.
- Tensor-detection case: again using B-modes, but now with $r = 0.01$, representing an intermediate regime between the previous two. Indeed, given that we assume the same noise level of the noise-dominated case, we allow for a detectable r without being completely signal-dominated.

For both the noise-dominated and tensor-detection cases, we assume the presence of the so-called lensing B-mode contribution due to the deflection of E-modes caused by mass distributions. On top of this, maps are smoothed with a cosine beam to limit their harmonic content without sharp cuts [18, 58].

The fiducial values of the cosmological parameters are chosen to be consistent with current constraints [13, 17, 22, 59, 60] and are summarized in table 1. Of course, the tensor-

detection case is an exception, as we set $r = 0.01$ instead. All theoretical CMB angular power spectra are computed with CAMB [61].

3.2 Simulations

For each fiducial power spectrum, we generate $N = 10000$ simulations using the `healpy` package. Maps are created at $N_{\text{side}} = 16$, corresponding to an angular resolution of approximately $220'$ and a total of 3072 equal-area pixels over the full sky. Harmonic information is extracted using a Quadratic Maximum Likelihood (QML) algorithm. Similarly to the PB likelihood, the QML requires the computation of the pixel covariance matrix, understood as the sum of a signal contribution from CMB and a diagonal noise contribution (isotropic white noise). While not required for full-sky coverage, QML ensures consistency with our cut-sky analysis. In particular, we adopt the `pse_qml`⁴ algorithm.

The first $N_{\text{sims}} = 9000$ simulations are used to estimate the harmonic covariance matrix, the fiducial spectra, and eventually the offset (see equation (2.3)). The remaining $N_{\text{data}} = 1000$ form the validation set. To avoid pixelization effects, we set the maximum multipole to $\ell_{\text{max}} = 2 \times N_{\text{side}} = 32$, in line with typical large-scale CMB analyses that use $\ell_{\text{max}} \simeq 30$.

As mentioned above, each simulation is characterized by some level of white-noise (see section 3.1). To mimic misestimation of these noise levels, we generate additional simulations in which the first N_{sims} simulations use incorrect noise amplitudes. Specifically, we apply a $\pm 20\%$ scaling to the true noise at the map level, both over- and underestimating it. We denote the estimated noise power spectrum as \hat{N}_ℓ and the true one as N_ℓ^{True} . This 20% error is a deliberately pessimistic assumption.⁵ Note that in this work we focus on the likelihood characterization of the CMB sky. Thus, this mismatch is never introduced in the power spectrum estimation step of our analysis. In a real-life scenario, these two would combine; however, here we isolate and study solely the likelihood effect.

To evaluate stability with respect to the fiducial cosmological value assumed, we also vary the fiducial value of r in the noise-dominated and tensor-detection cases, swapping them between 0 and 0.01, and consider an additional scenario with $r = 0.02$. For the signal-dominated case, we explore values $\tau = 0.03$ and $\tau = 0.09$ to test both under- and overestimation. Of course, this choice affects just the first N_{sims} simulations to mimic once again a wrong estimate. These values of the fiducial are deliberately extreme tests, as in real data analyses, strongly biased fiducials would be obvious and can be corrected through iterative re-analysis. As such, we relegate the results of these tests to appendix A.

We also repeat all analyses under partial-sky conditions. Although the full-sky case highlights the importance of proper likelihood usage under ideal conditions, we are also interested in more realistic scenarios with incomplete sky coverage. We use the publicly available HFI masks that retain 40%, 60%, and 80% of the sky, respectively.⁶ Each simulation is masked, and its power spectrum is estimated using QML. This allows us to examine any difference in likelihood behavior caused by the cut-sky. To isolate the effects of mismatches in the likelihood, we always assume the correct noise and fiducial values when estimating the

⁴https://baltig.infn.it/cosmology_ferrara/lowell-likelihood-analysis/-/tree/master/pse_qml

⁵Note that a 20% misestimation at map-level brings to a $\simeq 40\%$ mismatch in harmonic space.

⁶<https://pla.esac.esa.int/#maps>

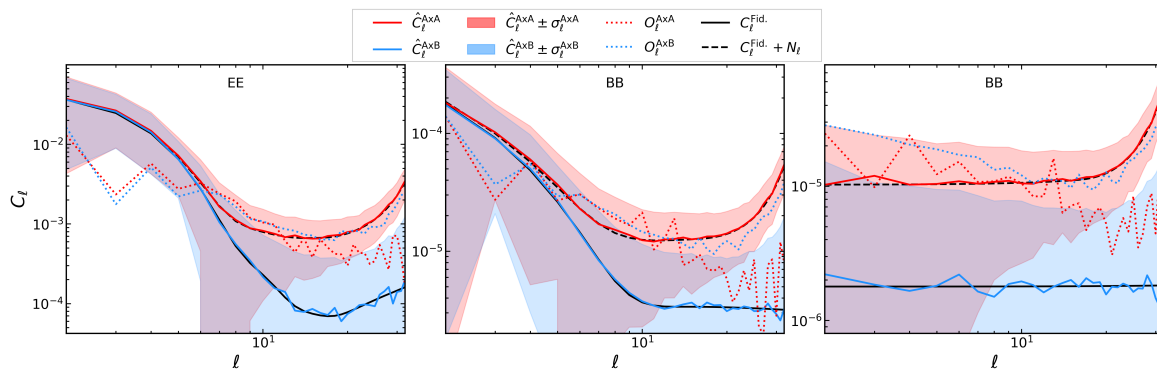


Figure 1. Auto- (red) and cross-spectra (blue) for the signal-dominated, tensor-detection, and noise-dominated cases, respectively. Solid lines indicate the mean over the N_{data} simulations, while shaded regions refer to the 1σ scatter around the mean and dotted ones to the corresponding offset O_ℓ . As a reference, we also show the fiducial power spectrum with and without noise bias with black solid and dashed lines.

power spectra. Studying the effect of incorrect assumptions during power spectrum estimation lies beyond the scope of this work and has been addressed elsewhere in the literature [62–65].

Before discussing how we recover the posterior of each case, we show in figure 1 an example of the spectra we are using for the signal-dominated, tensor-detection, and noise-dominated cases, respectively. In particular, we choose $A \times A$ and $A \times B$ as representatives of auto- and cross-correlations, plotting their recovered spectra with $f_{\text{sky}} = 40\%$.⁷ In solid lines, we show the mean of the N_{data} simulations, the shaded regions indicate the one-sigma scatter, and the dotted lines the offsets. In addition, we show the fiducial spectrum with and without the noise bias as a reference. Here, the effect of the beam function mentioned above is clearly visible as a degradation of the noise bias towards small scales. This plot shows that our QML estimation is unbiased. Also, one can appreciate that the offsets of $A \times A$ are comparable with the ones of $A \times B$, given the large number of negative spectra found assuming $f_{\text{sky}} = 40\%$. For the sake of brevity, we do not show this here, but while f_{sky} increases, the two offsets decrease and depart from each other, with $O_\ell^{A \times A}$ tending to zero faster than $O_\ell^{A \times B}$ (see definition in section 2.1.1).

3.3 Posterior evaluation

To validate all these HL-based likelihoods, we aim to estimate τ from E-modes and r from B-modes, separately and with all other cosmological parameters held fixed (see table 1). Thus, we evaluate the likelihood over a fixed grid of values on the considered parameter.

We define a uniform grid for τ in the range $[0.03, 0.09]$ using 960 steps, and for r in the range $[0.0, 0.02]$ using 800 steps. This is equivalent to assuming a flat prior over the specified ranges [23].

Using this setup, we compute the posterior for each of the N_{data} simulations. Doing the product over these posteriors yields a single mean posterior, which allows us to assess whether any bias has been introduced during the analysis pipeline.

⁷We do not show the full-sky case as it is trivial.

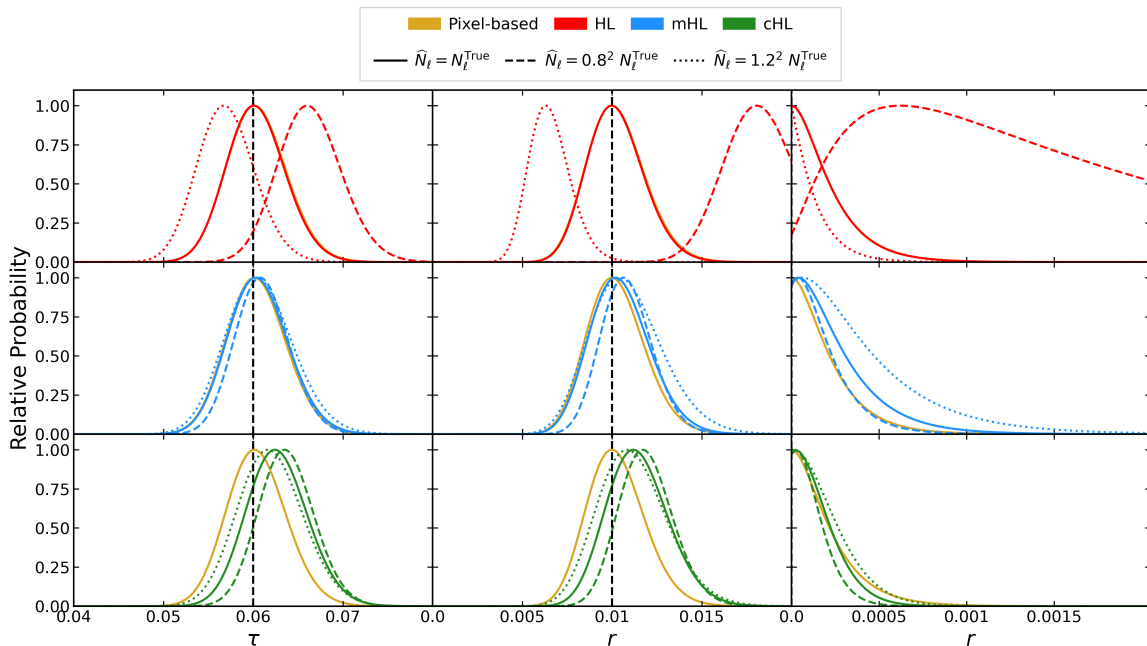


Figure 2. Posterior distributions from pixel-based (gold), HL (red), mHL (blue), and cHL (green) likelihoods under full-sky assumptions. Dashed (dotted) curves indicate under- (over-) estimation of the noise bias. Vertical black lines denote the true parameter values.

The code developed and used to perform this analysis is publicly available as a `Python` package named `CMB Harmonic Analysis with Robust Multi-field Likelihood` ([CHARM-Like](#)). Providing a set of precomputed angular power spectra for some channels, `CHARM-Like` allows us to quickly compute the posteriors of HL, mHL and cHL, together with all the single field HL.

4 Results

4.1 Full-sky

We start from the full-sky case, where every component of the analysis is perfectly characterized. In particular, the HL result is exact, and no mode coupling arises from masking effects. Figure 2 summarizes the results for the HL, mHL, and cHL likelihoods, including cases where the noise bias is misestimated. Each row (color-coded) corresponds to a different likelihood, while each column refers to the signal-dominated, tensor-detection, and noise-dominated cases, respectively. Dashed lines represent an underestimation of the noise bias, and dotted lines represent an overestimation.

Starting with the HL, all posteriors are unbiased when the correct noise bias is used. In fact, they appear exactly identical to the PB results. The same applies to the single-field HL, which also overlaps perfectly with the PB, and so is not shown here for clarity.

However, when a mismatch is introduced, a bias appears.⁸ As expected, the bias is less severe when the signal dominates (left column of figure 2). Both mHL and cHL mitigate

⁸Here, we do not show the PB and the single-field HL with noise bias misestimation. In fact, those two serve for comparison with the other HL-based likelihoods, and it is known that they are quite sensitive to such a mismatch.




Figure 3. Dependence of cHL on the number of cross-spectra used. Results are shown for one, two, and three spectra for the tensor-detection case in full-sky.

this bias to varying degrees in all cases: in the signal-dominated case the mHL shows a smaller dispersion with respect to the cHL, thus performing better. In the tensor-detection case, they perform similarly, while in the noise-dominated case the cHL behaves better. Indeed, even though auto-correlations are marginalized over, the incorrect noise still enters the HL transformation algebra.

The mHL generally yields more conservative results. This stems from the marginalization, effectively discarding some information captured by the HL. However, this degradation is significant only in the noise-dominated case (and affects marginally the tensor-detection case), suggesting a dependence on the signal-to-noise ratio.

A striking feature in figure 2 is the bias observed in the cHL even when the correct noise bias is used. This originates from applying the HL transformation independently to each cross-spectrum (see equation (2.9)). This bias increases with the number of cross-spectra used; this is shown in figure 3 for the tensor-detection case.

With a single cross-spectrum, as in [12], the bias is null, but it becomes more pronounced when multiple channels are included.

4.2 Cut-sky

Having validated the likelihoods in an ideal full-sky context, we now turn to a more realistic scenario with partial sky coverage. In this case, even the HL likelihood is no longer exact, and the PB approach must serve as our benchmark.

For brevity, we present results for $f_{\text{sky}} = 40\%$, the most extreme case considered. Additional results for $f_{\text{sky}} = 60\%$ are provided in appendix B, while the $f_{\text{sky}} = 80\%$ case closely resembles the full-sky scenario and is therefore omitted.

Figure 4 shows the results under noise bias mismatch for $f_{\text{sky}} = 40\%$ and summarizes the main results of this work. The first three rows, analogous to figure 2, clearly show broader posteriors, a direct consequence of information loss due to the mask.⁹

Focusing on the solid lines (i.e., with correct noise bias) and leaving aside the noise-dominated case for a moment, all likelihoods perform reasonably well in recovering the pixel-based results. Nonetheless, the mHL stands out as the best-performing approximation,

⁹Here, we do not plot the single-field HL posteriors for visualization purposes. Instead, we relegate it to the bottom panel, which is described below.

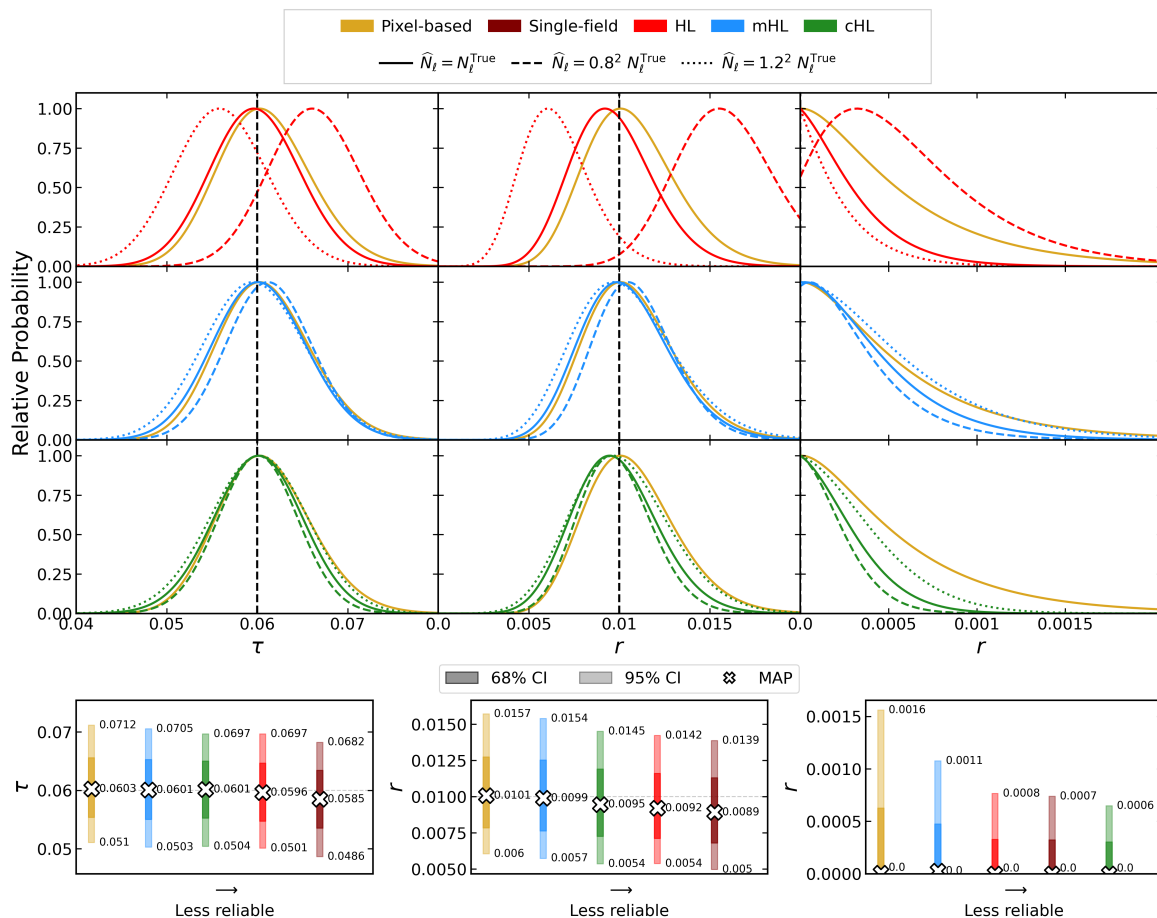


Figure 4. Posterior distributions from pixel-based (gold), HL (red), mHL (blue), and cHL (green) likelihoods for $f_{\text{sky}} = 40\%$. The columns correspond to signal-dominated, tensor-detection, and noise-dominated cases. In the first three rows, dashed (dotted) curves correspond to under- (over-) estimated noise bias. Vertical black lines mark the true parameter values ($\tau = 0.06$, $r = 0.01$ and $r = 0$). The bottom row provides a quantitative comparison of the solid-line posteriors above, together with the single-field HL (maroon). Crosses show MAP estimates, while shaded bars indicate 68% and 95% credible intervals.

consistently matching the PB across all cases, without showing the same biases affecting the others in the tensor-detection case, for example. In the noise-dominated case, the better accuracy of the mHL is particularly evident when comparing it to the others. We emphasize the crucial role played by the offset, as further discussed in appendix B, where we study the effect of removing it in the computation of both the HL and the mHL. In addition, we test in the same appendix two different offset prescriptions [12, 53]. Appendix B highlights another strength of the mHL: its robustness to the arbitrary choice of offset, owing to the fact that it marginalizes over the auto-spectra where the offset is introduced.

The HL shows a strong signal-to-noise dependence. While τ is well recovered in the signal-dominated case, it significantly underestimates uncertainty on r in the noise-dominated case. In contrast, mHL remains robust, even improving over its full-sky performance in the noise-dominated case. This can be pictured as the cut-sky introducing an additional

scatter in the estimated power spectra — an extra “noise contribution” not accounted for, which the mHL marginalizes over.

cHL shows a systematic shift to the left due to sky masking, increasing in magnitude as the noise becomes more important.¹⁰ Thus, in the signal-dominated case, the leftward bias due to the mask is balancing the rightward bias affecting the cHL in full-sky (see figure 2). In other words, this is serendipitous as shown in appendix B for $f_{\text{sky}} = 60\%$. Similarly to the HL, the cHL strongly underestimates the uncertainty on r in the noise-dominated case. Indeed, similarly to the HL, the posterior approaches $r = 0$ with a negative derivative, indicating that the actual peak of the posterior would be in the non-physical region $r < 0$ (a similar leftward bias is seen also in the tensor-detection case).

As for the impact of a mismatch in the noise bias, the results align with those in figure 2, suggesting that the effects of noise mismatch and partial sky coverage interact in a trivial way. The only notable difference is that the noise mismatches lead here to smaller biases. This is simply understood as the mask itself is producing an extra scatter, making the actual noise bias less important than before.

Finally, the bottom row of figure 4 presents a quantitative comparison between the posteriors with correctly estimated noise bias (solid lines in the upper panel of the same figure). Here, we also show the results of the single-field HL, of which we do not plot the posterior for visualization purposes. The approximations are ranked by their ability to recover the PB results. While this is qualitatively evident, we compute the Kullback-Leibler (KL) divergence between the PB and each approximation to assess quantitatively this correspondence. Thus, the KL divergence allows us to rank them accordingly. In particular, we show the Maximum A Posteriori (MAP) estimate and the 95% credible interval for each case.

This ranking confirms that mHL offers the best agreement with PB while maintaining robustness to noise bias mismatches. The worst performance occurs in the noise-dominated case, where even the mHL underestimates the 95% upper bound by $\simeq 28.8\%$. The HL and cHL perform significantly worse, with underestimations of approximately $\simeq 49.5\%$ and $\simeq 57.1\%$, respectively.

It is also worth noting that in all noise mismatch cases, the mHL provides the best KL divergence with respect to the true posterior. The only exception is the cHL in the signal-dominated case, which scores slightly better than the mHL. This means that, despite a possible misestimation of the noise bias, mHL still gives the most accurate and robust posterior among the ones studied in almost all the cases considered.

In addition, we note that the single-field HL tends to align with the HL. This suggests that applying the HL likelihood on the full-mission dataset, or on multiple data splits, is nearly equivalent when one employs all fields. This emphasizes the strength of the mHL, which allows for improvements in accuracy with respect to the PB, while offering robustness against an eventual mismatch of the noise bias. Thus, this suggests that a multi-field mHL likelihood should be preferable to a single-field HL for CMB analyses.

To conclude this section, we repeat the analysis while adopting the familiar, though notoriously incorrect, assumption that the C_ℓ distribution is Gaussian. Even though this

¹⁰Not shown here for brevity, but reducing f_{sky} from 100% to 40% progressively shifts the cHL leftward with respect to the full-sky result.

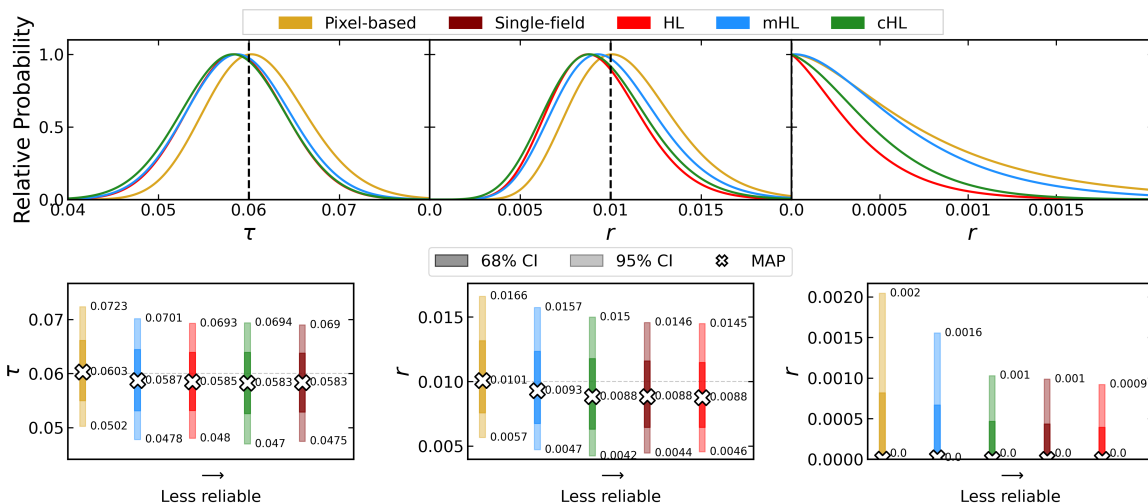


Figure 5. Comparison between pixel-based (gold), HL (red), mHL (blue), and cHL (green) likelihoods with $f_{\text{sky}} = 40\%$ using only channels A and B. In this case, cHL and oHL are identical. The columns correspond to signal-dominated, tensor-detection, and noise-dominated cases. The black vertical lines represent the true values of the parameters ($\tau = 0.06$, $r = 0.01$ and $r = 0$). The bottom row provides a quantitative comparison of the posteriors above, together with the single-field HL (maroon). Crosses show MAP estimates, while shaded bars indicate 68% and 95% credible intervals.

is a rougher approximation with respect to the HL on large scales [28], it is interesting to see how that performs. In appendix C, we show that this approximation brings distortions of the posterior, as expected.

4.3 Using only two channels

In the previous section, we focused on a multi-field case with three “channels”, or data splits, observing the same CMB signal with different noise realizations. This setup gives us three cross-spectra to build the mHL and cHL approximations.

In section 4.1, we mentioned a notable feature of the cHL approximation that depends on the number of cross-spectra used. While we do not explore this effect further, it is useful to examine the two-channel case more closely. This setup is common in the literature and is the only multi-field configuration with a single cross-spectrum [12, 23].

We repeated the analysis from section 3 using only channels A and B. In this case, the cHL likelihood definition is identical to the oHL, so we will use the latter name to remark this fact.

Figure 5 shows the results for $f_{\text{sky}} = 40\%$. There are two notable differences compared to the three-channel case. First, in the signal-dominated and tensor-detection cases, the HL and mHL do not match the PB results perfectly. Still, all three approximations — HL, mHL, and oHL — give similar posteriors, broadly consistent with the true one. The mHL performs better overall.

In the noise-dominated case, the HL and oHL underestimate the uncertainty on r , as seen in section 4. In contrast, the mHL performs better than the others. In particular, they all underestimate the 95% upper bound: the mHL is off by $\simeq 23.9\%$, the oHL by $\simeq 49.7\%$, and the HL by $\simeq 55\%$.

Finally, the single-field HL tends to align with the HL also in this case. This suggests once again that it is equivalent to apply the HL approach to one field or to multiple ones. Let us emphasize that here the full-mission dataset used for this single-field HL and the PB is obtained by combining solely channels A and B, to be consistent with the other approximations.

These results, however, are not directly comparable with the literature [12]. Indeed, here we use a different prescription for the offset. Similarly to appendix B, we repeat this two-channel analysis using LoLLiPoP offset [12]. For the sake of brevity, we do not show the results here. However, in this context, oHL and mHL provide very similar results in the noise-dominated case, with the oHL performing slightly better. On one side, this supports the use of oHL in single-cross-spectrum analyses of noise-dominated scenarios (e.g., [12]), and shows that the mHL is a competitive alternative. On the other hand, this emphasizes once more the importance of the choice of the offset.

5 Conclusions

In this work, we investigated the performance of several likelihood approximations based on the Hamimeche-Lewis (HL) prescription [27], focusing in particular on two extensions: the newly introduced *marginalized HL* (mHL) and the cross-spectra HL (cHL) — an extension of the offset HL (oHL) likelihood [53]. We studied the impact of realistic complications, such as partial sky coverage, mismatches in noise bias estimation, and incorrect fiducial assumptions. Our analysis was carried out in a multi-field context, where three observational channels, or data splits, of CMB polarization were jointly used for parameter estimation. Given that the distribution of data is more easily approximated by the HL likelihood on intermediate and small scales [27], thanks to the central limit theorem, we focused on the large-scale part of CMB, where the harmonic approximation represents a bigger challenge.

We began by validating each method in the idealized full-sky case (figure 2). As expected, the HL delivered unbiased and precise results when the noise bias was correctly modeled, but suffered from significant biases when it was not. Both the mHL and cHL mitigated these biases, with the mHL providing more conservative but robust constraints. We also identified a bias in the cHL when combining multiple cross-spectra, linked to the number of spectra considered: the more one uses, the worse the bias becomes (see section 4.1). In appendix A, we showed that the HL is particularly stable under fiducial mismodeling as expected from ref. [27]. The mHL displayed similar robustness, while the cHL exhibited stronger sensitivity due to its reliance on the fiducial spectrum through the offset.

In the cut-sky scenario, where even the HL becomes approximate, we benchmarked all methods against the pixel-based (PB) likelihood [28], computed using the “full-mission” dataset obtained by combining the three data splits. The mHL showed excellent agreement with the PB posterior, benefiting from its marginalization over auto-spectra, which helps absorb inaccuracies from noise modeling and sky cuts (figure 4). Both the HL and cHL tended to underestimate uncertainties, particularly in the noise-dominated case (see section 3). In this context, we also used an offset for the HL and mHL, and we showed in appendix B that the mHL remains stable under different offset prescriptions, thanks to the marginalization over auto-correlations.

To quantitatively compare the methods, we used the Kullback-Leibler divergence as a ranking metric. This confirmed the superior performance of the mHL in recovering the PB posterior across all tested scenarios. Even in the challenging noise-dominated regime, the mHL maintained acceptable accuracy, significantly outperforming both the HL and the cHL.

Using the same full-mission dataset, we also computed a “single-field HL” likelihood, corresponding to a single-field case where all data splits are combined into one map. Comparing results in figure 4, we found that the multi-field mHL approach performs significantly better than the single-field HL, suggesting that a multi-field treatment should be preferred for CMB data analysis. It not only improves accuracy but also enhances robustness against misestimations of the noise bias.

Overall, the marginalized HL emerges as a promising alternative to both the HL and the cHL. It is particularly well-suited for scenarios with imperfect noise knowledge and limited sky coverage, offering a strong balance between robustness and precision.

In section 4.3, we tested these methods in the two-channel case, a configuration often used in the literature [12, 23], which yields a single cross-spectrum. This is particularly interesting because the cHL definition (see section 2.1.2) is identical to the oHL one, allowing a comparison with ref. [53] and [12]. The mHL and the oHL provide very similar results in this case, with slight differences brought by the specific implementation of the offset (see also appendix B). This supports the use of oHL in single-cross-spectrum analyses of noise-dominated scenarios (e.g., [12]), and shows that the mHL is a viable and robust alternative.

As shown in [27] for the full-sky case, the distribution of the angular power spectrum estimator built from cross-correlations of many maps tends toward a Wishart distribution as the number of maps increases. Specifically, this number scales approximately as $n_{\text{ch}} \gg 1 + N_\ell / (C_\ell + N_\ell)$, depending on the signal-to-noise ratio. When noise is negligible, a single cross-spectrum may be sufficient, while in noise-dominated regimes, many more maps are needed, with $n_{\text{ch}} \gg 2$. Since the HL transformation is designed for Wishart-distributed variables, working closer to this regime is desirable. In this context, both the HL and the mHL likelihoods follow a similar behavior: in the signal-dominated case, they recover the exact solution accurately, but their performance degrades progressively toward the noise-dominated case. Additionally, using $n_{\text{ch}} = 2$ rather than $n_{\text{ch}} = 3$ yields worse results overall, as expected from the logic presented in [27]. The same does not apply to the cHL, which is more and more biased as the number of cross-spectra increases (see appendix B). Although the mHL currently provides the most accurate approximation to the exact solution among the tested methods, these observations suggest that increasing the number of available maps could further improve performance.

Finally, in appendix C, we tested a Gaussian likelihood approximation [28], often used for its simplicity. As expected, this assumption fails to capture the non-Gaussian features of the true C_ℓ distribution, leading to various distortions of the posteriors (see the appendix for more details).

Before concluding, it is also important to emphasize that the analysis presented here is applicable to any set of fields, not only data splits of a single observable — e.g. E-modes in the signal-dominated case. Thus, instead of $\{E_{\text{chA}}, E_{\text{chB}}, E_{\text{chC}}\}$, one could consider

$\{T_{\text{chA}}, T_{\text{chB}}, T_{\text{chC}}, E_{\text{chA}}, E_{\text{chB}}, E_{\text{chC}}\}$, or any other combination. The only rule to correctly apply the mHL prescription is to first build the full matrix as equation (2.2) starting from the spherical harmonic coefficients of the desired fields. Then, one should apply the HL transformation and marginalize on combinations considered unreliable.

We also emphasize that our analysis assumes an idealized scenario with uncorrelated white noise. In practice, the noise structure can be more complex — particularly in the presence of foregrounds and systematics — which may introduce non-trivial correlations. In this context, the ℓ - ℓ' independence is violated due to the mask; however, we already showed in section 4.2 that the mHL handles this deviation well, yielding accurate and robust results. However, one could also have $m - m'$ couplings and/or m -dependence of the covariance matrix, breaking statistical isotropy. Depending on how strong these deviations are, the HL likelihood could cease to be a good approximation, even in the full-sky case. A detailed study of such effects is beyond the scope of this work and is left for future exploration.

Several directions for future work are now open. A more detailed study of the offset computation would be valuable. Although the mHL appears robust against different offset definitions, a systematic investigation is needed to fully assess its stability, together with the dependence of other HL-based likelihoods. It would also be useful to explore the behavior of these likelihoods when increasing the number of fields. Moving from three to four channels is straightforward in principle, but studying the effects systematically could provide further insight. In addition, applying the mHL approach to real datasets could test its practical advantages over current analysis pipelines. Finally, while we focused here on large angular scales, where statistical uncertainties are significant, the same marginalized approach can be applied to smaller scales, e.g. in the context of ground-based experiments such as BICEP/Keck array [21], Simons Observatory [66], CMB-S4 [67].

Acknowledgments

The authors thank Loris P.L. Colombo, Matthieu Tristram and W.L. Kimmy Wu for their valuable comments and discussions. Some of the results in this paper have been derived using the following packages: CAMB [61], healpy [68], Matplotlib [69], SciPy [70], Numba [71] and NumPy [72]. We acknowledge the financial support from the INFN InDark initiative and from the COSMOS network through the ASI (Italian Space Agency) Grants 2016-24-H.0 and 2016-24-H.1-2018. M.G. and M.L. are funded by the European Union (ERC, RELiCS, project number 101116027). M.G. acknowledges support from the PRIN (Progetti di ricerca di Rilevante Interesse Nazionale) number 2022WJ9J33. This work was supported by the MUR PRIN2022 Project “BROWSEPOL: Beyond standaRd mOdel With coSmic microwavE background POLarization”-2022EJNZ53 financed by the European Union - Next Generation EU. P.C.’s work has been supported by the Fondazione ICSC, Spoke 3 Astrophysics and Cosmos Observations. National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) Project ID CN_00000013 “Italian Research Center on High-Performance Computing, Big Data and Quantum Computing” funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di “campioni nazionali di R&S (M4C2-19)” - Next Generation EU (NGEU).

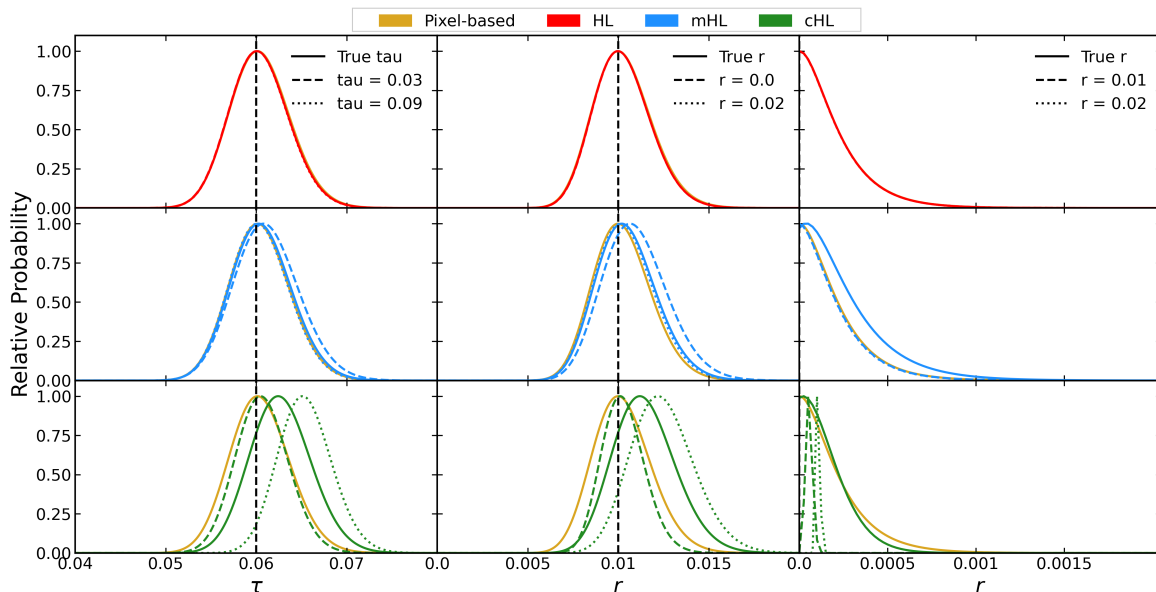


Figure 6. Comparison between pixel-based (gold), HL (red), mHL (blue), and cHL (green) likelihoods while assuming different fiducial cosmological values in full-sky. The columns correspond to signal-dominated, tensor-detection, and noise-dominated cases. Dashed and dotted curves indicate a misestimation of the fiducial spectra. The specific values for the fiducial parameters are shown in each column. The black vertical lines refer to the true values of the parameters ($\tau = 0.06$, $r = 0.01$ and $r = 0$).

A Assuming a wrong fiducial model

Figure 6 presents the results for a misestimation of the fiducial spectrum. Note that here we focus on the likelihood effects of a wrong estimation, not on the power spectrum estimation step.

The structure of the plot is analogous to figure 2, but here the dashed and dotted lines reflect errors in the fiducial assumptions for either τ or r .

Notably, the HL is extremely stable against such a mismatch. While this stability is expected [27], it is interesting to observe that it holds even for significantly erroneous fiducial values.

In contrast, the mHL demonstrates a discrete robustness against fiducial mismatches. In the noise-dominated case, when a wrong fiducial r , the posterior seems to collapse to the correct result, irrespective of the fiducial value assumed. However, this is just serendipitous. We tested other values of the fiducial between $r = 0$ and $r = 0.01$, and the posterior shifts continuously from the correct result to the dashed line in the plot.

On the other hand, the cHL likelihood shows a strong dependence on the assumed fiducial. This behavior is linked to the offset, which directly relies on the chosen fiducial spectrum. Indeed, focusing on the noise-dominated case, when r is assumed to be 0.01 or 0.02, the offset gets significantly suppressed, to the point that on the very first multipoles it drops to zero. As a consequence, nothing is preventing negative values from generating numerical problems, that manifest here in two very sharp peaks.

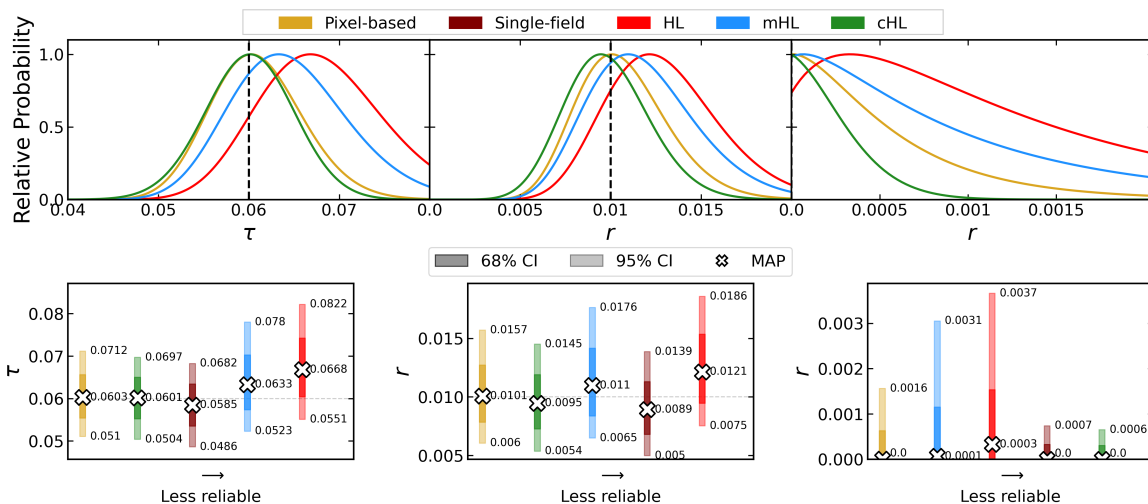


Figure 7. Comparison between pixel-based (gold), HL (red), mHL (blue), and cHL (green) likelihoods with $f_{\text{sky}} = 40\%$ omitting the offset. The columns correspond to signal-dominated, tensor-detection, and noise-dominated cases. The black vertical lines represent the true values of the parameters ($\tau = 0.06$, $r = 0.01$ and $r = 0$). The bottom row provides a quantitative comparison of the posteriors above, together with the single-field HL (maroon). Crosses show MAP estimates, while shaded bars indicate 68% and 95% credible intervals.

It is worth reiterating that a mismatch in the fiducial spectrum can be “detected” and resolved iteratively. The key takeaway here is that the cHL requires careful mitigation of this effect, whereas the HL and mHL remain more stable in the presence of such mismatches.

B The role of the offset

Building on the results shown in figure 4, we now explore the role of the offset in the HL and mHL results. As mentioned in section 3, the issue of negative spectra manifests in the HL transformation as negative eigenvalues in the matrix $\mathbf{C}_\ell^{-1/2} \hat{\mathbf{C}}_\ell \mathbf{C}_\ell^{-1/2}$. To mimic the oHL (and cHL), we add an offset to the diagonal elements of the various matrices involved to regularize their eigenvalues. This brings us to the main results of this work, i.e. figure 4, while figure 7 shows the equivalent results without adding an offset to the HL and the mHL.¹¹

Note that here we report solely the case in which $\hat{N}_\ell = N_\ell^{\text{True}}$ as the mismatch of the noise bias behaves as expected. Thus, we collapse the 3×3 grid of figure 4 into a single row comparing directly the three approximations.

We observe that the cut-sky introduces a bias into the HL approximation, causing it to significantly diverge from the pixel-based likelihood. In fact, the partial sky acts as additional noise, shifting the distribution rightward.

The mHL likelihood also exhibits a bias rightward when applied to the cut-sky; however, this effect is less pronounced compared to the HL, as the “extra noise” induced by the partial sky is somewhat marginalized over. The only robust case is the noise-dominated case, where the peak of the posterior remains largely unchanged.

¹¹We emphasize that both the cHL and the single-field HL are always computed adding the corresponding offset.

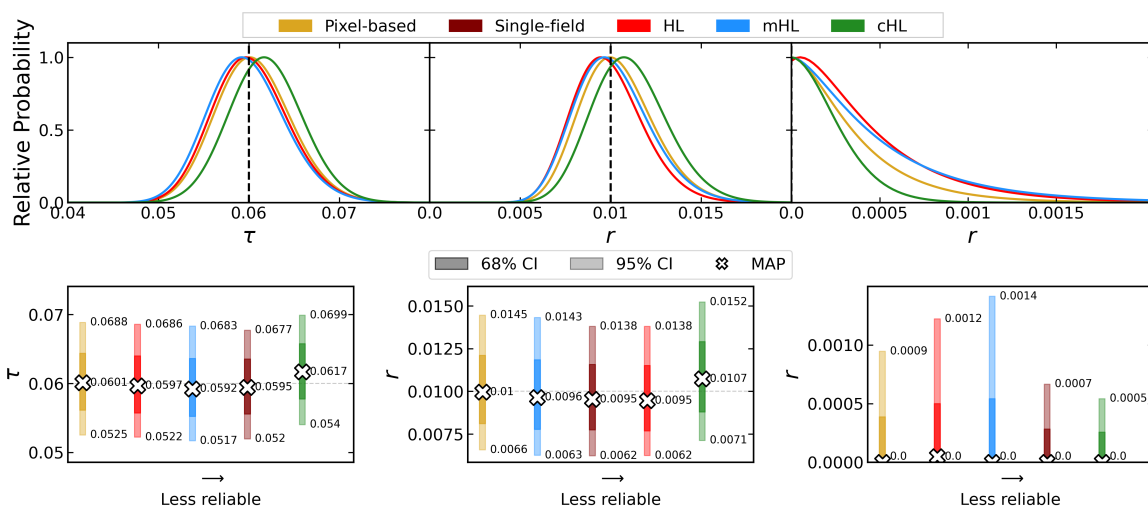


Figure 8. Comparison between pixel-based (gold), HL (red), mHL (blue), and cHL (green) likelihoods with $f_{\text{sky}} = 60\%$ omitting the offset. The columns correspond to signal-dominated, tensor-detection, and noise-dominated cases. The black vertical lines represent the true values of the parameters ($\tau = 0.06$, $r = 0.01$ and $r = 0$). The bottom row provides a quantitative comparison of the posteriors above, together with the single-field HL (maroon). Crosses show MAP estimates, while shaded bars indicate 68% and 95% credible intervals.

The cHL is already discussed in section 4.

In the bottom line of the plot, we can see that the KL divergence ranks the distributions very differently from figure 4. In the signal-dominated case, the most reliable one is the cHL due to the strong bias experienced by the other two. In the tensor-detection case, the mHL and the cHL obtain a very similar KL divergence that slightly favors the mHL. Finally, in the noise-dominated case, although mHL and HL largely overestimate the uncertainty, they obtain a better KL divergence. Still, it is clear that none of the approximations is doing a great job in recovering the PB here.

As shown in figure 8, for a less aggressive sky fraction of $f_{\text{sky}} = 60\%$, neither the HL nor the mHL shows very significant biases. This highlights a key strength of these likelihoods: for $f_{\text{sky}} \geq 60\%$, the power spectrum estimation is stable enough, and there is no need to define an offset to obtain unbiased results.

Here we can also notice that the cHL appears biased, while in the 40% it was not. As mentioned before, the cHL is biased rightward in full-sky, but the mask is causing a pull leftward. For $f_{\text{sky}} = 40\%$, the two effects balance, while in this case with $f_{\text{sky}} = 60\%$, they do not.

Once again the KL divergence ranks differently the various approximations. In the signal-dominated case, the HL gives the best results without doubts. Whether one has to resort to using one of mHL and cHL, they give very similar scores and appear almost symmetric with respect to the true value of τ . In the tensor-detection case, instead, the mHL gives the best agreement to the PB; finally, in the noise-dominated case, the most reliable is the HL, followed by the mHL, and both tend to overestimate the 95% upper bound.




Figure 9. Comparison between pixel-based (gold), HL (red), mHL (blue), and cHL (green) likelihoods with $f_{\text{sky}} = 40\%$ obtained with different offsets. Each panel corresponds to signal-dominated, tensor-detection, and noise-dominated cases. Solid lines refer to the computation presented in section 3, while the dashed lines correspond to the alternative definition introduced here. The black vertical lines indicate the true values of the parameters ($\tau = 0.06$, $r = 0.01$ and $r = 0$).

B.1 Changing offset

In section 3, we anticipated that the offset used in this work follows the prescription of [53]. While this reflects the original definition of the oHL, it is not the only possible way to compute this quantity. For instance, the *Planck* PR4 low- ℓ polarization results adopt a different approach within the so-called LoLLiPoP likelihood [12]. Instead of requiring 99% of spectra to be positive, this method compares the empirical variance of C_ℓ with the expected variance given by the Knox formula [73]

$$\sigma_\ell^2 = \frac{2(C_\ell + O_\ell)^2}{(2\ell + 1)f_{\text{sky}}}, \quad (\text{B.1})$$

and extracts the offset iteratively. We re-compute the posteriors using this alternative definition to test the sensitivity of our results to the offset choice. As an example, we consider the tensor-detection case.

The posteriors for the different likelihoods are shown in figure 9. As discussed in section 4, the mHL provides the best agreement with the PB. Since the offset is applied only to the auto-correlations, which are marginalized over in the mHL (see section 3), its results are largely unaffected by the choice of offset. In contrast, both the HL and cHL show greater sensitivity to the offset computation.

Overall, a larger offset appears to “Gaussianize” the distributions, making them less skewed. Although we do not show it explicitly, the LoLLiPoP-based prescription consistently produces larger offsets than the original method from [53] in the tensor-detection case.

A more detailed study of how different offset definitions affect the analysis is left for future work.

C Approximating the likelihood with a Gaussian

A common approximation for the distribution of C_ℓ values is the Gaussian. In this case, the likelihood has the same form as equation (2.3), where this time the X_ℓ vector is the

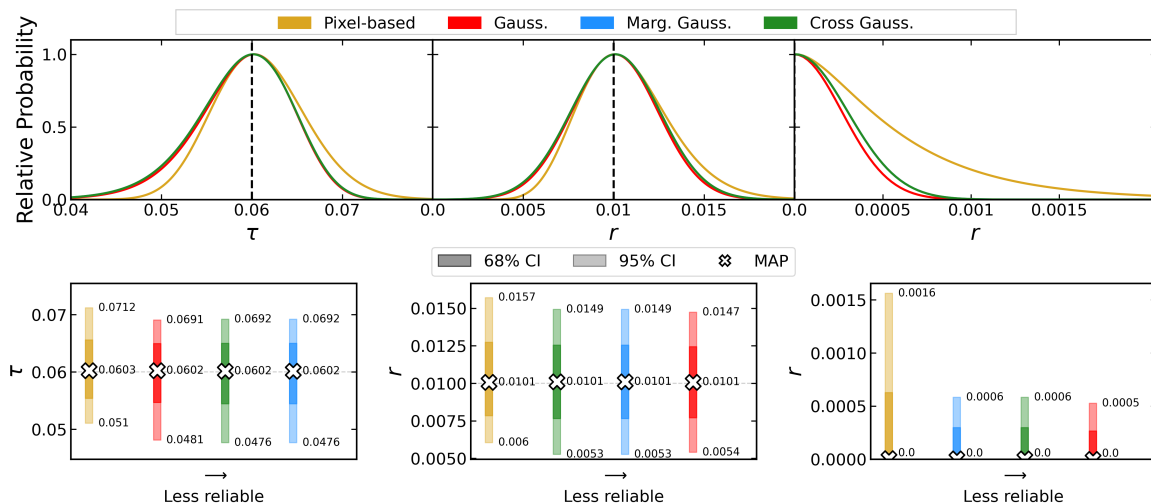


Figure 10. Comparison between pixel-based (gold), Gaussian (red), marginalized Gaussian (blue), and cross-spectra Gaussian (green) likelihoods with $f_{\text{sky}} = 40\%$. The columns correspond to signal-dominated, tensor-detection, and noise-dominated cases. The black vertical lines represent the true values of the parameters ($\tau = 0.06$, $r = 0.01$, and $l = 0$). The bottom row provides a quantitative comparison of the posteriors above: crosses show MAP estimates, while shaded bars indicate 68% and 95% credible intervals.

concatenation of the $n(n+1)/2$ independent $(\hat{C}_\ell - C_\ell)$ components. This approach is justified at high multipoles, where the number of degrees of freedom is large enough for the central limit theorem to apply. However, it is sometimes also used at large scales, where the goal is to approximate the effect of the mask through the covariance.

Figure 10 shows the results. The naming convention here matches the HL-based case; however, the marginalized Gaussian and cross-spectra Gaussian approximations are now identical: without the HL transformation, the two are mathematically the same, and this is confirmed by the results. The difference with the Gaussian lies in the inclusion of auto-spectra.

All posteriors are unbiased in this case. Since the data vector is the direct difference between the observed and model spectra, it is harder to get biased results when using an unbiased estimator like the QML.

However, the Gaussian approximation fails to capture the skewed shape of the true distribution. This leads to underestimated uncertainties in the noise-dominated case, and distorted distributions for τ and r in the signal-dominated and tensor-detection cases, favoring low values.

These results emphasize that non-Gaussian approximations, such as those in the “HL family”, remain necessary for accurate parameter inference on large scales.

References

- [1] U. Seljak and M. Zaldarriaga, *Signature of gravity waves in polarization of the microwave background*, *Phys. Rev. Lett.* **78** (1997) 2054 [[astro-ph/9609169](#)] [[INSPIRE](#)].
- [2] M. Kamionkowski, A. Kosowsky and A. Stebbins, *Statistics of cosmic microwave background polarization*, *Phys. Rev. D* **55** (1997) 7368 [[astro-ph/9611125](#)] [[INSPIRE](#)].

- [3] M. Kamionkowski and E.D. Kovetz, *The Quest for B Modes from Inflationary Gravitational Waves*, *Ann. Rev. Astron. Astrophys.* **54** (2016) 227 [[arXiv:1510.06042](#)] [[INSPIRE](#)].
- [4] W. Hu, *Wandering in the Background: A CMB Explorer*, Ph.D. Thesis, University of California, Berkeley, Berkeley, CA, U.S.A. (1995) [[astro-ph/9508126](#)] [[INSPIRE](#)].
- [5] D.J. Eisenstein, *Dark energy and cosmic sound*, *New Astron. Rev.* **49** (2005) 360 [[INSPIRE](#)].
- [6] D.J. Eisenstein et al., *Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies*, *Astrophys. J.* **633** (2005) 560 [[astro-ph/0501171](#)] [[INSPIRE](#)].
- [7] R. Sachs and A. Wolfe, *Perturbations of a cosmological model and angular variations of the microwave background*, *Astrophys. J.* **147** (1967) 73 [[INSPIRE](#)].
- [8] M. Rees and D. Sciama, *Large scale Density Inhomogeneities in the Universe*, *Nature* **217** (1968) 511 [[INSPIRE](#)].
- [9] E.T. Vishniac, *Reionization and small-scale fluctuations in the microwave background*, *Astrophys. J.* **322** (1987) 597 [[INSPIRE](#)].
- [10] M.C. Guzzetti, N. Bartolo, M. Liguori and S. Matarrese, *Gravitational waves from inflation*, *Riv. Nuovo Cim.* **39** (2016) 399 [[arXiv:1605.01615](#)] [[INSPIRE](#)].
- [11] U. Seljak and M. Zaldarriaga, *A Line of sight integration approach to cosmic microwave background anisotropies*, *Astrophys. J.* **469** (1996) 437 [[astro-ph/9603033](#)] [[INSPIRE](#)].
- [12] M. Tristram et al., *Improved limits on the tensor-to-scalar ratio using BICEP and Planck data*, *Phys. Rev. D* **105** (2022) 083524 [[arXiv:2112.07961](#)] [[INSPIRE](#)].
- [13] G. Galloni et al., *Updated constraints on amplitude and tilt of the tensor primordial spectrum*, *JCAP* **04** (2023) 062 [[arXiv:2208.00188](#)] [[INSPIRE](#)].
- [14] G. Galloni, S. Henrot-Versillé and M. Tristram, *Robust constraints on tensor perturbations from cosmological data: A comparative analysis from Bayesian and frequentist perspectives*, *Phys. Rev. D* **110** (2024) 063511 [[arXiv:2405.04455](#)] [[INSPIRE](#)].
- [15] N. Boggess et al., *The COBE mission — Its design and performance two years after launch*, *Astrophys. J.* **397** (1992) 420 [[INSPIRE](#)].
- [16] C.L. Bennett et al., *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results*, *Astrophys. J. Suppl.* **208** (2013) 20 [[arXiv:1212.5225](#)] [[INSPIRE](#)].
- [17] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [18] PLANCK collaboration, *Planck 2018 results. V. CMB power spectra and likelihoods*, *Astron. Astrophys.* **641** (2020) A5 [[arXiv:1907.12875](#)] [[INSPIRE](#)].
- [19] M. Hazumi et al., *LiteBIRD: A Satellite for the Studies of B-Mode Polarization and Inflation from Cosmic Background Radiation Detection*, *J. Low Temp. Phys.* **194** (2019) 443 [[INSPIRE](#)].
- [20] LITEBIRD collaboration, *Probing Cosmic Inflation with the LiteBIRD Cosmic Microwave Background Polarization Survey*, *Prog. Theor. Exp. Phys.* **2023** (2023) 042F01 [[arXiv:2202.02773](#)] [[INSPIRE](#)].
- [21] BICEP2 collaboration, *Detection of B-Mode Polarization at Degree Angular Scales by BICEP2*, *Phys. Rev. Lett.* **112** (2014) 241101 [[arXiv:1403.3985](#)] [[INSPIRE](#)].

- [22] BICEP/KECK collaboration, *Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season*, *Phys. Rev. Lett.* **127** (2021) 151301 [[arXiv:2110.00483](#)] [[INSPIRE](#)].
- [23] L. Pagano et al., *Reionization optical depth determination from Planck HFI data with ten percent accuracy*, *Astron. Astrophys.* **635** (2020) A99 [[arXiv:1908.09856](#)] [[INSPIRE](#)].
- [24] U. Natale et al., *A novel CMB polarization likelihood package for large angular scales built from combined WMAP and Planck LFI legacy maps*, *Astron. Astrophys.* **644** (2020) A32 [[arXiv:2005.05600](#)] [[INSPIRE](#)].
- [25] M. Lattanzi et al., *On the impact of large angle CMB polarization data on cosmological parameters*, *JCAP* **02** (2017) 041 [[arXiv:1611.01123](#)] [[INSPIRE](#)].
- [26] R. de Belsunce, S. Gratton, W. Coulton and G. Efstathiou, *Inference of the optical depth to reionization from low multipole temperature and polarization Planck data*, *Mon. Not. Roy. Astron. Soc.* **507** (2021) 1072 [[arXiv:2103.14378](#)] [[INSPIRE](#)].
- [27] S. Hamimeche and A. Lewis, *Likelihood Analysis of CMB Temperature and Polarization Power Spectra*, *Phys. Rev. D* **77** (2008) 103013 [[arXiv:0801.0554](#)] [[INSPIRE](#)].
- [28] M. Gerbino et al., *Likelihood methods for CMB experiments*, *Front. Phys.* **8** (2020) 15 [[arXiv:1909.09375](#)] [[INSPIRE](#)].
- [29] M. Remazeilles et al., *LiteBIRD science goals and forecasts. Mapping the hot gas in the Universe*, *JCAP* **12** (2024) 026 [[arXiv:2407.17555](#)] [[INSPIRE](#)].
- [30] LITEBIRD collaboration, *Tensor-to-scalar ratio forecasts for extended LiteBIRD frequency configurations*, *Astron. Astrophys.* **676** (2023) A42 [[arXiv:2302.05228](#)] [[INSPIRE](#)].
- [31] R. Stompor, S. Leach, F. Stivoli and C. Baccigalupi, *Maximum Likelihood algorithm for parametric component separation in CMB experiments*, *Mon. Not. Roy. Astron. Soc.* **392** (2009) 216 [[arXiv:0804.2645](#)] [[INSPIRE](#)].
- [32] V. Stolyarov, M.P. Hobson, M.A.J. Ashdown and A.N. Lasenby, *All-sky component separation for the Planck mission*, *Mon. Not. Roy. Astron. Soc.* **336** (2002) 97 [[astro-ph/0105432](#)] [[INSPIRE](#)].
- [33] H.K. Eriksen et al., *Joint Bayesian component separation and CMB power spectrum estimation*, *Astrophys. J.* **676** (2008) 10 [[arXiv:0709.1058](#)] [[INSPIRE](#)].
- [34] J. Delabrouille, J.-F. Cardoso and G. Patanchon, *Multi-detector multi-component spectral matching and applications for CMB data analysis*, *Mon. Not. Roy. Astron. Soc.* **346** (2003) 1089 [[astro-ph/0211504](#)] [[INSPIRE](#)].
- [35] J. Delabrouille et al., *A full sky, low foreground, high resolution CMB map from WMAP*, *Astron. Astrophys.* **493** (2009) 835 [[arXiv:0807.0773](#)] [[INSPIRE](#)].
- [36] WMAP collaboration, *First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Foreground emission*, *Astrophys. J. Suppl.* **148** (2003) 97 [[astro-ph/0302208](#)] [[INSPIRE](#)].
- [37] A. Carones, M. Migliaccio, D. Marinucci and N. Vittorio, *Analysis of Needlet Internal Linear Combination performance on B-mode data from sub-orbital experiments*, *Astron. Astrophys.* **677** (2023) A147 [[arXiv:2208.12059](#)] [[INSPIRE](#)].
- [38] LITEBIRD collaboration, *Multiclustering needlet ILC for CMB B-mode component separation*, *Mon. Not. Roy. Astron. Soc.* **525** (2023) 3117 [[arXiv:2212.04456](#)] [[INSPIRE](#)].
- [39] G.F. Smoot et al., *Structure in the COBE differential microwave radiometer first year maps*, *Astrophys. J. Lett.* **396** (1992) L1 [[INSPIRE](#)].

- [40] G.F. Smoot, *COBE observations and results*, in the proceedings of the *International Conference on 3 K Cosmology*, Rome, Italy, 5–10 October 1998, *AIP Conf. Proc.* **476** (1999) 1 [[astro-ph/9902027](#)] [[INSPIRE](#)].
- [41] K.M. Gorski et al., *On determining the spectrum of primordial inhomogeneity from the COBE dmr sky maps. 2. Results of two year data analysis*, *Astrophys. J. Lett.* **430** (1994) L89 [[astro-ph/9403067](#)] [[INSPIRE](#)].
- [42] K.M. Gorski, *On Determining the spectrum of primordial inhomogeneity from the Cobe DMR sky maps. 1. Method*, *Astrophys. J. Lett.* **430** (1994) L85 [[astro-ph/9403066](#)] [[INSPIRE](#)].
- [43] K.M. Gorski et al., *Power spectrum of primordial inhomogeneity determined from the four year COBE DMR sky maps*, *Astrophys. J. Lett.* **464** (1996) L11 [[astro-ph/9601063](#)] [[INSPIRE](#)].
- [44] K.M. Górski, *Cosmic microwave background anisotropy in the COBE DMR 4-year sky maps*, in the proceedings of the *31st Rencontres de Moriond: Microwave Background Anisotropies*, Les Arcs, France, 16–23 March 1996, pp. 77–84 [[astro-ph/9701191](#)] [[INSPIRE](#)].
- [45] J.B. Jewell et al., *A Markov Chain Monte Carlo Algorithm for analysis of low signal-to-noise CMB data*, *Astrophys. J.* **697** (2009) 258 [[arXiv:0807.0624](#)] [[INSPIRE](#)].
- [46] J.F. Taylor, M.A.J. Ashdown and M.P. Hobson, *Fast optimal CMB power spectrum estimation with Hamiltonian sampling*, *Mon. Not. Roy. Astron. Soc.* **389** (2008) 1284 [[arXiv:0708.2989](#)] [[INSPIRE](#)].
- [47] H.K. Eriksen et al., *Power spectrum estimation from high-resolution maps by Gibbs sampling*, *Astrophys. J. Suppl.* **155** (2004) 227 [[astro-ph/0407028](#)] [[INSPIRE](#)].
- [48] W.J. Percival and M.L. Brown, *Likelihood methods for the combined analysis of CMB temperature and polarisation power spectra*, *Mon. Not. Roy. Astron. Soc.* **372** (2006) 1104 [[astro-ph/0604547](#)] [[INSPIRE](#)].
- [49] M. Tegmark and A. de Oliveira-Costa, *How to measure CMB polarization power spectra without losing information*, *Phys. Rev. D* **64** (2001) 063001 [[astro-ph/0012120](#)] [[INSPIRE](#)].
- [50] R. de Belsunce, S. Gratton and G. Efstathiou, *B-mode constraints from Planck low-multipole polarization data*, *Mon. Not. Roy. Astron. Soc.* **518** (2022) 3675 [[arXiv:2207.04903](#)] [[INSPIRE](#)].
- [51] J.M. Delouis et al., *SRoll2: an improved mapmaking approach to reduce large-scale systematic effects in the Planck High Frequency Instrument legacy maps*, *Astron. Astrophys.* **629** (2019) A38 [[arXiv:1901.11386](#)] [[INSPIRE](#)].
- [52] PLANCK collaboration, *Planck intermediate results. XLVI. Reduction of large-scale systematic effects in HFI polarization maps and estimation of the reionization optical depth*, *Astron. Astrophys.* **596** (2016) A107 [[arXiv:1605.02985](#)] [[INSPIRE](#)].
- [53] A. Mangilli, S. Plaszczynski and M. Tristram, *Large-scale cosmic microwave background temperature and polarization cross-spectra likelihoods*, *Mon. Not. Roy. Astron. Soc.* **453** (2015) 3174 [[arXiv:1503.01347](#)] [[INSPIRE](#)].
- [54] M. Tristram et al., *Planck constraints on the tensor-to-scalar ratio*, *Astron. Astrophys.* **647** (2021) A128 [[arXiv:2010.01139](#)] [[INSPIRE](#)].
- [55] S. Vanneste, *Constraints on primordial gravitational waves from the large scales CMB data*, Ph.D. Thesis, Université Paris-Saclay (2019) [[INSPIRE](#)].
- [56] BICEP2/KECK and PLANCK collaborations, *Joint Analysis of BICEP2/Keck Array and Planck Data*, *Phys. Rev. Lett.* **114** (2015) 101301 [[arXiv:1502.00612](#)] [[INSPIRE](#)].

- [57] A. Mangilli, N. Bartolo, S. Matarrese and A. Riotto, *The impact of cosmic neutrinos on the gravitational-wave background*, *Phys. Rev. D* **78** (2008) 083517 [[arXiv:0805.3234](#)] [[INSPIRE](#)].
- [58] K. Benabed, J.F. Cardoso, S. Prunet and E. Hivon, *TEASING: a fast and accurate approximation for the low multipole likelihood of the Cosmic Microwave Background temperature*, *Mon. Not. Roy. Astron. Soc.* **400** (2009) 219 [[arXiv:0901.4537](#)] [[INSPIRE](#)].
- [59] E. Rosenberg, S. Gratton and G. Efstathiou, *CMB power spectra and cosmological parameters from Planck PR4 with CamSpec*, *Mon. Not. Roy. Astron. Soc.* **517** (2022) 4620 [[arXiv:2205.10869](#)] [[INSPIRE](#)].
- [60] M. Tristram et al., *Cosmological parameters derived from the final Planck data release (PR4)*, *Astron. Astrophys.* **682** (2024) A37 [[arXiv:2309.10034](#)] [[INSPIRE](#)].
- [61] A. Lewis, A. Challinor and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys. J.* **538** (2000) 473 [[astro-ph/9911177](#)] [[INSPIRE](#)].
- [62] M. Tegmark, *A Method for extracting maximum resolution power spectra from microwave sky maps*, *Mon. Not. Roy. Astron. Soc.* **280** (1996) 299 [[astro-ph/9412064](#)] [[INSPIRE](#)].
- [63] M. Tegmark, *How to measure CMB power spectra without losing information*, *Phys. Rev. D* **55** (1997) 5895 [[astro-ph/9611174](#)] [[INSPIRE](#)].
- [64] J.R. Bond, A.H. Jaffe and L. Knox, *Estimating the power spectrum of the cosmic microwave background*, *Phys. Rev. D* **57** (1998) 2117 [[astro-ph/9708203](#)] [[INSPIRE](#)].
- [65] G. Efstathiou, *Myths and truths concerning estimation of power spectra*, *Mon. Not. Roy. Astron. Soc.* **349** (2004) 603 [[astro-ph/0307515](#)] [[INSPIRE](#)].
- [66] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](#)] [[INSPIRE](#)].
- [67] CMB-S4 collaboration, *CMB-S4 Science Book, First Edition*, [arXiv:1610.02743](#) [[DOI:10.2172/1352047](#)] [[INSPIRE](#)].
- [68] K.M. Górski et al., *HEALPix — A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[INSPIRE](#)].
- [69] J.D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9** (2007) 90 [[INSPIRE](#)].
- [70] P. Virtanen et al., *SciPy 1.0 — Fundamental Algorithms for Scientific Computing in Python*, *Nat. Methods* **17** (2020) 261 [[arXiv:1907.10121](#)] [[INSPIRE](#)].
- [71] S.K. Lam, A. Pitrou and S. Seibert, *Numba: a LLVM-based Python JIT compiler*, in the proceedings of the *International Conference for High Performance Computing, Networking, Storage and Analysis: HPC Transforms*, Austin, U.S.A., 15–20 November 2015 [[DOI:10.1145/2833157.2833162](#)] [[INSPIRE](#)].
- [72] C.R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357 [[arXiv:2006.10256](#)] [[INSPIRE](#)].
- [73] L. Knox, *Determination of inflationary observables by cosmic microwave background anisotropy experiments*, *Phys. Rev. D* **52** (1995) 4307 [[astro-ph/9504054](#)] [[INSPIRE](#)].