

Likelihoods for stochastic gravitational wave background data analysis

Gabriele Franciolini^{1,*} Mauro Pieroni^{1,2,†} Angelo Ricciardone^{3,4,‡} and Joseph D. Romano^{5,§}

¹*CERN, Theoretical Physics Department, Esplanade des Particules 1, Geneva 1211, Switzerland*

²*Instituto de Estructura de la Materia (IEM), CSIC, Serrano 121, 28006 Madrid, Spain*

³*Dipartimento di Fisica “Enrico Fermi”, Università di Pisa,
Largo Bruno Pontecorvo 3, Pisa I-56127, Italy*

⁴*INFN, Sezione di Pisa, Largo Bruno Pontecorvo 3, Pisa I-56127, Italy*

⁵*Department of Physics and Astronomy, University of Texas Rio Grande Valley, One West University
Boulevard, Brownsville, Texas 78520, USA*



(Received 11 June 2025; accepted 22 October 2025; published 12 November 2025)

We present a systematic study of likelihood functions used for stochastic gravitational wave background (SGWB) searches. By dividing the data into many short segments, one customarily takes advantage of the central limit theorem to justify a Gaussian cross-correlation likelihood. We show, with a hierarchy of ever more realistic examples—beginning with a single frequency bin and one detector, and then moving to two and three detectors with white and colored signal and noise—that approximating the exact Whittle likelihood by various Gaussian alternatives can induce systematic biases in the estimation of the SGWB parameters. We derive several approximations for the full likelihood and identify regimes where Gaussianity breaks down. We also discuss the possibility of conditioning the full likelihood on fiducial noise estimates to produce unbiased SGWB parameter estimation. We show that for some segment durations and bandwidths, particularly in space-based and pulsar-timing arrays, the bias can exceed the statistical uncertainty. Our results provide practical guidance for segment choice, likelihood selection, and data-compression strategies to ensure robust SGWB inference in current and next-generation gravitational wave detectors.

DOI: [10.1103/sjtb-gnz5](https://doi.org/10.1103/sjtb-gnz5)

I. INTRODUCTION

The detection of a stochastic gravitational wave background (SGWB) is a primary goal of current and future GW detectors. Arising from the superposition of unresolved astrophysical and/or cosmological sources, the SGWB provides a unique opportunity to investigate the population properties of different astrophysical compact objects [1] and/or probe the physics of the early Universe [2].

Ground-based detectors such as LIGO [3], Virgo [4], and KAGRA [5] (LVK) at their design sensitivity, together with future detectors such as Einstein Telescope (ET) [6,7] and Cosmic Explorer (CE) [8], and space-based interferometers like LISA [9], will significantly improve the sensitivity to the SGWB. A key technique for SGWB searches with

current ground-based detectors consists of cross-correlating the output of multiple detectors (or channels) [10,11]. Assuming (at the lowest order) that the noise is not correlated between different detectors,¹ this procedure gives access to the signal that would otherwise be buried in, or be indistinguishable from, the detector noise [13].

In such searches, the data streams are typically cut into segments, cross-correlated on a segment-by-segment basis, and the result is averaged over segments to output an estimate of the cross-correlation statistic. This procedure greatly simplifies the likelihood function and reduces the computational cost of its evaluation. Indeed, due to the central limit theorem (CLT), when a large number of segments is considered, the estimator for the cross-correlation statistic is expected to follow a Gaussian distribution. Moreover, by segmenting the data stream into short segments, we can more safely assume stationarity to hold within each segment, which allows us to neglect correlations between different frequencies. Despite all these benefits, following this procedure requires some assumptions (e.g., a sufficiently large number of segments is used to construct the estimator to

*Contact author: gabriele.franciolini@cern.ch

†Contact author: mauro.pieroni@csic.es

‡Contact author: angelo.ricciardone@unipi.it

§Contact author: joseph.romano@utrgv.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by CERN.

¹See [12] for a study of the impact of correlated seismic noise on the reconstruction of the SGWB with ET.

TABLE I. This table summarize all the likelihoods considered in this work, highlighting the assumptions, the level of compression in frequency f or number of segments N_s , the noise assumptions, and the estimation accuracy, together with some implications. For each likelihood, we report the corresponding equations in the main text. We use the symbols “✓, ✗” to indicate the property of the likelihoods, which sometime can depend on some features specified in the most right column. For more details, see main text.

Likelihood	N_s -compressed	N_f -compressed	Noise agnostic	Unbiased	Coverage	Further info
$\mathcal{L}^W(\theta D_k^s)$, Eq. (11)	✗	✗	✗	✓	✓	...
$\mathcal{L}^G(\theta D_k)$, Eq. (14)	✓	✗	✓	✗	✓	...
$\mathcal{L}^{\text{WMA}^P}(\theta D_k)$, Eq. (37)	✓	Depends	✗	✓	✓	See [14,15], for N_f compression
$\mathcal{L}^W(\theta D_k)$, Eq. (64)	✓	✗	✗	✓	✓	...
$\mathcal{L}^F(\theta D_k)$, Eq. (65)	✓	✓	✗	✓	✓	Gaussian uncertainties in θ
$\mathcal{L}^{W,\text{red}}(\theta \bar{D}_k)$, Eq. (66)	✓	✗	✓	Depends	✗ (under)	Bias depends on $\bar{P}_{II,k}$ estimates
$\mathcal{L}^{G,\text{cc}}(\theta \bar{D}_k)$, Eq. (67)	✓	✗	✓	Depends	✗ (under)	Bias depends on $\bar{P}_{II,k}$ estimates
$\mathcal{L}^{G,\text{opt}}(\theta \bar{D}_k)$, Eq. (77)	✓	✓	✓	Depends	✗ (under)	Bias depends on $\bar{P}_{II,k}$ estimates

achieve Gaussianity) whose validity and impact must be checked thoroughly. This becomes more important for future detectors, where, in addition to the SGWB, numerous resolvable sources will be present in the data streams, and they will need to be characterized and subtracted.

In this paper, we aim to characterize the biases introduced by the Gaussian likelihood assumption and identify the regimes in which it remains a valid approximation for different GW detectors. We also propose alternative likelihoods that may offer a more accurate description of the underlying statistics. We provide a simple quantitative comparison between these approaches in order to improve the robustness of the SGWB parameter estimation and guide the analysis strategies of future GW detectors. To highlight the results of our work, in Table I, we summarize all the likelihoods considered in this work, reporting the degree of compression in frequency or number of segments, the assumptions concerning noise characterization, the expected accuracy and precision, and mention some additional features as needed.

The paper is organized as follows. We begin in Sec. II by setting up the full frequency-domain framework, introducing the covariance matrices, data-compression steps, and the exact Whittle likelihood. Then, in Sec. III, we consider a simple “toy” example consisting of a single detector and single frequency bin, to illustrate how the Gaussian approximation can break down when the number of segments is small, and to motivate other likelihoods. Section IV generalizes the analysis to multiple detectors (including both white and colored signal and noise models), deriving both the exact and reduced Whittle forms, the Fisher (Gaussian) approximation, and optimal estimators for cross-correlation. In Sec. V, we compute statistical uncertainties and systematic biases across segmentation choices for realistic detectors like current LVK, future ET/CE, LISA, and pulsar timing arrays detectors. Moreover, we briefly discuss the situation in anisotropic SGWB searches. Finally, in Sec. VI we draw our conclusions and outline some future directions for our work. Finally, in the Appendix, we perform further tests on the likelihoods presented in this work for different

signal regimes, number of frequencies, and number of segments.

II. SETTING UP THE FRAMEWORK

Schematically, we can model the GW detector output in the frequency domain by writing the data stream in the channel (or detector) $I = 1, 2, \dots, N_{\text{det}}$ as $\tilde{d}_{I,k} = \tilde{s}_{I,k} + \tilde{n}_{I,k}$, where k denotes the discrete frequency bin $f_k \equiv k/T$, with $k = 0, 1, \dots, N/2 - 1$ (non-negative frequencies) and $T \equiv N\Delta t$ is the duration of the observation. Here, $\tilde{s}_{I,k}$ represents the GW signal, while $\tilde{n}_{I,k}$ refers to the detector/channel noise component. A common assumption in GW data analysis is that the signal and noise components are zero-mean Gaussian processes, uncorrelated with each other; i.e., $\langle \tilde{s} \rangle = \langle \tilde{n} \rangle = \langle \tilde{s} \tilde{n} \rangle = 0$. In addition, the noise is often taken to be stationary and uncorrelated across distinct frequency bins.

Under these Gaussian and stationarity assumptions, the joint probability distribution of the data $\tilde{d} \equiv \{\tilde{d}_{I,k}\}$ can be written explicitly. In particular, one can write the log-likelihood function for a set of frequency-domain measurements as²

²In this paper, we will denote the likelihood function as $\mathcal{L}(\text{params}|\text{data})$, interpreting the model parameters as the variable quantities and the data as fixed (corresponding to a particular measurement). Alternatively, we could have used the notation $p(\text{data}|\text{params})$ for the likelihood function, interpreting the data as random variables, conditioned on a fixed set of parameter values. This latter form of the likelihood function is important if one needs to simulate data to test a data analysis pipeline, for example. Since we will be primarily interested in maximum-likelihood estimates and posterior distributions for the parameters, we will ignore additive terms to the log-likelihood function that are independent of the parameter values, even if these terms depend on the data. But beware that these data-dependent terms are important for evidence and Bayes factor calculations if different likelihood functions are used to describe the different models. If the same likelihood function is used for the different models, then the data-dependent terms will cancel out of the Bayes factor calculation.

$$-2 \ln \mathcal{L}^W(\theta|\tilde{d}) = \sum_{k=0}^{N/2-1} \left\{ 2 \ln \det [2\pi C_k(\theta)] + \sum_{I,J=1}^{N_{\text{det}}} \tilde{d}_{I,k}^* [C_k^{-1}(\theta)]_{IJ} \tilde{d}_{J,k} \right\}. \quad (1)$$

Here $C_k^{-1}(\theta)$ denotes the inverse of the $N_{\text{det}} \times N_{\text{det}}$ covariance matrix $C_k(\theta)$ (for a specific discrete frequency bin f_k), whose components can be expressed as

$$C_{IJ,k}(\theta) = \frac{T}{4} P_{IJ,k}(\theta), \quad (2)$$

where $P_{IJ,k}(\theta)$ are the (one-sided) auto and cross-power spectral densities (PSDs) of the data

$$\langle \tilde{d}_{I,k} \rangle = 0, \quad \langle \tilde{d}_{I,k} \tilde{d}_{J,k'}^* \rangle = \frac{T}{2} \delta_{kk'} P_{IJ,k}(\theta). \quad (3)$$

These power spectral densities can be split into signal and noise components P_h and P_n

$$P_{IJ,k}(\theta) \equiv P_{hIJ,k}(\theta_h) + P_{nIJ,k}(\theta_n), \quad (4)$$

which depend on some signal and noise parameters, θ_h, θ_n , respectively. In the literature, Eq. (1) is also known as the Whittle likelihood [16].

It is often convenient to write the log-likelihood in a slightly more compact form by defining the quadratic data combination matrices³

$$D_{IJ,k} \equiv \frac{2}{T} \text{Re}[\tilde{d}_{I,k} \tilde{d}_{J,k}^*]. \quad (5)$$

In terms of $D_{IJ,k}$ and $P_{IJ,k}$, we have

$$-2 \ln \mathcal{L}^W(\theta|D) = 2 \sum_k \{ \ln \det [P_k(\theta)] + \text{Tr}[P_k^{-1}(\theta) D_k] \}, \quad (6)$$

where we have ignored any additive terms that are independent of θ , noting that $D_{IJ,k}$ is a sufficient statistic for this likelihood.⁴ Equation (6) is a useful starting point for data compression, which is possible, e.g., if the data can be combined over frequency or segments (as discussed in some simple examples below).

When analyzing data from current ground-based detectors (i.e., LVK), each data channel corresponds to a

³The real part of the product $\tilde{d}_{I,k} \tilde{d}_{J,k}^*$ enters the trace in Eq. (6), since $P_{IJ,k}$ is a real symmetric matrix on account of the summation in Eq. (1) ranging over only non-negative frequencies.

⁴Since the data matrices $D_{IJ,k}$ follow a Wishart distribution [17], the log-likelihood function given in Eq. (6) is also called the Wishart likelihood.

different detector. Since the noises for different detectors are expected to be uncorrelated,⁵ we have

$$P_{nIJ,k}(\theta_n) = \delta_{IJ} P_{nI,k}(\theta_n), \quad (7)$$

where $P_{nI,k}$ denotes the noise power spectral density of the I th detector evaluated at frequency bin f_k . On the other hand, the gravitational wave background is correlated across multiple detectors, so that

$$P_{hIJ,k}(\theta_h) = \Gamma_{IJ}(f_k) S_h(f_k; \theta_h). \quad (8)$$

Here, $\Gamma_{IJ}(f)$ denotes the overlap reduction function [11,20,21] (which encodes the reduction in sensitivity when correlating physically separated and misaligned detectors) and

$$S_h(f) \equiv \frac{3H_0^2 \Omega_{\text{gw}}(f)}{2\pi^2 f^3}, \quad (9)$$

is the (one-sided) strain power spectral density for the GWB, with $\Omega_{\text{gw}}(f)$ denoting the fractional energy-density spectrum in gravitational waves [11].

When data are acquired over an extended period, the instrumental noise properties can change significantly over time, so the assumption of strict stationarity cannot be used. For this reason, the data stream is typically divided into a number N_s of much shorter segments, and the likelihood is expressed as a product,

$$\ln \mathcal{L}^W(\theta|D) = \sum_{s=1}^{N_s} \ln \mathcal{L}^W(\theta^s|D^s), \quad (10)$$

where s indexes the segments. Note that the parameter vector $\theta^s \equiv (\theta_h, \theta_n^s)$ now carries a segment index, since the noise parameters will generally change from segment to segment. This increases the total number of noise parameters that would need to be inferred from the data. Thus, the most general expression for the log-likelihood function involving multiple channels (detectors), time segments, and frequencies is

⁵Noise correlation can affect GW detectors in different ways; in the triangular configuration, colocated detectors share correlated seismic noise—both Rayleigh and body waves—along with local gravitational-field fluctuations (Newtonian noise). These effects can persist up to tens of Hz, posing a challenge to measuring an isotropic SGWB. Separated interferometers are affected by broadband environmental noise, originating from Schumann resonances, which leads to coherent magnetic-field fluctuations [18,19]. In this paper, we disregard both these correlated noise effects.

$$-2 \ln \mathcal{L}^W(\theta|D) = 2 \sum_s \sum_k \{ \ln \det[P_k(\theta^s)] + \text{Tr}[P_k^{-1}(\theta^s) D_k^s] \}, \quad (11)$$

where the quadratic data combination given in Eq. (5) now depends on the data segment, $D_{IJ,k}^s$.

For current ground-based detectors, the signal is typically too weak to overcome the detector noise, a scenario commonly referred to as the *weak-signal* (or *noise-dominated*) regime. In this regime, the individual detector noise power spectra $P_{nI,k}^s$ are well-approximated by the detector auto-power $P_{II,k}^s$, and may be replaced in the likelihood by (fixed) fiducial estimates $\bar{P}_{II,k}^s$ obtained, e.g., from spline fits of the quadratic data combinations. Under the assumption of uncorrelated noise, the GW signal appears uniquely in the cross-correlated data ($I \neq J$). Hence, searches for GWBs using ground-based detectors focus primarily on correlating data from pairs of distinct detectors [10,11].

It is then customary to define for each segment s the following cross-correlation estimator of the fractional energy-density spectrum,

$$\hat{\Omega}_{IJ,k}^s \equiv \left(\frac{2\pi^2}{3H_0^2} \right) \frac{f_k^3}{\Gamma_{IJ}(f_k)} D_{IJ,k}^s, \quad (12)$$

where $I \neq J$. The numerical factors on the rhs are chosen so that $D_{IJ,k}^s / \Gamma_{IJ}(f_k)$ is an estimator of $S_h(f_k)$ [see Eqs. (3), (5), and (8)], while $f_k^3 (2\pi^2 / 3H_0^2)$ converts strain power $S_h(f_k)$ to energy-density spectrum units [see Eq. (9)]. In the weak-signal regime, one can further combine these segment-dependent estimators by weighting them by the inverse of the (approximate) variances,

$$(\sigma^2)_{IJ,k}^s \approx \left(\frac{2\pi^2}{3H_0^2} \frac{f_k^3}{\Gamma_{IJ}(f_k)} \right)^2 \bar{P}_{II,k}^s \bar{P}_{JJ,k}^s. \quad (13)$$

If the number of segments is sufficiently large, this averaging over segments yields data combinations $\hat{\Omega}_{IJ,k}$ that (according to the central limit theorem) are approximately Gaussian distributed. In this case,

$$-2 \ln \mathcal{L}^G(\theta_h|\hat{\Omega}) = \sum_{I < J} \sum_k \left\{ \ln(2\pi \sigma_{IJ,k}^2) + \frac{[\hat{\Omega}_{IJ,k} - \Omega_{\text{gw}}(f_k, \theta_h)]^2}{\sigma_{IJ,k}^2} \right\}, \quad (14)$$

where $\Omega_{\text{gw}}(f, \theta_h)$ is the model spectrum and

$$1/\sigma_{IJ,k}^2 \equiv \sum_s 1/(\sigma^2)_{IJ,k}^s \quad (15)$$

is the variance of the estimator $\hat{\Omega}_{IJ,k}$. Note that the above summation is over only distinct pairs of detectors $I \neq J$,

and that the only parameters we are trying to infer from the data are those for the gravitational-wave background θ_h .

Following this procedure significantly reduces the amount of data, leading to a significant reduction of the computational cost, which, in turn, speeds up the analysis. However, with the enhanced sensitivity expected from future GW detectors, the precision of these approximations [i.e., the assumption of Gaussianity (14) and the weak-signal approximation (13)] might become insufficient due to improved measurement accuracy. In the following sections, we explore the conditions under which this approximate likelihood yields reliable maximum likelihood (ML) estimates, as well as scenarios in which ML biases are likely to become evident.

III. ONE-DETECTOR CASE

We start by presenting a very simple one-detector analysis, treating each frequency bin independently, which requires no assumptions on the frequency dependence of the signal and noise. We postpone the more complex scenario, where several frequency bins are combined, to the following sections.

A. A simple example

As a very simple example, consider N_s segments of stationary Fourier-domain data $\tilde{d}_{I,k}^s$ from a single detector I , and focus attention on a single frequency bin f_k . We will assume that the segments are statistically equivalent and independent of one another. Since I and k are fixed for this example, we will drop them as subscripts and write $\tilde{d}_s$ for the data and $D_s \equiv 2|\tilde{d}_s|^2/T$ for the quadratic data combination Eq. (5), where $s = 1, 2, \dots, N_s$. The frequency domain data satisfy the quadratic expectation values,

$$\langle \tilde{d}_s \rangle = 0, \quad \langle \tilde{d}_s \tilde{d}_{s'}^* \rangle = \frac{T}{2} \delta_{ss'} P, \quad (16)$$

where our assumption regarding stationarity across segments means that the value P of the power spectrum at frequency f_k is independent of the segment s . For this example, we will take P as the parameter we want to infer from the data. Note that, at this level, we do not try to separate signal and noise, meaning P represents the total autocorrelated power.

The normalized probability distribution for the N_s data samples $\tilde{d} \equiv \{\tilde{d}_s | s = 1, 2, \dots, N_s\}$ is then

$$p(\tilde{d}|P) \equiv \left(\frac{4}{2\pi TP} \right)^{N_s} \exp \left[-\frac{1}{2} \sum_s \frac{4|\tilde{d}_s|^2}{TP} \right]. \quad (17)$$

Note that the argument of the exponential can be written more simply in terms of the quadratic data combination

$$D \equiv \frac{1}{N_s} \sum_s D_s \equiv \frac{1}{N_s} \sum_s \frac{2}{T} |\tilde{d}_s|^2, \quad (18)$$

where we compressed D_s over segments. Indeed, one can show that D is the ML estimator of P , which we will denote as

$$\hat{P} \equiv D, \quad (19)$$

and has mean and variance

$$\langle D \rangle = P, \quad \langle D^2 \rangle - \langle D \rangle^2 = \frac{P^2}{N_s}. \quad (20)$$

The latter can be calculated directly from Eq. (16) using Isserlis's (or Wick's) theorem [22] to expand the expectation value of a product of four $\tilde{d}_s$'s. D is obviously non-Gaussian, being proportional to a sum of squared Gaussian random variables. In particular, it follows a scaled chi-squared distribution, with scaling factor $P/2N_s$ and $2N_s$ real degrees of freedom.

By inspection of Eqs. (17) and (18), one sees that the quadratic data combination D is a sufficient statistic for Eq. (17). Starting from Eq. (17), and making a change of variables from the original dataset $\{\tilde{d}_s | s = 1, \dots, N_s\}$ to D , one obtains the properly normalized conditional probability distribution for D given P as

$$p^W(D|P) = \frac{N_s^{N_s} D^{N_s-1} e^{-N_s D/P}}{\Gamma(N_s) P^{N_s}}, \quad (21)$$

where $\Gamma(x)$ is the Gamma function. The corresponding log-likelihood function (ignoring additive terms that are independent of the parameter P) is simply,

$$-2 \ln \mathcal{L}^W(P|D) = 2N_s (\ln P + D/P). \quad (22)$$

A plot of $\mathcal{L}^W(P|D)$ is given in Fig. 1 (black curve) together with several other approximate likelihood functions, which we discuss in the next two subsections.

B. Approximate Gaussian likelihood functions

Using the mean and variance of D given in Eq. (20), one can construct different *approximate Gaussian-distributed* likelihood functions, depending on how we want to treat the variance term. Following the notation in [23], we consider the following ‘‘Gaussian-like’’ approximations to the Whittle likelihood function. Plots of these likelihood functions are given in Fig. 1, all normalized to have unit area. Thus, one can think of these plots as *posterior probability distributions* for P , calculated assuming a flat prior probability distribution.

- (i) The *symmetric* Gaussian approximate likelihood function is defined by




FIG. 1. Plots of $\mathcal{L}^W(P|D)$ (black) together with several approximate likelihood functions (all normalized to have unit area), plotted as a function of the unknown parameter P for fixed data D . For this particular example, $D = 1.5$ and the number of data segments is $N_s = 4$.

$$-2 \ln \mathcal{L}^S(P|D) = N_s \left(\frac{D - P}{D} \right)^2, \quad (23)$$

where the theoretical variance P^2/N_s is replaced by the data-dependent expression D^2/N_s . This likelihood peaks at D (which is the ML estimate of P from the full Whittle likelihood) and is symmetric about around the peak.

- (ii) Alternatively, one can define the *quadratic* Gaussian approximate likelihood function, whose logarithm is quadratic in the data combination D :

$$-2 \ln \mathcal{L}^Q(P|D) = N_s \left(\frac{D - P}{P} \right)^2. \quad (24)$$

This approximate likelihood uses the theoretical expression for the variance but does not include the standard log-determinant contribution from the variance (which is used below). This likelihood again peaks at D , but it is no longer symmetric around the peak.

- (iii) Finally, we can consider the standard Gaussian likelihood function, which uses theoretical expressions for the variance and includes the log-determinant term (hence the ‘D’ superscript for *determinant*):

$$-2 \ln \mathcal{L}^D(P|D) = N_s \left(\frac{D - P}{P} \right)^2 + 2 \ln P. \quad (25)$$

The shape of this likelihood agrees more closely with the Whittle likelihood than those associated with the other two Gaussian approximations, but its peak does

not agree with D . By maximizing $\mathcal{L}^D(P|D)$ with respect to P , one can show that $\mathcal{L}^D$ is peaked at

$$\begin{aligned}\hat{P}^D &\equiv D(\sqrt{1 + 4/N_s} - 1)N_s/2 \\ &= D[1 - 1/N_s + \mathcal{O}(1/N_s^2)].\end{aligned}\quad (26)$$

This difference is indicated in Fig. 1 by the dashed orange vertical line. Since $\langle D \rangle = P$, it follows that

$$\langle \hat{P}^D \rangle = P[1 - 1/N_s + \mathcal{O}(1/N_s^2)], \quad (27)$$

which means that $\hat{P}^D$ (the ML estimator of P associated with the Gaussian likelihood $\mathcal{L}^D(P|D)$) is a *biased* estimate of the parameter P . For this case, the fractional bias factor is given by

$$\mathcal{B} \equiv \frac{\langle \hat{P}^D \rangle - P}{P} \simeq -\frac{1}{N_s}, \quad (28)$$

which shows that the Gaussian-based estimator systematically underestimates the true variance by a factor $1/N_s$. This bias will clearly diminish in the limit where $N_s \gg 1$.

We further illustrate this with a slightly more complicated example. Consider data dominated by a single SGWB component parametrized by P , measured over N_f frequency bins and averaged across N_s segments. Under the Gaussian approximation, one expects that by sampling over a total number of $N_f \times N_s$ independent realizations, the fractional statistical uncertainty of the ML estimator $\hat{P}$ of the full Whittle likelihood around its true value P will scale as

$$\mathcal{S} \equiv \Delta \hat{P} \equiv \left| \frac{\hat{P} - P}{P} \right| \simeq \frac{1}{\sqrt{N_f N_s}}. \quad (29)$$

The number of frequencies can be roughly estimated from the duration of the observational run T and the maximum frequency $f_{\max}$ as $N_f \sim f_{\max} T$. Note that, depending on the values of N_f and N_s , the resulting statistical uncertainty $\mathcal{S}$ may be smaller than the intrinsic bias $\mathcal{B}$. In such a case, the bias can become relevant and affect the reconstruction. Furthermore, in realistic scenarios where detector noise is present and we are not in a signal-dominated regime, a practical criterion to estimate $\mathcal{S}$ would be to count only frequency bins where the noise PSD is much smaller than the signal power, in order to improve the estimator. In the opposite case, the presence of the noise would increase the level of uncertainty.

C. A CMB-inspired alternative likelihood

The Gaussian-approximate likelihood functions described in the previous subsection are quadratic in the fractional difference,

$$\delta \equiv \frac{D - P}{P} \Leftrightarrow D = P(1 + \delta). \quad (30)$$

It turns out that by considering a linear combination of the quadratic Gaussian likelihood $\ln \mathcal{L}^Q(P|D)$ defined in Eq. (24), and a *log-normal* likelihood

$$-2 \ln \mathcal{L}^{\text{LN}}(P|D) = N_s \left[\ln \left(\frac{D}{P} \right) \right]^2, \quad (31)$$

we can construct an approximate likelihood that agrees with the full Whittle likelihood $\mathcal{L}^W(P|D)$ given in Eq. (22) to third order in δ .

To demonstrate this, we follow the discussion presented in [24] in the context of CMB physics. We start by rewriting the Whittle likelihood $\mathcal{L}^W(P|D)$ by adding to the rhs of Eq. (22) a couple of terms that are independent of the parameter P and hence do not change the shape of $\mathcal{L}^W(P|D)$ viewed as a function of P :

$$-2 \ln \mathcal{L}^W(P|D) \simeq 2N_s \left[\ln \left(\frac{P}{D} \right) + \frac{D}{P} - 1 \right]. \quad (32)$$

In terms of δ , we have

$$\begin{aligned}-2 \ln \mathcal{L}^W(P|D) &= 2N_s [-\ln(1 + \delta) + (1 + \delta) - 1] \\ &= N_s \left[\delta^2 - \frac{2}{3} \delta^3 + \mathcal{O}(\delta^4) \right].\end{aligned}\quad (33)$$

Similarly, we can express $\mathcal{L}^Q$ and $\mathcal{L}^{\text{LN}}$ in terms of δ :

$$-2 \ln \mathcal{L}^Q(P|D) = N_s \delta^2 \quad (34)$$

and

$$-2 \ln \mathcal{L}^{\text{LN}}(P|D) = N_s [\ln(1 + \delta)]^2 = N_s [\delta^2 - \delta^3 + \mathcal{O}(\delta^4)]. \quad (35)$$

In this way, the full Whittle likelihood can be written as

$$\ln \mathcal{L}^W(P|D) = \ln \mathcal{L}^{\text{WMAP}}(P|D) + \mathcal{O}(\delta^4), \quad (36)$$

where the WMAP likelihood approximation is defined as

$$\ln \mathcal{L}^{\text{WMAP}}(P|D) \equiv \frac{1}{3} \ln \mathcal{L}^Q(P|D) + \frac{2}{3} \ln \mathcal{L}^{\text{LN}}(P|D). \quad (37)$$

As shown in Sec. III A, the fractional deviation δ has a typical magnitude that scales as $\mathcal{O}(N_s^{-1})$ once the data are averaged over N_s (statistically independent) segments. The next correction—arising from the quartic term in the expansion—is $\mathcal{O}(N_s^{-2})$ and therefore subdominant for practical GWB searches. As evident from inspection of Fig. 1 (compare gray and black curves), the WMAP

likelihood approximation agrees very closely with that of the full Whittle likelihood function.

IV. MORE DETECTORS

We now consider a series of more complicated examples, which involve multiple detectors having either white or colored signal and noise models. For all cases, we will assume that the data are taken over a total observation period T_{obs} , and are split into N_s segments of duration $T \equiv T_{\text{obs}}/N_s$. We will assume that the data are stationary across segments, which means that the parameters θ describing the signal and noise are not segment dependent. This assumption will allow us to combine the data over segments as was done in Sec. III. We will also assume that individual frequency bins are statistically independent, which is a good approximation when the data are correlated over time scales much shorter than the segment duration T .

If we additionally assume that the signal and noise are white with variances $\sigma_h^2, \sigma_{n_I}^2$, for $I = 1, 2, \dots, N_{\text{det}}$, then the $C_k(\theta)$ matrices are also independent of frequency since the one-sided power spectral densities are given by

$$P_{n_{I,k}} = 2\sigma_{n_I}^2 \Delta t, \quad S_h(f_k) = 2\sigma_h^2 \Delta t, \quad (38)$$

where $\Delta t \equiv T/N$, with N being the number of time domain samples in a single segment. (Recall that the number of non-negative frequency bins is $N_f = N/2$.) Also, assuming that the overlap function $\Gamma_{IJ}(f)$ is independent of frequency and equals unity when $I = J$, then

$$C_k(\theta) = \frac{T}{4} P, \quad (39)$$

where the components P_{IJ} of the (constant) P matrix are given by

$$P_{IJ} = \begin{bmatrix} P_1 & \Gamma_{12} S_h & \cdots & \Gamma_{1N_{\text{det}}} S_h \\ \Gamma_{12} S_h & P_2 & \cdots & \Gamma_{2N_{\text{det}}} S_h \\ \vdots & \vdots & \ddots & \vdots \\ \Gamma_{1N_{\text{det}}} S_h & \Gamma_{2N_{\text{det}}} S_h & \cdots & P_{N_{\text{det}}} \end{bmatrix}, \quad (40)$$

with

$$P_I \equiv P_{II} \equiv P_{n_I} + \Gamma_{II} S_h = 2(\sigma_{n_I}^2 + \sigma_h^2) \Delta t. \quad (41)$$

For this white signal and noise case, we will take $\theta \equiv (P_1, \dots, P_{N_{\text{det}}}, S_h)$ as the set of parameters that we want to infer from the data; they are simply related to σ_h^2 and $\sigma_{n_I}^2$ using Eqs. (38) and (41).

Making the above simplifying substitutions into Eq. (11), we obtain

$$-2 \ln \mathcal{L}^W(\theta|D) = 2N_s N_f (\ln \det(P) + \text{Tr}[P^{-1}D]), \quad (42)$$

where the quadratic data matrix D is

$$D_{IJ} \equiv \frac{1}{N_s N_f} \sum_s \sum_k D_{IJ,k}^s, \quad (43)$$

and where we have ignored terms on the rhs of Eq. (42) that are independent of the parameters θ . Note that Eq. (42) is a *compressed* version of Eq. (11), since the data combination D_{IJ} above has been compressed over both segments and frequency bins given the assumptions of stationarity across segments and white signal and noise models. These equations will be the starting point for the analyses in Secs. IV A and IV B. All frequency-dependent considerations are postponed to Sec. IV C.

A. 2-detectors, white signal, white noise

For the case of 2-detectors with white signal and white noise, the log-likelihood function given in Eq. (42) takes the following explicit form:

$$\begin{aligned} & -2 \ln \mathcal{L}^W(\theta|D) \\ &= 2N_s N_f \left[\ln(P_1 P_2 - \Gamma_{12}^2 S_h^2) \right. \\ & \quad \left. + \left(1 - \frac{\Gamma_{12}^2 S_h^2}{P_1 P_2} \right)^{-1} \left(\frac{D_{11}}{P_1} + \frac{D_{22}}{P_2} - 2 \frac{D_{12} \Gamma_{12} S_h}{P_1 P_2} \right) \right]. \quad (44) \end{aligned}$$

If we maximize this log-likelihood function with respect to the parameters $\theta \equiv (P_1, P_2, S_h)$, we obtain the following ML estimates:

$$\hat{P}_1 = D_{11}, \quad \hat{P}_2 = D_{22}, \quad \hat{S}_h = \frac{D_{12}}{\Gamma_{12}}. \quad (45)$$

These estimators are quadratic in the Fourier domain data $\tilde{d}_{I,k}$. In general, ML estimators of the parameters need not be quadratic, as we shall see for the 3-detector examples studied in Secs. IV B and IV C.

The variance of these estimators can be computed starting from the quadratic expectation values of the Fourier-domain data

$$\langle \tilde{d}_{I,k}^s \tilde{d}_{J,k'}^{s'*} \rangle = \frac{T}{2} \delta_{ss'} \delta_{kk'} P_{IJ}. \quad (46)$$

Higher-order moments required for the variance calculation can then be evaluated by applying Isserlis's theorem [22], which allows one to express fourth-order expectation values of $\tilde{d}$ in terms of products of second-order moments. Alternatively, one can obtain the variances from the inverse of the *Fisher information matrix* (FIM) defined as

$$F_{ij} \equiv \left\langle -\frac{\partial^2 \ln \mathcal{L}^W(\theta|D)}{\partial \theta_i \partial \theta_j} \right\rangle, \quad (47)$$

which, for this particular example, takes the form

$$F_{ij} = N_s N_f \text{Tr} \left[P^{-1} \frac{\partial P}{\partial \theta_i} P^{-1} \frac{\partial P}{\partial \theta_j} \right]. \quad (48)$$

The variance for parameter θ_i is given by

$$\sigma_{\theta_i}^2 = (F^{-1})_{ii}. \quad (49)$$

More generally, the covariance between pairs of parameters θ_i and θ_j is given by $\text{Cov}(\theta_i, \theta_j) = (F^{-1})_{ij}$. For this example,

$$(F^{-1})_{ij} = \frac{1}{N_s N_f} \begin{bmatrix} P_1^2 & \Gamma_{12}^2 S_h^2 & P_1 S_h \\ \Gamma_{12}^2 S_h^2 & P_2^2 & P_2 S_h \\ P_1 S_h & P_2 S_h & \frac{1}{2} \left(S_h^2 + \frac{P_1 P_2}{\Gamma_{12}^2} \right) \end{bmatrix}. \quad (50)$$

Note that the variance of $\hat{S}_h$ is valid here even if the signal power is large compared to that of the noise.

By definition, the FIM and its inverse involve the model parameters, whose values are unknown. In practice, one evaluates the components of these matrices, F_{ij} and $(F^{-1})_{ij}$, at the ML estimates⁶ given in Eq. (45). We will denote these data-dependent expressions with “hats”, i.e., $\hat{F}_{ij} \equiv F_{ij}|_{\theta=\hat{\theta}}$ and similarly for $(\hat{F}^{-1})_{ij}$.

Using the ML estimates of the parameters and the associated covariance matrix $\hat{F}^{-1}$, we can construct a multivariate-Gaussian approximate to the Whittle likelihood function Eq. (42), which we will denote as $\mathcal{L}^F(\theta|D)$, with ‘F’ for Fisher. Ignoring terms that are independent of the parameters θ_i , we define

$$-2 \ln \mathcal{L}^F(\theta|D) = \sum_{i,j} (\hat{\theta}_i - \theta_i) \hat{F}_{ij} (\hat{\theta}_j - \theta_j), \quad (51)$$

where the indices i, j run over the total number of parameters, and the estimators $\hat{\theta}_i$ are expressed in terms of the data combinations D_{IJ} by Eq. (45).⁷ This is a symmetric Gaussian approximation for each parameter similar in spirit to Eq. (23). By construction, the ML estimators of the parameters and the magnitude of their uncertainties for this approximate likelihood function agree with those from the full Whittle likelihood Eq. (44).

⁶For more complicated examples where the ML estimators of the full Whittle likelihood are not given by simple quadratic combinations of the data, like in Eq. (45), one can calculate the ML estimators using the iterative FIM method described in Sec. IV B, see Eq. (56).

⁷See previous footnote regarding calculating $\hat{\theta}_i$ for the more general scenario.

It is also possible to construct a “reduced” non-Gaussian likelihood starting with $\ln \mathcal{L}^W(\theta|D)$ in Eq. (44), focusing attention on just the GWB parameter S_h . For this case, we replace the parameters P_1 , P_2 and data combinations $2D_{11}/T$, $2D_{22}/T$ by fixed quantities $\bar{P}_1$, $\bar{P}_2$, which are not inferred from the analysis. Rather, they can be either *fiducial estimates* of the autopower constructed, e.g., using spline fits of the detector output, or they may be estimated from the data segment itself.⁸ These replacements give rise to *reduced* power spectrum and quadratic data combination matrices, $\bar{P}_{IJ}$ and $\bar{D}_{IJ}$, which differ from P_{IJ} and D_{IJ} defined in Eqs. (40) and (43) only along the diagonal. With these replacements, we can define the *reduced Whittle* likelihood function for an arbitrary number of detectors, assuming here that the signal and noise are both white

$$\begin{aligned} & -2 \ln \mathcal{L}^{W,\text{red}}(S_h | \bar{D}; \bar{P}_1, \dots, \bar{P}_{N_{\text{det}}}) \\ & = 2N_s N_f [\ln \det(\bar{P}) + \text{Tr}(\bar{P}^{-1} \bar{D})]. \end{aligned} \quad (52)$$

For two detectors, the above likelihood depends only on the cross-correlated data D_{12} and is conditioned on $\bar{P}_1$ and $\bar{P}_2$. Explicitly,

$$\begin{aligned} & -2 \ln \mathcal{L}^{W,\text{red}}(S_h | D_{12}; \bar{P}_1, \bar{P}_2) \\ & = 2N_s N_f \left[\ln(\bar{P}_1 \bar{P}_2 - \Gamma_{12}^2 S_h^2) \right. \\ & \quad \left. + 2 \left(1 - \frac{\Gamma_{12}^2 S_h^2}{\bar{P}_1 \bar{P}_2} \right)^{-1} \left(1 - \frac{D_{12} \Gamma_{12} S_h}{\bar{P}_1 \bar{P}_2} \right) \right]. \end{aligned} \quad (53)$$

One can show that this reduced likelihood function is peaked at D_{12}/Γ_{12} , which is the ML estimator of S_h of the Whittle likelihood function as indicated in Eq. (45). Since this likelihood function is conditioned on $\bar{P}_1$ and $\bar{P}_2$, the uncertainty in the recovery of S_h will be underestimated compared to that obtained by marginalizing over P_1 and P_2 .

Alternatively, a reduced likelihood function, which is Gaussian in the cross-correlated data D_{12} , can be constructed by again using fiducial estimates $\bar{P}_1$, $\bar{P}_2$ for the parameters P_1 , P_2 . One defines (again up to additive terms independent of S_h)

$$-2 \ln \mathcal{L}^{G,\text{cc}}(S_h | D_{12}; \bar{P}_1, \bar{P}_2) = \frac{(\hat{S}_{h,\text{cc}} - S_h)^2}{\sigma_{S_{h,\text{cc}}}^2}, \quad (54)$$

where

⁸We do not recommend using data-dependent estimators of the autopower that are constructed from the analysis segment, since they may lead to biases in parameter estimates, as shown for example in Fig. 5. Using power spectrum estimates from neighboring segments avoids this problem, see, e.g., [25].




FIG. 2. Comparison of the posterior distributions for S_h using the Whittle, Fisher, reduced Whittle, and cross-correlation approximate likelihood functions for the two-detector, white signal, white noise example. The plots for $\mathcal{L}^W$ and $\mathcal{L}^F$ were marginalized over the parameters P_1 and P_2 . The marginalization was done via simple numerical integration over a grid of P_1 and P_2 values, assuming flat priors for all the parameters. For this analysis, $N_s = 4$ and $N_f = 32$ corresponding to a total compression factor of $N_s N_f = 128$.

$$\hat{S}_{h,cc} \equiv \frac{D_{12}}{\Gamma_{12}}, \quad \sigma_{\hat{S}_{h,cc}}^2 \equiv \frac{1}{2N_s N_f} \frac{\bar{P}_1 \bar{P}_2}{\Gamma_{12}^2}. \quad (55)$$

This is a symmetric Gaussian likelihood function, quadratic in the data $\tilde{d}$, depending only on the cross-correlations D_{12} . The corresponding variance is valid only for sufficiently weak GWB signals. This is evident from inspecting the 33-component of F^{-1} in Eq. (50), which shows that the signal-signal contribution S_h^2 is missing from $\sigma_{\hat{S}_{h,cc}}^2$.

A comparison of the likelihood functions $\mathcal{L}^W$, $\mathcal{L}^F$, $\mathcal{L}^{W,red}$, $\mathcal{L}^{G,cc}$, plotted as functions of the GW power spectrum parameter S_h , is given in Fig. 2. The likelihood functions $\mathcal{L}^W$ and $\mathcal{L}^F$ have been marginalized over the other parameters P_1 , P_2 , while $\mathcal{L}^{W,red}$, $\mathcal{L}^{G,cc}$ are conditioned on fiducial estimates $\bar{P}_1$, $\bar{P}_2$ of these parameters. All the plots have been normalized to have unit area, so they can be interpreted as posterior distributions for S_h , under the assumption of flat priors. The marginalization over P_1 , P_2 for $\mathcal{L}^W$ and $\mathcal{L}^F$ have been done via simple numerical integration as opposed to MCMC methods, which will be used in Sec. IV B for the 3-detector examples. For these plots, the simulated data $\tilde{d}_{I,k}^s$ had $N_s = 4$ and $N_f = 32$ corresponding to a total compression factor of $N_s N_f = 128$.

B. 3-detectors, white signal, white noise

We now extend the analyses of the previous subsection to the case of 3-detectors, again assuming stationary, white

signal, and white noise power spectra. The relevant parameters for this problem are the constant power spectral densities $\theta \equiv (P_1, P_2, P_3, S_h)$. Since the following analysis extends easily to an arbitrary number of detectors having this simple signal + noise model, we will write down the general form of these expressions whenever possible.

The general Whittle likelihood function $\mathcal{L}^W(\theta|D)$ for the multidetector, stationary, white signal, white noise case was already written down in Eq. (42). Similarly, the multivariate-Gaussian likelihood function $\mathcal{L}^F(\theta|D)$ for this case has the same form as that given in Eq. (51). However, for three or more detectors, the ML estimators $\hat{\theta}_i$, $i = 1, 2, \dots, N_{\text{det}} + 1$ of the Whittle likelihood in Eq. (42) are no longer guaranteed to be given by simple extensions of the quadratic estimators Eq. (45). Instead, the ML estimators of the parameters $\theta = (P_1, \dots, P_{N_{\text{det}}}, S_h)$ can be obtained by iterating the equations,

$$\begin{aligned} \Delta \hat{\theta}_i^{(n)} &= \sum_j [F^{-1}(\theta)]_{ij} \frac{\partial \ln \mathcal{L}^W(\theta|D)}{\partial \theta_j} \Big|_{\theta=\hat{\theta}^{(n)}}, \\ \hat{\theta}_i^{(n+1)} &= \hat{\theta}_i^{(n)} + \Delta \hat{\theta}_i^{(n)}, \end{aligned} \quad (56)$$

for $n = 0, 1, \dots$ until the absolute value of the deviations between successive iterations $|\Delta \hat{\theta}_i^{(n)}|$ drop below some prescribed tolerance. The definition of the FIM F_{ij} and its expression for the case of stationary, white signal, and white noise are given in equations Eqs. (47) and (48). An explicit expression for the partial derivative of the log-likelihood function entering Eq. (56) for this simple signal + noise model is

$$\frac{\partial \ln \mathcal{L}^W(\theta|D)}{\partial \theta_i} = N_s N_f \text{Tr} \left[P^{-1} \frac{\partial P}{\partial \theta_i} P^{-1} (D - P) \right]. \quad (57)$$

Typically, one starts the iteration by assuming that the parameter values are given by the quadratic estimators⁹

$$\hat{P}_I^{(0)} = D_{II} \quad \text{for } I = 1, \dots, N_{\text{det}}, \quad (58)$$

$$\hat{S}_h^{(0)} \equiv \frac{1}{N_{\text{pairs}}} \sum_{I=1}^{N_{\text{det}}} \sum_{J>I}^{N_{\text{det}}} \frac{D_{IJ}}{\Gamma_{IJ}}, \quad (59)$$

where $N_{\text{pairs}} \equiv N_{\text{det}}(N_{\text{det}} - 1)/2$ is the number of distinct detector pairs. The final outcome of the iteration gives

⁹Equation (59) for $\hat{S}_h^{(0)}$ is a simple cross-correlation estimator that is most appropriate when the total autocorrelated power for all detectors is the same. If that is not the case, one can construct a smaller variance estimator by forming a $1/\sigma^2$ weighting of D_{IJ}/Γ_{IJ} with their variances (which we would either need to know *a priori* or approximate using fiducial estimates). A discussion of optimal (i.e., minimum-variance unbiased) cross-correlation estimators of S_h is discussed in Sec. IV C; see in particular Eqs. (75) and (76).




FIG. 3. Similar to Fig. 2, but for the 3-detector, stationary, white signal, white noise example. For this 3-detector example, the marginalizations are done using MCMC methods, assuming flat priors for all the parameters P_1 , P_2 , P_3 , S_h . Note that for 3-detectors, the ML estimate of the parameter S_h does not agree with the quadratic estimator for S_h .

nonquadratic estimators $\hat{\theta}_i$ of the parameters θ_i . An illustration of the difference between the quadratic and iterated ML estimators of the parameter S_h for a single realization of data can be found in Fig. 3 (gray and black dashed vertical lines). But even for this simple 3-detector example, we were not able to find an analytic expression for the ML estimators.

In addition to the Whittle and Fisher likelihood functions, one can also write down the “reduced” Whittle likelihood approximation $\mathcal{L}^{W, \text{red}}(S_h | \bar{D}; \bar{P}_1, \dots, \bar{P}_{N_{\text{det}}})$, see Eq. (52). As discussed in Sec. IV A, the reduced Whittle likelihood is conditioned on the parameters P_I , assuming fixed, fiducial values $\bar{P}_I$ for these parameters. Recall that the matrices $\bar{P}_{IJ}$ and $\bar{D}_{IJ}$ differ from P_{IJ} and D_{IJ} only by having $\bar{P}_I \equiv \bar{P}_{II}$ as their diagonal elements.

Finally, the “reduced” cross-correlation Gaussian likelihood approximation $\mathcal{L}^{G, \text{cc}}(S_h | \bar{D}; \bar{P}_1, \dots, \bar{P}_{N_{\text{det}}})$ has exactly the same form as $\mathcal{L}^{G, \text{cc}}(S_h | D_{12}; \bar{P}_1, \bar{P}_2)$ given in Eq. (54), with $\hat{S}_{h, \text{cc}}$ and $\sigma_{\hat{S}_{h, \text{cc}}}^2$ given by

$$\hat{S}_{h, \text{cc}} \equiv \frac{1}{N_{\text{pairs}}} \sum_{I < J} \frac{D_{IJ}}{\Gamma_{IJ}}, \quad \sigma_{\hat{S}_{h, \text{cc}}}^2 \equiv \frac{1}{2N_s N_f} \frac{1}{N_{\text{pairs}}^2} \sum_{I < J} \frac{\bar{P}_I \bar{P}_J}{\Gamma_{IJ}^2}. \quad (60)$$

As noted earlier, the above variance is strictly valid only in the weak-signal limit.

A comparison of the likelihood functions $\mathcal{L}^W$, $\mathcal{L}^F$, $\mathcal{L}^{W, \text{red}}$, $\mathcal{L}^{G, \text{cc}}$, for the 3-detector, stationary, white signal, and white noise case, is given in Fig. 3. As we did in Fig. 2, we plot the marginalized likelihoods as functions of the single GW parameter S_h , normalizing the distributions so that they correspond to posterior distributions for S_h assuming flat priors for all the parameters. The difference between the values of the quadratic and iterated ML estimators of S_h is apparent for this particular realization of the data. By construction, the peak of the cross-correlation Gaussian approximation using fiducial estimates $\bar{P}_I$ of the autopower agrees with the quadratic cross-correlation estimator $\hat{S}_{h, \text{cc}}$ and not $\hat{S}_h$, where the latter is obtained by iterating equations Eq. (56) for the full Whittle likelihood. Nonetheless, since both $\hat{S}_{h, \text{cc}}$ and $\hat{S}_h$ satisfy $\langle \hat{S}_{h, \text{cc}} \rangle = S_h$ and $\langle \hat{S}_h \rangle = S_h$, there is no intrinsic bias in these analyses. However, as mentioned earlier, the reduced analyses will underestimate the uncertainty in the recovery of S_h , since they are conditioned on the autopower $\bar{P}_I$, and the Gaussian cross-correlation analysis explicitly uses a fiducial estimate of the variance, which is strictly only valid in the weak-signal regime.

C. 3-detectors, nonwhite case

We now extend the 3-detector example to allow for nonwhite signal and noise models. Since the data are no longer white, we cannot naively compress the data $\tilde{d}_{I,k}^s$ over frequencies f_k as we did in Eq. (43). But, given our assumption of stationarity, we can still compress over segments:

$$D_{IJ,k} \equiv \frac{1}{N_s} \sum_s \frac{2}{T} \text{Re}[\tilde{d}_{I,k}^s \tilde{d}_{J,k}^{s*}]. \quad (61)$$

For simplicity, we will assume that both the signal and noise power spectra $S_h(f_k)$ and $P_{n_I}(f_k)$ have known shapes $H(f_k)$ and $N_I(f_k)$, i.e.,

$$S_h(f_k) = S_h H(f_k), \quad P_{n_I}(f_k) = P_{n_I} N_I(f_k), \quad (62)$$

where $H(f_k)$ and $N_I(f_k)$ are normalized to unity at some reference frequency. For illustrative purposes, in our analysis, we modeled both the signal and noise spectra as simple parabolic functions of $x \equiv \log(f/f_*)$ with $f_* = \sqrt{f_{\text{min}} f_{\text{max}}}$ for signal and noise spectra. In Fig. 4 we plot the resulting power spectral densities for these injected signal and noise models. The parameters that we will need to infer for this example are then the overall amplitudes of the power spectra, $\theta \equiv (P_{n_1}, P_{n_2}, P_{n_3}, S_h)$.

If we assume that the noise is uncorrelated across detectors, then the expected value of $D_{IJ,k}$ is




FIG. 4. Power spectral density for every pair of detectors (channels) in a three-detector network. The three left panels show autospectra, while the rightmost panel displays all cross-spectra. Each colored line corresponds to one channel combination, with dashed lines indicating the injected signal model (black), noise (red), and their sum (purple). For this analysis we use $N_s = 3$ and $N_f = 10^3$.

$$\langle D_{IJ,k} \rangle = P_{IJ,k}, \quad \text{where } P_{IJ,k} \equiv \delta_{IJ} P_{n_I} N_{I,k} + S_h \Gamma_{IJ} H_k. \quad (63)$$

In terms of the matrices D_k and P_k (with components $D_{IJ,k}$ and $P_{IJ,k}$ defined above), the four likelihood functions we defined in previous subsections become

$$-2 \ln \mathcal{L}^W(\theta | D_k) = 2N_s \sum_k [\ln \det(P_k) + \text{Tr}(P_k^{-1} D_k)], \quad (64)$$

$$-2 \ln \mathcal{L}^F(\theta | D_k) = \sum_{i,j} (\hat{\theta}_i - \theta_i) \hat{F}_{ij} (\hat{\theta}_j - \theta_j), \quad (65)$$

$$\begin{aligned} -2 \ln \mathcal{L}^{W,\text{red}}(S_h | \bar{D}_k; \bar{P}_{11,k}, \dots) \\ = 2N_s \sum_k [\ln \det(\bar{P}_k) + \text{Tr}(\bar{P}_k^{-1} \bar{D}_k)], \end{aligned} \quad (66)$$

$$-2 \ln \mathcal{L}^{G,\text{cc}}(S_h | \bar{D}_k; \bar{P}_{11,k}, \dots) = \sum_k \frac{(\hat{S}_{h,k} - S_h)^2}{\sigma_{\hat{S}_{h,k}}^2}, \quad (67)$$

which are similar in form to what we had in previous subsections. But now the data combination and power spectra depend on frequency, and there is a summation over the discrete frequencies. The expressions for $\bar{D}_k$ and $\bar{P}_k$ are also similar to what we had earlier; namely, we replace the diagonal elements of $D_{IJ,k}$ and $P_{IJ,k}$ with estimates of the total autopower $\bar{P}_{II,k}$. This replacement means that, for the reduced analyses, we no longer need to infer the amplitudes of the noise power spectra P_{n_I} . In addition, for the reduced Gaussian cross-correlation likelihood, we have

$$\hat{S}_{h,k} \equiv \frac{1}{N_{\text{pairs}}} \sum_{I < J} \frac{D_{IJ,k}}{\Gamma_{IJ}}, \quad \sigma_{\hat{S}_{h,k}}^2 \equiv \frac{1}{2N_s} \frac{1}{N_{\text{pairs}}^2} \sum_{I < J} \frac{\bar{P}_{II,k} \bar{P}_{JJ,k}}{\Gamma_{IJ}^2}. \quad (68)$$

which are similar to Eq. (60), but using the frequency-dependent quadratic data combinations $D_{IJ,k}$ for $I \neq J$ and fiducial autopower estimates $\bar{P}_{II,k}$. Notice that depending on the choice for $\bar{P}_{II,k}$, i.e., some fiducial estimate smoothed over neighboring frequency bins vs a purely frequency-by-frequency data-driven estimate, the reduced likelihood $\mathcal{L}^{W,\text{red}}$ and the cross-correlation likelihood $\mathcal{L}^{G,\text{cc}}$ can return unbiased/biased results. An explicit example of this behavior is shown explicitly in Fig. 5 for $\mathcal{L}^{G,\text{cc}}$ (note the difference between the green and dashed-green distribution).

For the reduced Gaussian cross-correlation likelihood, it is also possible to further compress the data over frequency, using information about the expected shape of the signal and noise models. In particular, we can form an *optimal* (i.e., minimum-variance,¹⁰ unbiased) estimate of S_h , which is summed over frequencies and distinct detector pairs using information about the mean and covariance of $D_{IJ,k}$, which we calculate below.

But, to first understand optimal filtering in the context of a simpler example, suppose we have N samples of random data x_i with mean and covariance,

$$\langle x_i \rangle = A \Gamma_i, \quad \langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle = C_{ij}. \quad (69)$$

We note that the x_i do not have to be Gaussian distributed. Let us further assume that the mean “shape” Γ_i and covariance matrix C_{ij} are known *a priori*, or that they can be estimated from neighboring data segments or via some other means, e.g., making spline fits to the data. Then the optimal quadratic estimator of the amplitude A is given by

¹⁰This estimator of S_h has minimum variance when restricted to the class of quadratic cross-correlation estimators. The unrestricted minimum-variance, unbiased estimator of S_h comes from the full Whittle likelihood (it saturates the Cramer-Rao bound [26]), using information from autocorrelations. As mentioned previously, the Whittle likelihood ML estimator is non-quadratic and requires iteration for detector networks having three or more detectors.




FIG. 5. Comparison of the posterior distributions for S_h using the Whittle, Fisher, reduced Whittle, and cross-correlation approximate likelihood functions for the three-detector, colored noise and signal example. For this analysis, we have injected a signal with $S_h = 0.1$, noise amplitude $P_{n_1} = 1.25$, $P_{n_2} = 1.08$ and $P_{n_3} = 1.25$ for each detector, and we have set $N_s = 3$ and $N_f = 1000$.

$$\hat{A} = \frac{\sum_{i,j} \Gamma_i (C^{-1})_{ij} x_j}{\sum_{k,l} \Gamma_k (C^{-1})_{kl} \Gamma_l}, \quad (70)$$

which can be thought of as an “inverse-noise weighted matched filter” statistic, normalized such that $\langle \hat{A} \rangle = A$. In practice, the inverse noise weighting “whitens” the data. The variance of the estimator $\hat{A}$ is given by

$$\sigma_{\hat{A}}^2 = \frac{1}{\sum_{k,l} \Gamma_k (C^{-1})_{kl} \Gamma_l}. \quad (71)$$

Note that Eq. (70) reduces to the standard $1/\sigma_i^2$ weighting if the covariance matrix C is diagonal.

We already calculated the mean value of $D_{IJ,k}$ in Eq. (63). The covariance matrix of $D_{IJ,k}$ is given by

$$\begin{aligned} C_{IJK,KLk'} &\equiv \langle D_{IJ,k} D_{KL,k'} \rangle - \langle D_{IJ,k} \rangle \langle D_{KL,k'} \rangle \\ &= \frac{1}{2N_s} \delta_{kk'} (P_{IK,k} P_{JL,k} + P_{IL,k} P_{JK,k}), \\ &\equiv C_{IJ,KL} \delta_{kk'}, \end{aligned} \quad (72)$$

where we used Isserlis’s theorem [22] to expand the fourth-order expectation values of products of the Gaussian data $\tilde{d}_{I,k}^s$ which enter in $D_{IJ,k}$.

Since the noise is uncorrelated, we can construct an optimal quadratic estimator for the amplitude S_h of the GW power spectrum using only the cross-correlated data combinations $D_{IJ,k}$ where $I \neq J$. This reduces the range of indices for the covariance matrix $C_{IJ,KL}$. Explicitly, for the case of 3 detectors, we can write

$$C_{IJK,KLk'} = \frac{1}{2N_s} \delta_{kk'} \begin{bmatrix} P_{11,k} P_{22,k} + P_{12,k}^2 & P_{11,k} P_{23,k} + P_{12,k} P_{13,k} & P_{12,k} P_{23,k} + P_{13,k} P_{22,k} \\ P_{11,k} P_{23,k} + P_{12,k} P_{13,k} & P_{11,k} P_{33,k} + P_{13,k}^2 & P_{12,k} P_{33,k} + P_{13,k} P_{23,k} \\ P_{12,k} P_{23,k} + P_{13,k} P_{22,k} & P_{12,k} P_{33,k} + P_{13,k} P_{23,k} & P_{22,k} P_{33,k} + P_{23,k}^2 \end{bmatrix}. \quad (73)$$

We will further restrict attention to the case of a “reduced analysis” where we condition on the fiducial estimates of $P_{II,k}$ and also assume that the GWB is weak compared to the noise. Then

$$\begin{aligned} \bar{C}_{IJK,KLk'} &= \frac{1}{2N_s} \delta_{kk'} \begin{bmatrix} \bar{P}_{11,k} \bar{P}_{22,k} & 0 & 0 \\ 0 & \bar{P}_{11,k} \bar{P}_{33,k} & 0 \\ 0 & 0 & \bar{P}_{22,k} \bar{P}_{33,k} \end{bmatrix} \\ &= \frac{1}{2N_s} \delta_{kk'} \delta_{IK} \delta_{JL} \bar{P}_{II,k} \bar{P}_{JJ,k}, \end{aligned} \quad (74)$$

which allows us to drop all cross-spectra terms $P_{IJ,k}$ (for $I \neq J$) relative to the autopower spectra $P_{II,k}$. We put a bar

over the covariance matrix to indicate it is conditioned on the power spectrum estimates.

In terms of this reduced covariance matrix, the optimal cross-correlation estimator of S_h is given by the data combination

$$\hat{S}_{h,\text{opt}} = \mathcal{N} \left[\sum_{I < J} \sum_k \frac{\Gamma_{IJ} H_k D_{IJ,k}}{\bar{P}_{II,k} \bar{P}_{JJ,k}} \right], \quad (75)$$

with variance

$$\sigma_{\hat{S}_{h,\text{opt}}}^2 \equiv \frac{1}{2N_s} \mathcal{N} \equiv \frac{1}{2N_s} \left[\sum_{I < J} \sum_k \frac{\Gamma_{IJ}^2 H_k^2}{\bar{P}_{II,k} \bar{P}_{JJ,k}} \right]^{-1}, \quad (76)$$

which is valid in the weak-signal limit. In terms of this optimal estimator, the reduced Gaussian cross-correlation likelihood function, denoted $\mathcal{L}^{\text{G,opt}}$ is then

$$-2 \ln \mathcal{L}^{\text{G,opt}}(S_h | \bar{D}_k; \bar{P}_{11,k}, \dots) = \frac{(\hat{S}_{h,\text{opt}} - S_h)^2}{\sigma_{\hat{S}_{h,\text{opt}}}^2}, \quad (77)$$

with $\hat{S}_{h,\text{opt}}$ and $\sigma_{\hat{S}_{h,\text{opt}}}^2$ given above.

In Fig. 5 we plot a comparison of the likelihood functions $\mathcal{L}^{\text{W}}$, $\mathcal{L}^{\text{F}}$, $\mathcal{L}^{\text{W,red}}$, $\mathcal{L}^{\text{G,cc}}$, and $\mathcal{L}^{\text{G,opt}}$, for the 3-detector, stationary, nonwhite signal, and nonwhite noise case. We have injected a signal with $S_h = 0.1$ and noise amplitude for each detector sampled from a normal distribution $\mathcal{N}(0, 1.25)$. As before, we plot the marginalized likelihoods as functions of the single GW parameter S_h , normalizing the distributions so that they correspond to posterior distributions for S_h assuming flat priors for all the parameters. For the reduced analyses, we obtain unbiased/biased results depending on whether $\bar{P}_{11,k}$ is given by some fiducial estimate smoothed over neighboring frequency bins or uses a purely frequency-by-frequency data-driven estimate. This is illustrated by the green and dashed-green distributions for the $\mathcal{L}^{\text{G,cc}}$ likelihood analyses in Fig. 5.

We conclude this section by referring the reader to the Appendix, where we perform a series of tests to assess the robustness of the different likelihoods in different signal, frequencies, and segments regimes. In particular, we show the posterior distribution in the strong signal regime and varying the number of frequencies/segments, and we produce some P-P plots to study the coverage accuracy of the different likelihoods.

V. REALISTIC GW DETECTORS

In the data analysis of current and future detectors, the choice of the number of time segments used in the data processing plays a crucial role in determining the accuracy and reliability of the likelihood function employed to detect and characterize the SGWB. Segmenting the continuous data stream into N_s segments ensures the validity of the stationarity and Gaussianity assumptions within each segment, significantly simplifying the analysis and eventually avoiding biases in the parameter estimation. In this section, we estimate the expected fractional statistical uncertainty $\mathcal{S}$ with present and future GW detectors and compare it to the expected bias $\mathcal{B}$ obtained through the Gaussian-based estimator defined in Eq. (28). The main results of this section are summarized in Table II.

A technical point to mention is that, typically, the analysis uses the full GW-sensitive band (up to a given f_{max}) to estimate signal and noise simultaneously. In practice, different analyses focus on narrower frequency (or time–frequency) regions, where the target signal is most probable (i.e., where the signal-to-noise ratio is more likely to be larger) or stationary. As a result, the true effective bandwidth used to estimate signal and noise is narrower than the entire detector range in most searches, which can alter both statistical uncertainties and systematic biases compared to the idealized, full-band estimates. To try to address this point, we also provide uncertainties using the FIM derived as in Eq. (47), i.e., taking derivatives of the detector network Whittle likelihood with respect to the model parameters and weighting them by the inverse of the signal and noise PSDs. For reference, we inject a flat, i.e., frequency-independent, SGWB with amplitude $\Omega_{\text{GW}} = 5.8 \times 10^{-9}$, corresponding to the 95% credible upper limit reported by the LVK

TABLE II. Representative number of segments N_s , independent frequency bins N_f , minimal statistical uncertainty $\mathcal{S}$, and potential systematic bias $\mathcal{B}$ for various detector configurations. In the last column, we report the relative uncertainty on the signal amplitude, estimated using the FIM, assuming a flat SGWB injected with an amplitude matching the current LVK upper bound [27], i.e., $\Omega_{\text{GW}} = 5.8 \times 10^{-9}$. Ground-based figures assume $f_{\text{max}} = 10^3$ Hz; space-based and PTA entries follow the specifications discussed in the text.

Detector config. (seg duration T)	T_{obs}	N_s	N_f	$\mathcal{S}$	$\mathcal{B}$	$\mathcal{S}$ from FIM
Ground-based detectors						
LVK (4 s)	18 mo	1.2×10^7	4.0×10^3	4.6×10^{-6}	8.0×10^{-8}	1.4×10^{-2}
LVK (10 m)		7.8×10^4	6.0×10^5		1.3×10^{-5}	
ET (CE) (4 s)	12 mo	7.9×10^6	4.0×10^3	5.6×10^{-6}	1.3×10^{-7}	$4.6(2.5) \times 10^{-5}$
(10 m)		5.3×10^4	6.0×10^5		1.9×10^{-5}	
Space-based detectors						
LISA (11.5 d)	4 yr	1.3×10^2	5.0×10^5	1.3×10^{-4}	8.0×10^{-3}	4.2×10^{-4}
Taiji (11.5 d)	6 yr	1.9×10^2	5.0×10^5	1.0×10^{-4}	5.2×10^{-3}	2.8×10^{-4}
Galactic-scale detectors						
PTAs current (single segment)	15 yr	1	~ 10	4×10^{-2}	...	1.3×10^{-1}
PTAs future (single segment)	15 yr	1	$\sim 10^2$	7×10^{-3}	...	3×10^{-2}

Collaboration [27]. These choices maximize the amplitude of the possible signal in the ground-based detector's frequency band, while indicating a representative benchmark for experiments in different frequency ranges, and thus give an optimistic estimate of the sensitivity. Clearly, signals with a lower amplitude would correspond to larger parameter uncertainties. The noise curves adopted for the individual detectors are the design-sensitivity strain power spectral density $P_n(f)$ for Advanced LIGO [28], the ET-D noise configuration for the Einstein Telescope [29], and the noise-equivalent strain spectral density $S_n(f)$ for the LISA detector [30]. For current and future PTA datasets, we model pulsar noises based on the currently observed pulsars and generate future catalogs with similar properties (see [31] for more details).

A. Possible biases in isotropic searches

Ground-based detectors, such as the current LVK network and the future ET and CE detectors, operate in the frequency band between 1 Hz and several kHz. The LVK detectors are interferometers with arm lengths of a few km designed to detect GWs produced by compact binary mergers (mainly black holes and neutron stars), supernovae, and, at design sensitivity, also the SGWB of astrophysical origin [27]. The current event rate is approximately one merger every 2 to 3 days [32]. LVK is now in its fourth run (O4), which is supposed to last until fall 2025 for a total of $T_{\text{obs}} \simeq 18$ months (excluding commissioning breaks).

Typical data segments in LVK analysis last from 4 seconds to several (10) minutes, depending on the nature of the targeted signal, where the noise is (roughly speaking) stationary. Assuming $f_{\text{max}} = 10^3$ Hz and $T_{\text{obs}} = 18$ months, we report in Table II the values for the statistical uncertainties $\mathcal{S}$, bias $\mathcal{B}$, and relative uncertainty on the signal amplitude for two typical segmentation time values. We notice that with 4 s segments, the systematic bias is below the statistical uncertainty, and more than three orders of magnitude below the FIM bound. Therefore, the measurement is statistics-limited. With $T = 10$ min, the bias increases but remains well below the statistical uncertainty estimated with the FIM.

ET [7] is a proposed third-generation GW detector with a triangular configuration of three 10-kilometer interferometers (see [29] for other possible configurations) situated underground to minimize the impact of seismic noise at low frequency. Similarly, CE will be a third-generation L-shaped interferometer with 40-km arm length [33]. They will have at least one order of magnitude better sensitivity than second-generation detectors and will extend the observable frequency range down to 1 Hz, which will allow for the detection of signals from a wide variety of astrophysical sources, including the SGWB. ET and CE are expected to significantly enhance detection rates of compact binary coalescences due to their advanced sensitivity.

Current forecasts estimate that ET could detect approximately 10^5 BBH mergers per year, along with a similar number of BNS mergers [7,34]. This large increase in event rates, compared to current detectors, will provide a more comprehensive understanding of the population of compact binary systems [6,7].

Even if a proper segmentation strategy is not yet defined for ET/CE, as long as all the noise contributions are not well characterized, assuming a similar duration as LVK (from a few seconds up to several minutes) will lead to a large number of N_s per year. The choice of such a large number of shorter segments helps to ensure that the noise can be treated as stationary and Gaussian within each segment. For ET and CE, a 4-sec segmentation yields a bias that is two orders of magnitude smaller than the statistical uncertainty and is entirely negligible compared to the FIM bound. However, when 10-min segments are considered, the bias exceeds the statistical error and becomes a dominant contribution to the total uncertainty, indicating a possible transition to a bias-limited regime.

LISA consists of a constellation of three spacecraft in an equilateral configuration, with arm lengths of 2.5 million kilometers [30]. This configuration is designed to detect low-frequency GW signals, such as those emitted by supermassive black hole binaries (SMBHBs) and extreme mass-ratio inspirals (EMRIs), which cannot be observed by ground-based detectors. LISA will operate in the milli-Hz frequency regime (from 3×10^{-5} Hz up to 0.5 Hz) for a baseline mission of at least 4.5 years, potentially extending up to 10 years. Motivated by the cadence of scheduled satellite operations, the data are commonly divided into relatively long segments. LISA will have an effective observation time of $T_{\text{obs}} = 4$ years, and the segment duration will be $T = 11.5$ days each. This corresponds to a frequency resolution $\Delta f = 10^{-6}$ Hz. In this way, we get $N_s \simeq 126$ and $N_f \simeq 5 \times 10^5$. In Table II, we report the values for the bias $\mathcal{B}$, the statistical uncertainty $\mathcal{S}$, and the FIM estimates for the LISA detector. We notice that the bias $\mathcal{B}$ can exceed both the statistical $\mathcal{S}$ and FIM uncertainty.

Other forthcoming space-based GW observatories, such as Taiji [35] or TianQin [36]—which are also designed to work in the milli-Hz band with comparable baselines—will face similar challenges. They will likely adopt analogous segmentation strategies and must likewise account for biases that can exceed statistical uncertainties in their data analysis. In Table II, we compute the expected uncertainties also for Taiji, considering $T_{\text{obs}} = 6$ years, $T = 11.5$ days, and the same frequency resolution as LISA.

Finally, we consider the case of Pulsar Timing Arrays (PTAs). PTAs represent an international effort of six collaborations—NANOGrav, EPTA, InPTA, PPTA, CPTA, and most recently MeerKAT—that monitor the pulse arrival times of ultrastable millisecond pulsars across the sky [37–42]. By analyzing timing residuals accumulated over more than a decade (15 yrs), they have independently measured the

distinctive quadrupolar correlation pattern predicted by Hellings and Downs—a result interpreted as the first detection of a nanohertz-band SGWB. However, the physical origin of such a background remains uncertain; it could arise from a population of inspiraling supermassive black-hole binaries or from cosmological sources, such as cosmic strings, first-order phase transitions, or scalar-induced gravitational waves [43–46].

Because the effective observation time is set by the entire data span, PTAs operate with essentially a single Fourier segment ($N_s \sim 1$). So the suppression of statistical noise is absent; therefore, the experiment is highly sensitive to biases.¹¹ Since these are of the same order of the recovered values, we do not report them explicitly in the $\mathcal{B}$ column of Table II, where we compute the expected uncertainties for current and future PTA detector configurations. For the future PTA configuration, we assumed $T_{\text{obs}} = 15$ yr observations of $N_p = 200$ pulsars with an SKA-like noise configuration (see, e.g., for more detail [31]). For the PTA result only, we estimate the fractional statistical uncertainties using $1/\sqrt{N_f N_s N_p^\alpha}$, including the scaling with the number of pulsars, which here is controlled by the parameter α . Depending on whether cross-correlations or autocorrelations carry more information, α can smoothly interpolate between $\alpha = 2$ and $\alpha = 1$, respectively. Notice that, upon replacing N_p with N_{det} , a similar scaling would also apply when combining the output of several interferometers.

B. A remark on anisotropic searches

Anisotropic or directional SGWB searches using LVK data use a data folding technique to exploit Earth’s rotation [47–49]. Instead of treating each of the N_s segments separately, the cross-correlation data are compressed into one sidereal day (23 h 56 m) of effective time. This is done by “folding” the time-series; segmenting data in 192s chunks aligned to sidereal time and stacking them modulo 1 sidereal day. Following this procedure, the summations over number of segments can be decomposed as $\sum_s \rightarrow \sum_{i_{\text{day}}} \sum_{t_s}$, where i_{day} runs over all sidereal days included in the data set, and t_s spans a single sidereal day. In this way, each sidereal day contains $\simeq 448$ segments of 192 s (since $24_{\text{sid}} \text{h} \simeq 86164$ s), which repeat every day. Thus, the data are mapped onto these $\simeq 448$ sidereal time bins.

For instance, in the O1 + O2 runs, the anisotropic search folded 129 days of data into one sidereal day, i.e., each sidereal-time bin averages over ≈ 129 independent segments (one from each day). Similarly, for the O3 anisotropic analysis, 169 days of data were folded into 1 sidereal

day. So, each of the sidereal bins contains the averaged cross-power from 169 different days at that sidereal time.

Given these remarks, it might appear that the effective number of independent time segments in a folded anisotropic search is of order only a few hundred (the number of sidereal bins) per baseline, much lower than in the isotropic case, which compresses all segments in one large average (losing sky-direction information). However, the folded anisotropy search avoids this issue by performing an additional operation, assuming a stationary distribution of GW power on the sky. Given the data folded onto the sidereal day, the time and direction-dependent (meaning no average over sky-directions) overlap reduction function is deconvolved from the data, allowing compression over the time bins of the sidereal day. At the end of this procedure, anisotropic searches effectively compress over the same number of segments as the isotropic search, which might preserve approximate Gaussianity at the level of accuracy set by the statistical uncertainty.

VI. CONCLUSIONS AND OUTLOOK

We have presented a detailed, end-to-end investigation of the Bayesian framework underpinning SGWB searches. We have tested the validity of using the standard Gaussian likelihood employed by current LVK searches to model the cross-correlation statistics.

Beginning with the exact Whittle likelihood in the simplest setting of a single detector and a single frequency bin, we have shown analytically that replacing the true likelihood with some approximated version can introduce a bias relative to the maximum-likelihood estimator. Then, we have extended these toy models to multiple (i.e., two and three) detectors for both white and realistic colored noise, and we have quantified how the bias $\mathcal{B}$ and the statistical uncertainty $\mathcal{S}$ depend on segment length T , total observation time, and bandwidth.

For second-generation ground-based detectors (LVK) and next-generation detectors (ET/CE), we find that very short segments $\mathcal{O}(1)$ s yield $\mathcal{B} \ll \mathcal{S}$, preserving a “statistics-limited” regime. However, pushing T into the $\mathcal{O}(10)$ minutes regime can bring $\mathcal{B}$ to grow larger than $\mathcal{S}$, inducing a significant bias in the estimation of the SGWB energy density. For space-based missions (LISA/Taiji), which typically divide their multi-year observation into $\mathcal{O}(10^2)$ segments of roughly ~ 11 days length, such a segmentation can lead to sizable bias, which typically exceeds the statistical uncertainty. Finally, for PTA $N_s \sim 1$, which places them in the bias-limited regime. We discussed the impact of our investigations for anisotropic SGWB searches. We concluded that, since sidereal-day folding reorganizes, but does not reduce, the effective sample count, the results derived for isotropic searches continue to be valid.

To better quantify the uncertainties, we have developed two classes of modified likelihoods. The “reduced” Whittle

¹¹Notice that the full analysis based on the time-of-arrival residuals does not rely on the definition of an estimator for the cross-correlations, meaning it will not suffer from the biases described in this work.

likelihood conditioned on fiducial data-driven autocorrelation spectra, and the “optimal-filter” one, which applies inverse-noise weighting across frequency and detector pairs to minimize variance without the problem of extra bias. Both methods recover unbiased maximum-likelihood estimates and could be easily integrated into existing stochastic pipelines.

As the detector sensitivities improve and SGWB detections become affordable, it will be essential to incorporate these refined likelihoods into the future LVK analyses, LISA/ET/CE data challenge, and PTA pipelines. Combining these likelihoods with optimized segment choices and data-compression strategies can ensure SGWB detection and robust, unbiased parameter reconstruction. Including non-Gaussian noise transients, frequency-frequency correlations, and/or time-varying noise covariances are interesting possibilities for further generalizations. Moreover, realistic mock data, including resolved and stochastic signals together with additional nonstationary features, would be important to validate our estimator’s performance. We leave all these investigations for future work.

ACKNOWLEDGMENTS

We thank James Alvey for his very valuable comments on a nearly final version of this manuscript. We also thank Jessica Lawrence and Kevin Turbang for providing the “dictionary” plot in Fig. 7. M. P. acknowledges the hospitality of Imperial College London, which provided office space during some parts of this project. The work of M. P. is supported by the Comunidad de Madrid under the Programa

de Atracción de Talento Investigador No. 2024-T1TEC-31343. J. D. R. acknowledges financial support from the National Science Foundation (NSF) Grant No. PHY-2207270 and startup funds from the University of Texas Rio Grande Valley.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX: FURTHER TESTS

In this appendix, we present a series of additional tests for addressing the robustness of the different likelihoods and to study their behavior in different signal-to-noise regimes. In Fig. 6, we present a comparison of the posterior distributions for the different likelihoods increasing the number of segments to $N_s = 32$ (left plot) or injecting a “very strong” signal, $S_h = 100$ (right plot), to compare with the behavior of the posteriors in the case of a low number of segments $N_s = 3$ and “intermediate” signal (i.e., comparable with noise) $S_h = 0.1$, presented in the main text in Fig. 5. In the left plot, we notice that, as N_s increases, the uncertainties become smaller and the bias decreases relative to the uncertainty. On the other hand, as shown in the right plot, by increasing the signal amplitude, the relative uncertainty (i.e., mean over standard deviation) shrinks, but the bias becomes more significant. Both of these behaviors are as expected.

To aid the interpretation of the probability-probability (P-P) plots, we refer to Fig. 7, which serves as a practical




FIG. 6. Left: Similar to Fig. 5, with an injected signal with $S_h = 0.1$, noise amplitude $P_{n_1} = 1.25$, $P_{n_2} = 1.08$ and $P_{n_3} = 1.25$ for each detector, but increasing the number of segments to $N_s = 32$. Right: Similar to Fig. 5, with an injected signal with $S_h = 10^2$, noise amplitude $P_{n_1} = 1.25$, $P_{n_2} = 1.08$ and $P_{n_3} = 1.25$ for each detector, and number of segments $N_s = 3$.




FIG. 7. “Dictionary” for interpreting P-P plots. The blue probability distribution in the inset is the true distribution; the red probability distribution is the recovered distribution. The different rows and columns correspond to different values of the mean (i.e., possible bias) and uncertainty (i.e., possible under or over coverage) of the recovered distributions. (Credit: Jessica Lawrence and Kevin Turbang.).

“dictionary” for diagnosing common types of patterns. Each panel in Fig. 7 shows the empirical CDF of the posterior quantiles plotted against that of the true CDF (blue line), compared to what you would expect if the two distributions agree (diagonal dotted black line), with shaded bands corresponding to 1 σ , 2 σ , and 3 σ confidence regions. Insets in each panel display the assumed true (blue) and recovered (red) probability density functions, with vertical dashed lines marking the 68% confidence interval. The rows and columns correspond to different values of the

mean μ and uncertainty scaling σ of the recovered distribution. When the P-P plot lies systematically above (below) the diagonal, the inferred mean is biased low (high). Additionally, underestimation of the posterior width leads to an “S-shaped” P-P plot, as the injected parameter preferentially falls in the tails of the distribution. On the other hand, an overestimation of the posterior width leads to an “inverted S” P-P plot, which indicates the injected parameter preferentially sits in the central quantiles of the distribution.



P-P Plots for Different Likelihoods and Signal Regimes, $N_s = 3$, $N_f = 1000$ 

FIG. 8. P-P plot for $N_s = 3$ and $N_f = 1000$ for the different likelihoods presented in the paper, for three different signal regimes: “weak”, i.e., $S_h = 0.01$ (left column), “intermediate”, i.e., $S_h = 0.1$ (middle column), “strong”, i.e., $S_h = 1$ (right column).

P-P Plots for Different Likelihoods and Signal Regimes, $N_s = 60$, $N_f = 50$

 FIG. 9. P-P plots as in Fig. 8 for $N_s = 60$ and $N_f = 50$ for the different likelihoods presented in the paper, for three different signal regimes: “weak”, i.e., $S_h = 0.01$ (left column), “intermediate”, i.e., $S_h = 0.1$ (middle column), “strong”, i.e., $S_h = 1$ (right column).

We proceed by discussing Figs. 8 and 9, which show the P-P plots for the cases where $N_s = 3$, $N_f = 1000$ and $N_s = 60$, $N_f = 50$, respectively, for the “weak” ($S_h = 0.01$), “intermediate” ($S_h = 0.1$), and “strong” ($S_h = 1$) signal regimes. For these plots, we generate 1000 different realizations of the data with the same injected signal value.

Let us start by discussing Fig. 8. We notice that the P-P plots for the Whittle $\mathcal{L}^W(\theta|D_k^s)$, Wishart $\mathcal{L}^W(\theta|\bar{D}_k)$, and Fisher $\mathcal{L}^F(\theta|D_k^s)$ likelihoods all lie almost exactly along the diagonal in the weak, intermediate and strong regimes, indicating nearly perfect agreement with the true probability distributions. On the other hand, the reduced Whittle likelihood $\mathcal{L}^{W,\text{red}}(\theta|\bar{D}_k)$ underestimates the variance and is biased low, with both of these deviations becoming more significant as the signal strength increases relative to the noise. The frequency-dependent cross-correlation Gaussian likelihood, $\mathcal{L}^{G,\text{cc}}(\theta|\bar{D}_k)$, is also biased low, so much so that for large signals the recovered distribution has minimal overlap with the true distribution. On the other hand, the “optimal” cross-correlation Gaussian $\mathcal{L}^{G,\text{opt}}(\theta|\bar{D}_k)$ is unbiased, but underestimates the uncertainty in the recovery of the mean.

Figure 9 is similar to Fig. 8, but shows what happens when we have more segments ($N_s = 60$ instead of $N_s = 3$) but fewer frequency bins ($N_f = 50$ instead of $N_f = 1000$). Basically, as N_s increases, the distributions become more Gaussian and all of the P-P plots become more diagonal, compared to what we saw in Fig. 8. Biases are still present for both the reduced Whittle $\mathcal{L}^{W,\text{red}}(\theta|\bar{D}_k)$ and frequency-dependent cross-correlation Gaussian $\mathcal{L}^{G,\text{cc}}(\theta|\bar{D}_k)$ likelihoods (increasing with increasing signal strength), but they are smaller than before because N_s is larger. The reduced likelihoods $\mathcal{L}^{W,\text{red}}(\theta|\bar{D}_k)$, $\mathcal{L}^{G,\text{cc}}(\theta|\bar{D}_k)$, and $\mathcal{L}^{G,\text{opt}}(\theta|\bar{D}_k)$ continue to underestimate the variance, since this is a consequence of conditioning the likelihoods on the auto-power estimates. Recall that these estimates are independent of the total number of segments and depend on the number of frequencies only via the smoothing performed when constructing the autopower estimates.

In short, Figs. 8 and 9 allow for a diagnosis of both the accuracy and precision of the inference procedure using different likelihoods, signal regimes, frequencies, and segments.

-
- [1] T. Regimbau, *Res. Astron. Astrophys.* **11**, 369 (2011).
 - [2] C. Caprini and D. G. Figueroa, *Classical Quantum Gravity* **35**, 163001 (2018).
 - [3] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
 - [4] F. Acernese *et al.* (Virgo Collaboration), *Classical Quantum Gravity* **32**, 024001 (2014).
 - [5] T. Akutsu *et al.* (KAGRA Collaboration), *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2021).
 - [6] M. Maggiore *et al.* (ET Collaboration), *J. Cosmol. Astropart. Phys.* **03** (2020) 050.
 - [7] A. Abac *et al.*, [arXiv:2503.12263](#).
 - [8] D. Reitze *et al.*, *Bull. Am. Astron. Soc.* **51**, 035 (2019).
 - [9] M. Colpi *et al.* (LISA Collaboration), [arXiv:2402.07571](#).
 - [10] B. Allen and J. D. Romano, *Phys. Rev. D* **59**, 102001 (1999).
 - [11] J. D. Romano and N. J. Cornish, *Living Rev. Relativity* **20**, 2 (2017).
 - [12] I. Caporali, G. Capurri, W. Del Pozzo, A. Ricciardone, and L. Valbusa Dall’Armi, *Phys. Rev. D* **112**, 022005 (2025).
 - [13] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **100**, 061101 (2019).
 - [14] C. Caprini, D. G. Figueroa, R. Flauger, G. Nardini, M. Peloso, M. Pieroni, A. Ricciardone, and G. Tasinato, *J. Cosmol. Astropart. Phys.* **11** (2019) 017.
 - [15] R. Flauger, N. Karnesis, G. Nardini, M. Pieroni, A. Ricciardone, and J. Torrado, *J. Cosmol. Astropart. Phys.* **01** (2021) 059.
 - [16] P. Artist Moran and P. Whittle, *J. R. Stat. Soc., Ser. A* **114**, 579 (1951).
 - [17] J. Wishart, *Biometrika* **20A**, 32 (1928).
 - [18] E. Thrane, N. Christensen, and R. M. S. Schofield, *Phys. Rev. D* **87**, 123009 (2013).
 - [19] M. W. Coughlin *et al.*, *Phys. Rev. D* **97**, 102007 (2018).
 - [20] É. É. Flanagan, *Phys. Rev. D* **48**, 2389 (1993).
 - [21] N. Christensen, *Phys. Rev. D* **55**, 448 (1997).
 - [22] L. Isserlis, *Biometrika* **12**, 134 (1918).
 - [23] S. Hamimeche and A. Lewis, *Phys. Rev. D* **77**, 103013 (2008).
 - [24] L. Verde *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **148**, 195 (2003).
 - [25] A. Matas and J. D. Romano, *Phys. Rev. D* **103**, 062003 (2021).
 - [26] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory* (Prentice Hall, Englewood Cliffs, NJ, 1993), Vol. 1.
 - [27] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), *Phys. Rev. D* **104**, 022004 (2021).
 - [28] L. Barsotti, S. Gras, M. Evans, and P. Fritschel, The updated Advanced LIGO design curve, LIGO (2018), LIGO Document T1800044.
 - [29] M. Branchesi, M. Maggiore *et al.*, *J. Cosmol. Astropart. Phys.* **07** (2023) 068.
 - [30] P. Amaro-Seoane *et al.* (LISA Collaboration), [arXiv:1702.00786](#).
 - [31] S. Babak, M. Falxa, G. Franciolini, and M. Pieroni, *Phys. Rev. D* **110**, 063022 (2024).
 - [32] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), *Phys. Rev. X* **13**, 041039 (2023).
 - [33] M. Evans *et al.*, [arXiv:2306.13745](#).
 - [34] M. Pieroni, A. Ricciardone, and E. Barausse, *Sci. Rep.* **12**, 17940 (2022).

- [35] W.-H. Ruan, Z.-K. Guo, R.-G. Cai, and Y.-Z. Zhang, *Int. J. Mod. Phys. A* **35**, 2050075 (2020).
- [36] J. Luo *et al.* (TianQin Collaboration), *Classical Quantum Gravity* **33**, 035010 (2016).
- [37] G. Agazie *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **951**, L8 (2023).
- [38] G. Agazie *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **951**, L9 (2023).
- [39] J. Antoniadis *et al.* (EPTA, InPTA Collaborations), *Astron. Astrophys.* **678**, A50 (2023).
- [40] D. J. Reardon *et al.*, *Astrophys. J. Lett.* **951**, L6 (2023).
- [41] H. Xu *et al.*, *Res. Astron. Astrophys.* **23**, 075024 (2023).
- [42] M. T. Miles *et al.*, *Mon. Not. R. Astron. Soc.* **536**, 1489 (2024).
- [43] A. Afzal *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **951**, L11 (2023); **971**, L27(E) (2024).
- [44] J. Antoniadis *et al.* (EPTA, InPTA Collaboration), *Astron. Astrophys.* **685**, A94 (2024).
- [45] D. G. Figueroa, M. Pieroni, A. Ricciardone, and P. Simakachorn, *Phys. Rev. Lett.* **132**, 171002 (2024).
- [46] J. Ellis, M. Fairbairn, G. Franciolini, G. Hütsi, A. Iovino, M. Lewicki, M. Raidal, J. Urrutia, V. Vaskonen, and H. Veermäe, *Phys. Rev. D* **109**, 023522 (2024).
- [47] A. Ain, P. Dalvi, and S. Mitra, *Phys. Rev. D* **92**, 022003 (2015).
- [48] A. Ain, J. Suresh, and S. Mitra, *Phys. Rev. D* **98**, 024001 (2018).
- [49] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), *Phys. Rev. D* **104**, 022005 (2021).