

Frequentist versus Bayesian analyses: Cross-correlation as an approximate sufficient statistic for LIGO-Virgo stochastic background searches

Andrew Matas^{1,*} and Joseph D. Romano^{2,†}

¹Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-14476 Potsdam, Germany

²Department of Physics and Astronomy, Texas Tech University,
Box 41051, Lubbock, Texas 79409-1051, USA



(Received 26 December 2020; accepted 8 February 2021; published 16 March 2021)

Sufficient statistics are combinations of data in terms of which the likelihood function can be rewritten without loss of information. Depending on the data volume reduction, the use of sufficient statistics as a preliminary step in a Bayesian analysis can lead to significant increases in efficiency when sampling from posterior distributions of model parameters. Here we show that the frequency integrand of the cross-correlation statistic and its variance are approximate sufficient statistics for ground-based searches for stochastic gravitational-wave backgrounds. The sufficient statistics are approximate because one works in the weak-signal approximation and uses measured estimates of the autocorrelated power in each detector. We perform analytic and numerical calculations to show that, in this approximation, LIGO-Virgo's hybrid frequentist-Bayesian parameter estimation analysis is equivalent to a fully Bayesian analysis. This work closes a gap in the LIGO-Virgo literature and suggests directions for additional searches.

DOI: [10.1103/PhysRevD.103.062003](https://doi.org/10.1103/PhysRevD.103.062003)

I. INTRODUCTION

Current searches for stochastic gravitational-wave backgrounds (GWBs) using ground-based laser interferometers (e.g., the Advanced LIGO [1] and Virgo [2] detectors) use a hybrid of frequentist and Bayesian analysis techniques [3–5]. Certain frequentist statistics (namely, frequency integrands for the cross-correlation statistic and its variance) are calculated for relatively short stretches of time-series data (of order a couple of minutes), and then combined using inverse-noise weighting to produce two final frequency series, valid for the whole observation period (of order several months to a year). These frequency series are then used as the fundamental input data and variance for a subsequent Bayesian analysis that calculates posterior probability distributions and Bayesian upper limits on the strength of potential correlated gravitational-wave signals [6]. These hybrid frequentist-Bayesian analyses have been used to place upper limits on GWBs with different amplitudes and spectral shapes [3,7]; GWBs having non-general relativity polarizations

predicted by alternative theories of gravity [3,8]; and potential contamination from correlated global magnetic fields—i.e., Schumann resonances [9–13].

On the other hand, most searches for GWBs using pulsar timing data [14,15] and proposed searches using space-based detectors such as LISA [16,17] are *fully* Bayesian, proceeding directly from the time-series data to posterior distributions and upper limits, without ever calculating a cross-correlation frequency integrand or statistic. So the question naturally arises as to whether LIGO and Virgo's current hybrid frequentist-Bayesian analysis is losing information relative to a fully Bayesian search.

As we shall show below, the answer is basically “no.” We assume that we can work in the weak-signal approximation, and we use measured estimates of the autocorrelated power in each detector, as opposed to trying to infer the noise power spectra as part of the full analysis. Under these simplifications, the frequency integrands of the cross-correlation statistic and its variance are approximately lossless combinations of the full time-series data in terms of which the full likelihood function can be rewritten. Hence, Bayesian posterior distributions produced from these frequency series agree quite well with those produced from the full time-series data. Said another way, the frequency integrands of the cross-correlation statistic and its variance are approximate *sufficient statistics* for the analysis, so the hybrid frequentist-Bayesian method is essentially equivalent to a full Bayesian analysis.

The calculations presented in this paper can also be thought of as providing an alternative conceptual starting

*andrew.matas@aei.mpg.de

†joseph.d.romano@ttu.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). 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 the Max Planck Society.

point for LIGO-Virgo's stochastic cross-correlation analysis. We use the full Bayesian likelihood function as an *organizing* principle, out of which the LIGO-Virgo stochastic analyses follows. Various steps in the stochastic analysis, such as coarse graining of cross-correlated data [18], estimating autocorrelated power spectra from neighboring data segments [19], inverse-noise weighting, optimal filtering [20], the inclusion of bias factors for the theoretical variance [19], and the use of the cross-correlation frequency integrand and its variance to do parameter estimation and model selection [6], all follow directly from the Bayesian likelihood under the assumptions of a weak signal and estimated autocorrelated detector power spectra. We also see opportunities to develop new methods to look for signals which violate our assumptions.

Our main result will be to show that the full Gaussian likelihood for a stochastic background [cf. Eqs. (31) and (32)] is equivalent to the reduced likelihood given in Eq. (48), under the approximations enumerated above. To do this, we will build up the tools necessary in a series of increasingly more realistic (and more complex) scenarios. In Sec. II, we give a simple example of a sufficient statistic in the context of a search for a constant deterministic signal in the output of a single detector. This example serves as a basis for the calculations in the following two sections for cross-correlation-based searches for stochastic gravitational waves. We first restrict attention in Sec. III to white signal + noise models and stationary data. We then extend our analyses in Sec. IV to colored signal + noise models, allowing for nonstationary noise. Both of these sections present results of different analyses performed on simulated data comparing posterior distributions produced by fully Bayesian and sufficient-statistic analyses. We conclude in Sec. V with a brief summary, followed by a discussion of other related approaches in the literature, and possible extensions of these results. We also include three Appendixes: Appendix A contains a simple, yet very useful, identity (A1) that we use repeatedly throughout the paper; Appendix B summarizes uncertainties in power spectrum estimation; and Appendix C gives an alternative derivation of a reduced likelihood function, but which makes different assumptions than those given in the main text.

II. SIMPLE EXAMPLE

Perhaps the simplest example of a nontrivial sufficient statistic is the sample mean of the data for a constant signal in white, Gaussian noise (see Sec. 3.5 of Ref. [5] for a more detailed treatment of this example). We suppose we record N time-series data samples $d \equiv \{d_i\}$ as

$$d_i = a + n_i, \quad i = 1, 2, \dots, N, \quad (1)$$

where $a > 0$ is the amplitude of the signal and n_i denotes the i th sample of the noise. For simplicity, we will assume

that the noise is white and has zero mean, i.e., $\langle n_i \rangle = 0$, $\langle n_i n_j \rangle = \sigma^2 \delta_{ij}$, and that the variance σ^2 is known *a priori*. The likelihood function is then

$$p(d|a) = \frac{1}{(2\pi)^{N/2} \sigma^N} \exp \left[-\frac{1}{2\sigma^2} \sum_i (d_i - a)^2 \right], \quad (2)$$

which has the interpretation of being the probability of observing the data d given a signal of amplitude a in white noise with known variance σ^2 .

It is fairly straightforward to show that the maximum-likelihood estimator of the amplitude a is given by the sample mean of the data

$$\hat{a} \equiv a_{\text{ML}}(d) = \frac{1}{N} \sum_i d_i. \quad (3)$$

This is an unbiased estimator of a and has variance

$$\sigma_{\hat{a}}^2 = \sigma^2/N. \quad (4)$$

In terms of \hat{a} , the data-dependent term in the likelihood becomes

$$\sum_i (d_i - a)^2 = \sum_i d_i^2 - N\hat{a}^2 + N(\hat{a} - a)^2. \quad (5)$$

This equation is a special case of the general identity (A1), which is discussed and proven in Appendix A, and which we will use in future sections. The likelihood can then be rewritten in the form

$$p(d|a) = \frac{1}{(2\pi)^{N/2} \sigma^N} \times \exp \left[-\frac{1}{2\sigma^2} \sum_i d_i^2 \right] \exp \left[\frac{\hat{a}^2}{2\sigma_{\hat{a}}^2} \right] \exp \left[-\frac{(\hat{a} - a)^2}{2\sigma_{\hat{a}}^2} \right]. \quad (6)$$

From this expression, we see that the likelihood function has factored into a product of components, the last of which depends on the data d_i only via the combination $\hat{a} \equiv a_{\text{ML}}(d)$, with all other components being *independent* of the parameter a . This means that \hat{a} is a *sufficient statistic* for the parameter a for this simple signal + noise model.

The posterior distribution for a , denoted $p(a|d)$, is calculated using Bayes' theorem

$$p(a|d) = \frac{p(d|a)p(a)}{p(d)}, \quad (7)$$

where $p(a)$ is the *prior probability distribution* for a , and

$$p(d) \equiv \int da p(d|a) p(a) \quad (8)$$

is the so-called *evidence* (or *marginalized likelihood*). Since $p(d)$ is independent of a , the evidence acts as a normalization factor as far as the posterior distribution of a is concerned. Thus, one often writes $p(a|d) \propto p(d|a)p(a)$ for posterior distribution calculations. For this example,

$$p(a|d) \propto \exp \left[-\frac{(\hat{a} - a)^2}{2\sigma_{\hat{a}}^2} \right] p(a), \quad (9)$$

where we have ignored all factors in the likelihood function (6) that are independent of the parameter a . Thus, the sufficient statistic $\hat{a} \equiv a_{\text{ML}}(d)$ is the only data combination we need to calculate the posterior. This observation has computational advantages, since if we use a Bayesian sampling algorithm to estimate the posterior, then by precomputing \hat{a} we obtain an $O(N)$ saving relative to the same posterior calculation but using the original form of the likelihood, Eq. (2).

Finally, we note that the prior probability distribution $p(a)$ is chosen based on expectations for a prior to observing the data. It is common to use *flat* or *log-uniform* priors,

$$p(a) = \text{const} \quad \text{or} \quad p(a) = \text{const}/a, \quad (10)$$

defined over some interval $[a_{\min}, a_{\max}]$, where the constants in these expressions are determined by the normalization condition $\int da p(a) = 1$. For the analyses that we will perform in the following sections, we will consider flat priors for simplicity, but note here that the results we obtain are valid for arbitrary, nonflat priors as well.

III. SUFFICIENT STATISTICS FOR CROSS-CORRELATION SEARCHES—WHITE SIGNAL + NOISE MODELS

Here we extend the calculations of the previous section to cross-correlation searches for stochastic gravitational waves. We restrict our attention in Sec. III A to a white signal + noise model, assuming stationary signal and noise (see also Sec. 4 of Ref. [5]). In Sec. III B, we consider a *reduced* version of this model, where we assume weak signals relative to the detector noise and use measured estimates of the autocorrelated power in each detector, as opposed to having these as model parameters to be determined by the search. Our analysis of the white noise case contains all of the important steps present in the colored noise case that we shall discuss in Sec. IV, but with considerably less complication.

A. White signal + noise, stationary data

To start, let us consider two coincident and coaligned detectors with uncorrelated noise. We take both the detector noise and correlated stochastic signal to be Gaussian, stationary, and white. We denote the variances by $\sigma_{n_1}^2$,

$\sigma_{n_2}^2$, σ_h^2 , respectively. We will not assume that we know the noise variances *a priori*, so we will try to recover $\sigma_{n_1}^2$ and $\sigma_{n_2}^2$ in addition to σ_h^2 .

Then the likelihood function is given by

$$p(d|\sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2) = \frac{1}{\sqrt{\det(2\pi C)}} e^{-\frac{1}{2}d^T C^{-1}d}, \quad (11)$$

where

$$C \equiv \begin{bmatrix} (\sigma_{n_1}^2 + \sigma_h^2)\mathbb{1}_{N \times N} & \sigma_h^2\mathbb{1}_{N \times N} \\ \sigma_h^2\mathbb{1}_{N \times N} & (\sigma_{n_2}^2 + \sigma_h^2)\mathbb{1}_{N \times N} \end{bmatrix} \quad (12)$$

is the covariance matrix and

$$d^T C^{-1}d \equiv \sum_{\alpha, \beta} \sum_{i, j} d_{\alpha i} (C^{-1})_{\alpha i, \beta j} d_{\beta j}. \quad (13)$$

The indices $i, j = 1, 2, \dots, N$ label individual time samples, and $\alpha, \beta = 1, 2$ label the two detectors. [The covariance matrix C and its inverse C^{-1} are thus 2×2 block matrices with $N \times N$ blocks; cf. (12) and (19).] The joint posterior distribution for the signal and noise variances is

$$p(\sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2|d) = \frac{p(d|\sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2)p(\sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2)}{p(d)}, \quad (14)$$

where $p(\sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2)$ is the joint prior probability distribution, and

$$p(d) = \int d\sigma_{n_1}^2 \int d\sigma_{n_2}^2 \int d\sigma_h^2 \times p(d|\sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2)p(\sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2) \quad (15)$$

is the evidence for this signal + noise model.

It is easy to show that the maximum-likelihood estimators for the parameters $\sigma_{n_1}^2$, $\sigma_{n_2}^2$, σ_h^2 are given by the following quadratic combinations of the data:

$$\begin{aligned} \hat{\sigma}_{n_1}^2 &\equiv \frac{1}{N} \sum_i d_{1i}^2 - \frac{1}{N} \sum_i d_{1i}d_{2i}, \\ \hat{\sigma}_{n_2}^2 &\equiv \frac{1}{N} \sum_i d_{2i}^2 - \frac{1}{N} \sum_i d_{1i}d_{2i}, \\ \hat{\sigma}_h^2 &\equiv \frac{1}{N} \sum_i d_{1i}d_{2i}. \end{aligned} \quad (16)$$

These expressions show that it is also convenient to define

$$\hat{\sigma}_1^2 \equiv \frac{1}{N} \sum_i d_{1i}^2, \quad \hat{\sigma}_2^2 \equiv \frac{1}{N} \sum_i d_{2i}^2, \quad (17)$$

which are estimators of the total autocorrelated variances in the two detectors:

$$\sigma_1^2 \equiv \sigma_{n_1}^2 + \sigma_h^2, \quad \sigma_2^2 \equiv \sigma_{n_2}^2 + \sigma_h^2. \quad (18)$$

In the weak-signal limit $\sigma_1^2 \approx \sigma_{n_1}^2$ and $\sigma_2^2 \approx \sigma_{n_2}^2$, but we will *not* make that approximation at this stage. We will discuss this approximation in Sec. III B.

To show that the above estimators are sufficient statistics for this problem, it suffices to show that the data enter the likelihood function (11) *solely* in the form of these estimators, up to an overall normalization that does not depend on the signal and noise parameters. Since the only part of the likelihood function that depends on the data is the argument of the exponential, we just need to show that $d^T C^{-1} d$ can be written in terms of $\hat{\sigma}_{n_1}^2$, $\hat{\sigma}_{n_2}^2$, $\hat{\sigma}_h^2$, or, equivalently, in terms of $\hat{\sigma}_1^2$, $\hat{\sigma}_2^2$, $\hat{\sigma}_h^2$.

Since the covariance matrix C is a 2×2 -block matrix with each block proportional to $\mathbb{1}_{N \times N}$, we can explicitly invert C , yielding

$$C^{-1} = \frac{1}{\beta} \begin{bmatrix} \sigma_2^2 \mathbb{1}_{N \times N} & -\sigma_h^2 \mathbb{1}_{N \times N} \\ -\sigma_h^2 \mathbb{1}_{N \times N} & \sigma_1^2 \mathbb{1}_{N \times N} \end{bmatrix}, \quad (19)$$

where

$$\begin{aligned} \beta &\equiv \sigma_1^2 \sigma_2^2 - (\sigma_h^2)^2 \\ &= \sigma_{n_1}^2 \sigma_{n_2}^2 + (\sigma_{n_1}^2 + \sigma_{n_2}^2) \sigma_h^2 > 0. \end{aligned} \quad (20)$$

Using this result, it is then straightforward to show that

$$\begin{aligned} d^T C^{-1} d &= \frac{1}{\beta} \left[\sigma_2^2 \sum_i d_{1i}^2 + \sigma_1^2 \sum_i d_{2i}^2 - 2\sigma_h^2 \sum_i d_{1i} d_{2i} \right] \\ &= \frac{N}{(1 - (\sigma_h^2)^2 / (\sigma_1^2 \sigma_2^2))} \left[\frac{\hat{\sigma}_1^2}{\sigma_1^2} + \frac{\hat{\sigma}_2^2}{\sigma_2^2} - 2 \frac{\sigma_h^2 \hat{\sigma}_h^2}{\sigma_1^2 \sigma_2^2} \right]. \end{aligned} \quad (21)$$

Thus,

$$\begin{aligned} p(d | \sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2) &= \frac{1}{(2\pi)^N (\sigma_1^2 \sigma_2^2 - (\sigma_h^2)^2)^{N/2}} \\ &\times \exp \left\{ -\frac{1}{2} \frac{N}{(1 - (\sigma_h^2)^2 / (\sigma_1^2 \sigma_2^2))} \left[\frac{\hat{\sigma}_1^2}{\sigma_1^2} + \frac{\hat{\sigma}_2^2}{\sigma_2^2} - 2 \frac{\sigma_h^2 \hat{\sigma}_h^2}{\sigma_1^2 \sigma_2^2} \right] \right\}, \end{aligned} \quad (22)$$

which depends on the data only via $\hat{\sigma}_1^2$, $\hat{\sigma}_2^2$, $\hat{\sigma}_h^2$, or, equivalently, $\hat{\sigma}_{n_1}^2$, $\hat{\sigma}_{n_2}^2$, $\hat{\sigma}_h^2$.

B. Weak-signal approximation, estimating detector autocorrelations

We now consider a *reduced* signal + noise model and its corresponding likelihood function, again for the case of white signal + noise, stationary data. The reduction in the model has two components: First, instead of treating the detector noise variances $\sigma_{n_1}^2$, $\sigma_{n_2}^2$ or the total autocorrelated variances σ_1^2 , σ_2^2 as free parameters for our analysis, we will use *measured estimates* $\bar{\sigma}_1^2$, $\bar{\sigma}_2^2$ in place of σ_1^2 , σ_2^2 . We use overbars instead of hats to denote these quantities to indicate that $\bar{\sigma}_1^2$, $\bar{\sigma}_2^2$ are not the same as the maximum-likelihood data combinations $\hat{\sigma}_1^2$, $\hat{\sigma}_2^2$ for the segment that we are analyzing. For example, for the actual LIGO-Virgo stochastic searches, the estimates of the autocorrelated power are constructed from two segments of data (each approximately one minute in duration) immediately preceding and following the analysis segment. In fact, we will show in Fig. 4 that if the measured estimates of the autocorrelated power equal the maximum-likelihood data combinations for the analysis segment, then we obtain biased recoveries of the GWB spectrum, due to covariances between the autocorrelation and cross-correlation estimators [19]. (But see also the discussion of bias at the end of Appendix C.)

Second, we will assume that the stochastic signal is *weak* compared to the detector noise and thus keep only the leading-order terms in expressions involving $(\sigma_h^2)^2 / \bar{\sigma}_1^2 \bar{\sigma}_2^2 \ll 1$. To zeroth order, the detector noise and total autocorrelated variances are equal to one another, $\sigma_{n_1}^2 \approx \sigma_1^2$, $\sigma_{n_2}^2 \approx \sigma_2^2$. This weak-signal approximation is valid for searches for GWBs using ground-based detectors such as LIGO and Virgo. It is not a good approximation for GWB searches using pulsar timing arrays, where the autocorrelated power in the GWB may exceed that due to pulsar noise and timing measurement noise at very low frequencies [15,21].

The likelihood function for the reduced signal + noise model can be obtained from (22) by making the simplifications described above. We first approximate the terms $\hat{\sigma}_1^2 / \sigma_1^2$ and $\hat{\sigma}_2^2 / \sigma_2^2$ by 1, given that we are replacing the parameters σ_1^2 and σ_2^2 by measured estimates of these quantities. We then replace the remaining factors of σ_1^2 , σ_2^2 , which appear in the likelihood function in the combination $1 / \sigma_1^2 \sigma_2^2$ by

$$\frac{1}{\bar{\Sigma}_{12}^4} \equiv \frac{1}{\bar{\sigma}_1^2 \bar{\sigma}_2^2 (1 + 2/N_{\text{avg}})}, \quad (23)$$

where N_{avg} (which we assume to be equal for both detectors) is the number of averages used in the construction of $\bar{\sigma}_1^2$, $\bar{\sigma}_2^2$, e.g., Welch power spectrum estimates [22]. The justification for including the factor of $(1 + 2/N_{\text{avg}})$ is given in Appendix B; the factor removes a bias that would otherwise exist in the estimation of $1 / \sigma_1^2 \sigma_2^2$, due to the use of a finite amount of data to construct $\bar{\sigma}_1^2$, $\bar{\sigma}_2^2$. Making all these replacements in (22), we obtain

$$p(d|\bar{\sigma}_1^2, \bar{\sigma}_2^2, \sigma_h^2) = \frac{1}{(2\pi)^N (\bar{\Sigma}_{12}^4)^{N/2} (1 - (\sigma_h^2)^2 / \bar{\Sigma}_{12}^4)^{N/2}} \exp \left\{ -\frac{N}{(1 - (\sigma_h^2)^2 / \bar{\Sigma}_{12}^4)} \left[1 - \frac{\sigma_h^2 \hat{\sigma}_h^2}{\bar{\Sigma}_{12}^4} \right] \right\}. \quad (24)$$

We then Taylor expand the right-hand side (RHS) keeping only the leading-order terms in $(\sigma_h^2)^2 / \sigma_1^2 \sigma_2^2 = (\sigma_h^2)^2 / \bar{\Sigma}_{12}^4$. The σ_h^2 factor in front of the exponential can be written as

$$\frac{1}{(1 - (\sigma_h^2)^2 / \bar{\Sigma}_{12}^4)^{N/2}} \approx 1 + \frac{N}{2} \frac{(\sigma_h^2)^2}{\bar{\Sigma}_{12}^4} \approx \exp \left[\frac{N}{2} \frac{(\sigma_h^2)^2}{\bar{\Sigma}_{12}^4} \right], \quad (25)$$

while the exponential factor itself becomes

$$\exp \left\{ -\frac{N}{(1 - (\sigma_h^2)^2 / \bar{\Sigma}_{12}^4)} \left[1 - \frac{\sigma_h^2 \hat{\sigma}_h^2}{\bar{\Sigma}_{12}^4} \right] \right\} \approx \exp \left\{ -N \left(1 + \frac{(\sigma_h^2)^2}{\bar{\Sigma}_{12}^4} - \frac{\sigma_h^2 \hat{\sigma}_h^2}{\bar{\Sigma}_{12}^4} \right) \right\}. \quad (26)$$

Combining these last two expressions gives

$$\exp \left[-\frac{N}{2} \left(2 + \frac{(\sigma_h^2)^2 - 2\sigma_h^2 \hat{\sigma}_h^2}{\bar{\Sigma}_{12}^4} \right) \right] = e^{-N} \exp \left[\frac{(\hat{\sigma}_h^2)^2}{2\text{var}(\hat{\sigma}_h^2)} \right] \exp \left[-\frac{(\hat{\sigma}_h^2 - \sigma_h^2)^2}{2\text{var}(\hat{\sigma}_h^2)} \right], \quad (27)$$

where we completed the square in σ_h^2 and $\hat{\sigma}_h^2$ and made the substitution

$$\text{var}(\hat{\sigma}_h^2) \equiv \bar{\Sigma}_{12}^4 / N = (1 + 2/N_{\text{avg}}) \bar{\sigma}_1^2 \bar{\sigma}_2^2 / N, \quad (28)$$

which is the leading-order expression for the variance of the maximum-likelihood estimator $\hat{\sigma}_h^2$. [The full expression for the variance would include an additional term proportional to the square of σ_h^2 ; cf. Eq. (4.18) in [5]. We also derive a closely related relationship involving estimators constructed from the data in (C12).] Thus,

$$p(d|\bar{\sigma}_1^2, \bar{\sigma}_2^2, \sigma_h^2) = \frac{e^{-N}}{(2\pi)^N (N\text{var}(\hat{\sigma}_h^2))^{N/2}} \exp \left[\frac{(\hat{\sigma}_h^2)^2}{2\text{var}(\hat{\sigma}_h^2)} \right] \exp \left[-\frac{(\hat{\sigma}_h^2 - \sigma_h^2)^2}{2\text{var}(\hat{\sigma}_h^2)} \right], \quad (29)$$

which shows that $\hat{\sigma}_h^2$ and $\text{var}(\hat{\sigma}_h^2)$ are sufficient statistics for this reduced signal + noise model. We can also see from (29) that $\hat{\sigma}_h^2$ is the maximum-likelihood estimator of σ_h^2 for this reduced signal + noise model. Note that the σ_h^2 -dependent term in the likelihood function has the same general form as that for the simple example described in Sec. II; see (6). (In Appendix C, we present an alternative derivation of the likelihood function for a reduced signal + noise model, but under different assumptions than used here.)

Figure 1 compares the marginalized posterior distributions for σ_h^2 calculated from the two likelihood functions, Eqs. (22) and (29), for a short stretch of simulated time-series data consisting of a white signal injected into white noise in two coincident and coaligned detectors. The simulated noise had variances $\sigma_{n_1}^2 = \sigma_{n_2}^2 = 1$, while the injected signal had $\sigma_h^2 = 0.3$. For 512 samples, these values correspond to an expected signal-to-noise ratio $\rho = \sqrt{N}\sigma_h^2 / \sqrt{\sigma_1^2 \sigma_2^2} = 5.22$ for the full set of data. Note, however, that $(\sigma_h^2)^2 / \sigma_1^2 \sigma_2^2 = 0.05 \ll 1$, consistent with the weak-signal approximation. The blue histogram is the marginalized posterior for σ_h^2 calculated from the full likelihood function (22); the orange histogram is the marginalized posterior for σ_h^2 calculated from the reduced likelihood function (29). For the reduced likelihood analysis, the detector autocorrelated power were estimated from additional, adjacent, simulated LIGO-Virgo analyses such as [23]. The dashed vertical gray line and the dashed vertical red line show the injected value of σ_h^2 and the maximum-likelihood value $\hat{\sigma}_h^2$.




FIG. 1. Comparison of the marginalized posteriors for the full likelihood function (22) (blue histogram) and its reduced version (29) (orange histogram), which substitutes measured estimates $\bar{\sigma}_1^2, \bar{\sigma}_2^2$ for the autocorrelated power σ_1^2, σ_2^2 and keeps only leading-order terms in $(\sigma_h^2)^2 / \sigma_1^2 \sigma_2^2$. (Note that since this posterior is based on a single realization of the data, we cannot conclude whether the offset of the peak of the posterior relative to the injected value is because of a statistical fluctuation or a systematic bias. We return to this point in Figs. 4 and 5.)

IV. SUFFICIENT STATISTICS FOR CROSS-CORRELATION SEARCHES—COLORED SIGNAL + NOISE MODELS

Here we extend the analysis of the previous section to a colored signal + noise model, starting with stationary data in Sec. IV A and then discussing the complications introduced by nonstationary noise in Sec. IV B. In Sec. IV C, we simplify the signal + noise model by again considering weak signals and measured estimates of the autocorrelated power, for which the frequency integrands of the cross-correlation statistic and its variance are sufficient statistics. In Sec. IV D we show that the full Bayesian analysis is approximately equivalent to LIGO-Virgo's hybrid frequentist-Bayesian analysis, and in Sec. IV E we construct percentile-percentile (pp) plots [24] to show that the reduced analyses have proper statistical coverage. (We do not assume that the detectors are coincident and coaligned in this section.)

A. Colored signal + noise, stationary data

For the case where the detector noise and GWB signal are colored, it is simpler to work in the frequency domain, since the Fourier components are *independent* of one another (assuming the data are stationary over the duration of the analysis segment). Assuming multivariate Gaussian distributions as before, the variances $\sigma_{n_1}^2$, $\sigma_{n_2}^2$, σ_h^2 , for the white signal + noise case are replaced by *power spectral densities* $P_{n_1}(f)$, $P_{n_2}(f)$, $P_h(f)$, defined by

$$\langle \tilde{h}(f)\tilde{h}^*(f') \rangle = \frac{1}{2}\delta(f-f')P_h(f), \quad \text{etc.}, \quad (30)$$

where $\tilde{h}(f)$ is the Fourier transform of the signal component $h(t)$ of the time-series data. The factor of 1/2 is included to make these *one-sided* power spectral densities, so that the total variance is given by an integral of the power spectral density over *positive frequencies*, $\sigma^2 = \int_0^\infty df P(f)$. Although the above expressions are written in terms of continuous functions of frequency, in practice we work with frequency series, e.g., $P_h(f_k)$, where the discrete frequencies take values $f_k \equiv k\Delta f$, where $\Delta f \equiv 1/T$ and $k = 0, 1, \dots, N/2 - 1$, for a data segment of duration $T \equiv N\Delta t$. For white data, $P(f) = 2\sigma^2\Delta t$ is constant between $f = 0$ and the Nyquist frequency $f_{\text{Nyq}} \equiv 1/2\Delta t$.

Even though we are assuming here that the data are stationary, it is convenient to divide the data from a large observation period into segments, which we will label by $I = 1, 2, \dots, N_{\text{seg}}$. The full likelihood function is then a product of the likelihood functions for the individual segments

$$p(d|P_{n_1}, P_{n_2}, P_h) = \prod_I p_I(d_I|P_{n_1}, P_{n_2}, P_h), \quad (31)$$

where

$$p_I(d_I|P_{n_1}, P_{n_2}, P_h) = \prod_k \frac{1}{\det(2\pi\tilde{C}(f_k))} e^{-\frac{1}{2}\sum_{\alpha,\beta} \tilde{d}_{\alpha I}^*(f_k)(\tilde{C}(f_k)^{-1})_{\alpha\beta}\tilde{d}_{\beta I}(f_k)}. \quad (32)$$

In the above expression

$$\tilde{C}(f) = \frac{T}{4} \begin{bmatrix} P_1(f) & \gamma(f)P_h(f) \\ \gamma(f)P_h(f) & P_2(f) \end{bmatrix} \quad (33)$$

is the covariance matrix of the data, with

$$\begin{aligned} P_1(f) &\equiv P_{n_1}(f) + P_h(f), \\ P_2(f) &\equiv P_{n_2}(f) + P_h(f), \end{aligned} \quad (34)$$

denoting the total autocorrelated detector power spectral densities. The inverse covariance matrix is simply

$$\begin{aligned} \tilde{C}(f)^{-1} &= \frac{4}{T} \frac{1}{(P_1(f)P_2(f) - \gamma^2(f)P_h^2(f))} \\ &\times \begin{bmatrix} P_2(f) & -\gamma(f)P_h(f) \\ -\gamma(f)P_h(f) & P_1(f) \end{bmatrix}. \end{aligned} \quad (35)$$

Note that although the data depend on segment I , the parameters P_{n_1} , P_{n_2} , P_h do not, since we are assuming that the noise and signal power spectra are the same in each segment.

The dimensionless function $\gamma(f)$, which appears in the off-diagonal elements of the covariance matrix, is the *overlap reduction function*, which accounts for the relative position and orientation of the detectors [25,26]. The functional form of $\gamma(f)$ is not relevant for the discussion that follows, other than the fact that $\gamma(f)$ equals unity for coincident and coaligned detectors in the long-wavelength limit.

If we want to estimate the values of the power spectral densities $P_{n_1}(f)$, $P_{n_2}(f)$, $P_h(f)$, at *each* discrete (positive) frequency $f_k \equiv k\Delta f$, then there are no simplifying sufficient statistics for this case as the data enter the likelihood function only through the combinations

$$|\tilde{d}_{1I}(f_k)|^2, \quad |\tilde{d}_{2I}(f_k)|^2, \quad \text{Re}(\tilde{d}_{1I}^*(f_k)\tilde{d}_{2I}(f_k)), \quad (36)$$

which does not correspond to a reduction in the number of data samples used in writing the likelihood function. However, if the power spectra are expected to be smooth over a *coarser* frequency resolution $\delta f \equiv 1/\tau > \Delta f$, where $\tau \equiv T/M$ is some fractional part of the segment duration T , then there is a reduction in the data combinations. This is because the relevant power spectra need only be estimated at fewer discrete frequencies $f_\ell \equiv \ell\delta f$, where $\ell = 0, 1, \dots, (N/M)/2 - 1$. (For typical LIGO-Virgo

searches, M is of order 10 or larger [23].) Hence the data combinations (36) can be averaged over a subset of M fine-grained frequencies f_k centered at f_ℓ :

$$\begin{aligned}\tilde{c}_{11,I}(f_\ell) &\equiv \frac{1}{M} \sum_{k=\ell-M/2}^{\ell+M/2-1} |\tilde{d}_{1I}(f_k)|^2, \\ \tilde{c}_{22,I}(f_\ell) &\equiv \frac{1}{M} \sum_{k=\ell-M/2}^{\ell+M/2-1} |\tilde{d}_{2I}(f_k)|^2, \\ \tilde{c}_{12,I}(f_\ell) &\equiv \frac{1}{M} \sum_{k=\ell-M/2}^{\ell+M/2-1} \text{Re}(\tilde{d}_{1I}^*(f_k)\tilde{d}_{2I}(f_k)).\end{aligned}\quad (37)$$

This leads to averaged (or *coarse-grained*) power spectral density estimators

$$\begin{aligned}\hat{P}_{1I}(f_\ell) &\equiv \frac{2}{T} \tilde{c}_{11,I}(f_\ell), \\ \hat{P}_{2I}(f_\ell) &\equiv \frac{2}{T} \tilde{c}_{22,I}(f_\ell), \\ \hat{P}_{hI}(f_\ell) &\equiv \frac{2}{T} \frac{\tilde{c}_{12,I}(f_\ell)}{\gamma(f_\ell)},\end{aligned}\quad (38)$$

in terms of which the likelihood function (32) can be written:

$$\begin{aligned}p_I(d_I|P_{n_1}, P_{n_2}, P_h) &= \prod_{\ell} \frac{1}{(\pi T/2)^{2M} (P_1(f_\ell)P_2(f_\ell) - \gamma^2(f_\ell)P_h^2(f_\ell))^M} \\ &\times \exp \left\{ -\frac{M}{(1 - \gamma^2(f_\ell)P_h^2(f_\ell)/(P_1(f_\ell)P_2(f_\ell)))} \left[\frac{\hat{P}_{1I}(f_\ell)}{P_1(f_\ell)} + \frac{\hat{P}_{2I}(f_\ell)}{P_2(f_\ell)} - 2\gamma^2(f_\ell) \frac{P_h(f_\ell)\hat{P}_{hI}(f_\ell)}{P_1(f_\ell)P_2(f_\ell)} \right] \right\}.\end{aligned}\quad (39)$$

We note that estimating a power spectrum by subdividing a segment of data into shorter duration subsegments is a standard practice in signal processing [22]. It is a way to reduce the variance in the power spectrum estimate at the expense of a coarser frequency resolution.

For this stationary case, we get a further level of data reduction in the sufficient statistics, as we can average the coarse-grained power spectrum estimators (38) over the number of segments. By construction, these segment-averaged estimators will give expected values of $P_h(f_\ell)$, $P_1(f_\ell)$, $P_2(f_\ell)$ over the *whole* observation. This is fine for stationary data. But if the detector noise levels can change from segment to segment, then this simple averaging will fail to capture the nonstationarity of the noise. We have to do something different for this more complicated, but realistic, scenario.

B. Colored signal + noise, nonstationary data

Real data from ground-based gravitational-wave interferometers is not stationary. As discussed, for instance, in [27], non-Gaussian noise transients (often referred to as *glitches*) can frequently occur, and additionally the noise power spectrum changes with time over the course of the run, for instance due to changes in detector configuration or in response to seismic activity. Therefore, analyses of real data must explicitly test that the data are sufficiently

stationary. One approach developed in [28] is to explicitly test that the estimator follows the expected distribution from Gaussian noise. It is also common for LIGO-Virgo analyses to reject segments where the noise is nonstationary between adjacent segments [23].

For nonstationary detector noise, we have to increase the number of model parameters from $P_{n_1}(f_\ell)$, $P_{n_2}(f_\ell)$, $P_h(f_\ell)$ to $P_{n_{1I}}(f_\ell)$, $P_{n_{2I}}(f_\ell)$, $P_h(f_\ell)$, where $I = 1, 2, \dots, N_{\text{seg}}$, since the noise levels can differ from segment to segment. We are assuming here that the power spectrum of the stochastic signal is stationary, which is not necessarily the case for a ‘‘popcornlike’’ background, such as that produced by stellar-mass binary black hole mergers [29]. The covariance matrix for this case is then

$$\tilde{C}_I(f_\ell) = \frac{T}{4} \begin{bmatrix} P_{1I}(f_\ell) & \gamma(f_\ell)P_h(f_\ell) \\ \gamma(f_\ell)P_h(f_\ell) & P_{2I}(f_\ell) \end{bmatrix},\quad (40)$$

where

$$\begin{aligned}P_{1I}(f) &\equiv P_{n_{1I}}(f) + P_h(f), \\ P_{2I}(f) &\equiv P_{n_{2I}}(f) + P_h(f),\end{aligned}\quad (41)$$

are the autocorrelated detector power spectra for segment I . Similar to what we found in (39), the corresponding likelihood function for a single segment of data can be written as

$$\begin{aligned}p_I(d_I|P_{n_{1I}}, P_{n_{2I}}, P_h) &= \prod_{\ell} \frac{1}{(\pi T/2)^{2M} (P_{1I}(f_\ell)P_{2I}(f_\ell) - \gamma^2(f_\ell)P_h^2(f_\ell))^M} \\ &\times \exp \left\{ -\frac{M}{(1 - \gamma^2(f_\ell)P_h^2(f_\ell)/(P_{1I}(f_\ell)P_{2I}(f_\ell)))} \left[\frac{\hat{P}_{1I}(f_\ell)}{P_{1I}(f_\ell)} + \frac{\hat{P}_{2I}(f_\ell)}{P_{2I}(f_\ell)} - 2\gamma^2(f_\ell) \frac{P_h(f_\ell)\hat{P}_{hI}(f_\ell)}{P_{1I}(f_\ell)P_{2I}(f_\ell)} \right] \right\},\end{aligned}\quad (42)$$

where the estimators $\hat{P}_{hI}(f_\ell)$, $\hat{P}_{1I}(f_\ell)$, $\hat{P}_{2I}(f_\ell)$ are the same as in (38). We emphasize that this likelihood differs from that in (39) only by our assumption that the noise is not stationary, which is reflected in the fact that the parameters $P_{1I}(f_\ell)$, $P_{2I}(f_\ell)$ carry I indices. Note that the above likelihood has the same form as the white noise case (22), but with both frequency and segment dependence in the estimators and parameters. The full expression for the likelihood function involves a further product over I , as specified in (31).

C. Reduced version of the colored, nonstationary likelihood function

To simplify the above analysis, we will proceed as we did in Sec. III B, where we considered a reduced signal + noise model and its corresponding likelihood by replacing the autocorrelated power spectral densities with measured estimates $\bar{P}_{1I}(f_\ell)$, $\bar{P}_{2I}(f_\ell)$, and working in the weak-signal approximation where $P_h^2(f_\ell)/\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell) \ll 1$. For a given discrete frequency f_ℓ and data segment I , the reduction in the likelihood function has exactly the same form as the white noise case, which allows us to immediately write down

$$p_I(d_I|\bar{P}_{1I}, \bar{P}_{2I}, P_h) = \prod_{\ell} \frac{e^{-2M}}{(\pi T/2)^{2M} (\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell)(1+2/N_{\text{avg}}))^M} \exp\left[\frac{\hat{P}_{hI}^2(f_\ell)}{2\bar{\sigma}_{hI}^2(f_\ell)}\right] \exp\left[-\frac{(\hat{P}_{hI}(f_\ell) - P_h(f_\ell))^2}{2\bar{\sigma}_{hI}^2(f_\ell)}\right], \quad (43)$$

where

$$\bar{\sigma}_{hI}^2(f_\ell) \equiv \frac{1}{2T\delta f} \frac{\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell)}{\gamma^2(f_\ell)} (1+2/N_{\text{avg}}), \quad (44)$$

which is the leading-order expression for the variance of $\hat{P}_{hI}(f_\ell)$. [Here we used the relation $M = T\delta f$ for the coarse-grained frequencies, and we included the factor of $(1+2/N_{\text{avg}})$ to account for imperfect estimation of the detector autocorrelated power spectra $\bar{P}_{1I}(f_\ell)$, $\bar{P}_{2I}(f_\ell)$.] Additionally, since the parameter $P_h(f_\ell)$ shows up only in the last exponential, we can use identity (A1) from Appendix A to perform the product over the number of data segments, which translates into a sum over I inside the exponential

$$\sum_I \frac{(\hat{P}_{hI}(f_\ell) - P_h(f_\ell))^2}{\bar{\sigma}_{hI}^2(f_\ell)} = \sum_I \frac{\hat{P}_{hI}^2(f_\ell)}{\bar{\sigma}_{hI}^2(f_\ell)} - \frac{\hat{P}_{h,\text{opt}}^2(f_\ell)}{\bar{\sigma}_{h,\text{opt}}^2(f_\ell)} + \frac{(\hat{P}_{h,\text{opt}}(f_\ell) - P_h(f_\ell))^2}{\bar{\sigma}_{h,\text{opt}}^2(f_\ell)}, \quad (45)$$

where

$$\frac{\hat{P}_{h,\text{opt}}(f_\ell)}{\bar{\sigma}_{h,\text{opt}}^2(f_\ell)} \equiv \sum_I \frac{2T\delta f \gamma(f_\ell)}{\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell)} \frac{1}{(1+2/N_{\text{avg}})} \frac{2}{T} \frac{1}{M} \sum_{k=\ell-M/2}^{\ell+M/2-1} \text{Re}(\tilde{d}_{1I}^*(f_k)\tilde{d}_{2I}(f_k)), \quad (46)$$

$$\frac{1}{\bar{\sigma}_{h,\text{opt}}^2(f_\ell)} \equiv \sum_I \frac{2T\delta f \gamma^2(f_\ell)}{\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell)} \frac{1}{(1+2/N_{\text{avg}})}. \quad (47)$$

Thus, summing over both I and ℓ ,

$$p(d|\{\bar{P}_{1I}\}, \{\bar{P}_{2I}\}, P_h) = \prod_{\ell} \frac{e^{-2MN_{\text{seg}}}}{(\pi T/2)^{2MN_{\text{seg}}} \prod_I (\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell))^M} \exp\left[\frac{\hat{P}_{h,\text{opt}}^2(f_\ell)}{2\bar{\sigma}_{h,\text{opt}}^2(f_\ell)}\right] \exp\left[-\frac{(\hat{P}_{h,\text{opt}}(f_\ell) - P_h(f_\ell))^2}{2\bar{\sigma}_{h,\text{opt}}^2(f_\ell)}\right]. \quad (48)$$

This is our main result. Note that it has the same basic form as (29), which we derived for the white signal + noise case. Thus, given some choice for the prior probability $p(P_h(f_\ell))$, the posterior distribution for $P_h(f_\ell)$ is given by

$$p(P_h(f_\ell)|d, \{\bar{P}_{1I}(f_\ell)\}, \{\bar{P}_{2I}(f_\ell)\}) \propto \exp\left[-\frac{(\hat{P}_{h,\text{opt}}(f_\ell) - P_h(f_\ell))^2}{2\bar{\sigma}_{h,\text{opt}}^2(f_\ell)}\right] p(P_h(f_\ell)). \quad (49)$$

This expression for the posterior shows that $\hat{P}_{h,\text{opt}}(f_\ell)$ and $\bar{\sigma}_{h,\text{opt}}^2(f_\ell)$ given by (46) and (47) are sufficient statistics for this colored signal + noise, nonstationary analysis, assuming weak signals and measured estimates of the autocorrelated power in two detectors for each data segment.¹

¹One might worry that the frequency-dependent factors multiplying the last exponential in (48) will affect the calculation of the posterior distribution for $P_h(f_\ell)$. But this is not the case since these factors can be lumped together as a single *frequency-independent* normalization factor that does not depend on $P_h(f_\ell)$. This is a consequence of the associative property $\prod_{\ell} (a_{\ell} b_{\ell}) = (\prod_{\ell} a_{\ell}) (\prod_{\ell} b_{\ell})$.




FIG. 2. Recovered posterior distributions for the amplitude of the GWB as obtained from the full and reduced versions of the Bayesian likelihood functions appropriate for both white and colored data. The simulated data consisted of a white signal and white detector noise. (As with Fig. 1, this plot is based on a single realization of the data and cannot be used to infer the presence or absence of a systematic bias. See Figs. 4 and 5 for further discussion.)

In Fig. 2, we compare recoveries of the amplitude of the GWB using the full and reduced versions of the Bayesian likelihood functions appropriate for both white and colored data. The simulated data are the same as that from the previous section, consisting of a white GWB signal injected into white detector noise, for two coincident and coaligned detectors. For the reduced likelihood functions, the detector autocorrelated power were estimated from an additional, adjacent, simulated data segment. We see that all analyses agree very well, demonstrating that the mathematical derivations capture the behavior of the fully Bayesian analysis in practice.

In Fig. 3, we compare recoveries of the GWB amplitude and spectral index using the full and reduced versions of the Bayesian likelihood function appropriate for a colored signal + noise model. For the full likelihood, we parametrize both the GWB signal and detector noise as power laws of the form

$$P_h(f) = A \left(\frac{f}{f_{\text{ref}}} \right)^\beta, \quad \text{etc.}, \quad (50)$$




FIG. 3. Recovered posterior distributions obtained from the full and reduced likelihood functions for a colored signal + noise model. The simulated data consisted of a white GWB signal injected into power-law detector noise. The solid red lines show the injected values of the GWB amplitude and spectral index.

where A is the amplitude, β is the spectral index, and f_{ref} is a reference frequency. (The amplitude and spectral indices for the detector noise will differ, in general, from those for the GWB.) For this simulation the injected noise power spectra have $A_{n_1} = A_{n_2} = 0.125$, and $\beta_1 = \beta_2 = 0.5$, while the injected signal has $A = A_{n_1}/2$, $\beta = 0$ (i.e., it is white). For the reduced likelihood function, the detector autocorrelated power were estimated from coarse-grained power spectral density estimators applied to an additional, adjacent simulated data segment.

Finally, we note that by using sufficient statistics, one improves the computational efficiency of a search relative to a fully Bayesian analysis that works with the raw *uncombined* time-domain (or frequency-domain) data samples. For example, for a stationary GWB, one can reduce the number of likelihood evaluations by a factor of MN_{seg} , where M , which is of order 10 or larger, corresponds to coarse-graining, and $N_{\text{seg}} \sim 10^5$ comes from averaging six months of data split into approximately one-minute data segments. One simply precomputes the sufficient statistic data combinations (46) and (47), and performs Bayesian inference (e.g., Markov-chain Monte Carlo (MCMC) sampling) for the corresponding likelihood function (48).

D. Connection to LIGO's hybrid frequentist-Bayesian analysis

We can use the reduced form of the likelihood (48) to search for GWB signals described by the power spectral density $P_h(f)$. But to make connection with the literature on GWBs, we should parametrize the background in terms of the (dimensionless) energy density spectrum [20]:

$$\Omega_{\text{gw}}(f) \equiv \frac{1}{\rho_c} \frac{d\rho_{\text{gw}}}{d \ln f}, \quad (51)$$

where ρ_c is the energy density needed to close the universe. Then

$$P_h(f) = \frac{3H_0^2}{10\pi^2} \frac{\Omega_{\text{gw}}(f)}{f^3}, \quad (52)$$

which differs from the one-sided strain spectral density $S_h(f) = (3H_0^2/2\pi^2)\Omega_{\text{gw}}(f)/f^3$ by a factor of 1/5, since we are interested here in the strain response of a laser interferometer with a 90° opening angle between the arms [5]. Thus, we can rewrite the likelihood function (48) in terms of $\Omega_{\text{gw}}(f)$ and its optimal estimator $\hat{\Omega}_{\text{gw}}(f)$ as

$$p(d|\{\bar{P}_{1l}\}, \{\bar{P}_{2l}\}, \Omega_{\text{gw}}) = \prod_{\ell} \frac{e^{-2MN_{\text{seg}}}}{(\pi T/2)^{2MN_{\text{seg}}}} \frac{1}{\prod_l (\bar{P}_{1l}(f_\ell) \bar{P}_{2l}(f_\ell))^M} \exp \left[\frac{\hat{\Omega}_{\text{gw}}^2(f_\ell)}{2\bar{\sigma}_{\text{gw}}^2(f_\ell)} \right] \exp \left[-\frac{(\hat{\Omega}_{\text{gw}}(f_\ell) - \Omega_{\text{gw}}(f_\ell))^2}{2\bar{\sigma}_{\text{gw}}^2(f_\ell)} \right], \quad (53)$$

where

$$\frac{\hat{\Omega}_{\text{gw}}(f_\ell)}{\bar{\sigma}_{\text{gw}}^2(f_\ell)} \equiv \sum_I \frac{2T\delta f}{\bar{P}_{1l}(f_\ell) \bar{P}_{2l}(f_\ell)} \left(\frac{3H_0^2 \gamma(f_\ell)}{10\pi^2 f_\ell^3} \right) \frac{1}{(1+2/N_{\text{avg}})} \frac{2}{TM} \sum_{k=\ell-M/2}^{\ell+M/2-1} \text{Re}(\tilde{d}_{1l}^*(f_k) \tilde{d}_{2l}(f_k)), \quad (54)$$

$$\frac{1}{\bar{\sigma}_{\text{gw}}^2(f_\ell)} \equiv \sum_I \frac{2T\delta f}{\bar{P}_{1l}(f_\ell) \bar{P}_{2l}(f_\ell)} \left(\frac{3H_0^2 \gamma(f_\ell)}{10\pi^2 f_\ell^3} \right)^2 \frac{1}{(1+2/N_{\text{avg}})}. \quad (55)$$

This is the form of the likelihood function that you will find in the LIGO-Virgo GWB literature, e.g., Ref. [6], which serves as the starting point for subsequent Bayesian parameter estimation analyses.

We can go one step further if we fix the spectral shape of the GWB and focus attention on estimating only its amplitude at some reference frequency f_{ref} , where we normalize the spectral shape to have unit value. (Typically $f_{\text{ref}} = 25$ Hz for LIGO-Virgo stochastic analyses.) For example, for a power-law background with spectral index α , we have

$$\Omega_{\text{gw}}(f) \equiv \Omega_\alpha \left(\frac{f}{f_{\text{ref}}} \right)^\alpha \quad (56)$$

(note that there is no implied sum over α in the above equation). The spectral index of $\Omega_{\text{gw}}(f)$ is related to the spectral index of $P_h(f)$ defined in the previous section by $\alpha = \beta + 3$. Then we can perform the product over frequencies f_ℓ , again using identity (A1) from Appendix A to do the relevant sum of the argument of the exponential. This yields a likelihood function and posterior distribution for the amplitude Ω_α that are both proportional to

$$\exp \left[-\frac{(\hat{\Omega}_\alpha - \Omega_\alpha)^2}{2\bar{\sigma}_\alpha^2} \right], \quad (57)$$

where

$$\frac{\hat{\Omega}_\alpha}{\bar{\sigma}_\alpha^2} \equiv \sum_I \frac{2T\delta f}{\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell)} \left(\frac{3H_0^2 \gamma(f_\ell)}{10\pi^2 f_{\text{ref}}^3} \left(\frac{f_\ell}{f_{\text{ref}}} \right)^{\alpha-3} \right) \frac{1}{(1+2/N_{\text{avg}})} \frac{2}{T} \frac{1}{M} \sum_{k=\ell-M/2}^{\ell+M/2-1} \text{Re}(\tilde{d}_{1I}^*(f_k)\tilde{d}_{2I}(f_k)), \quad (58)$$

$$\frac{1}{\bar{\sigma}_\alpha^2} \equiv \sum_I \frac{2T\delta f}{\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell)} \left(\frac{3H_0^2 \gamma(f_\ell)}{10\pi^2 f_{\text{ref}}^3} \left(\frac{f_\ell}{f_{\text{ref}}} \right)^{\alpha-3} \right)^2 \frac{1}{(1+2/N_{\text{avg}})}. \quad (59)$$

Note that the factor

$$Q(f) \equiv \frac{3H_0^2 \gamma(f_\ell)}{10\pi^2 f_{\text{ref}}^3} \left(\frac{f_\ell}{f_{\text{ref}}} \right)^{\alpha-3} \frac{1}{\bar{P}_{1I}(f_\ell)\bar{P}_{2I}(f_\ell)}, \quad (60)$$

which multiplies the correlated data in (58), is proportional to the standard expression for the *optimal filter* in the frequency domain; see, e.g., Refs. [5,20].

E. Rigorous comparison

As a check on the results of the previous section, we produce several pp plots [24] to verify that the LIGO-Virgo hybrid frequentist-Bayesian analysis has good statistical coverage.

To generate a pp plot, we first perform $N = 300$ injections and recover the injection with different likelihood functions, as described below. For each recovery, we record the percentile of the posterior distribution at which the injected value lies. We then plot the cumulative distribution function of the percentiles, along with a 90% credible interval determined using order statistics showing the expected range of the cumulative distribution for each percentile value. If the methods we use are unbiased, then the cumulative distribution should lie within the 90% credible interval, close to a diagonal line. Deviations from this line indicate poor coverage, showing the method is biased.

First, we consider a set of white signal + noise injections. We generate 300 noise and signal realizations, and we recover the signal amplitude using the full Bayesian likelihood, reduced likelihood for a white signal + noise model, and also the LIGO-Virgo hybrid frequentist-Bayesian analysis. As emphasized previously, for the reduced likelihood, it is important to use data segments different from the analysis segment to estimate the autocorrelated power in the two detectors, in order to avoid a bias in the recovered parameters. To illustrate the importance of this point, we recover the signal using the reduced likelihood analysis in two different ways: (i) using data segments different from the analysis segment to estimate the autocorrelated power, and (ii) using the same data segment as the analysis segment for estimating the autocorrelated power. We show the results of those analyses in Fig. 4. The analysis using the same data segment as the analysis segment to estimate the autocorrelated power is clearly biased. On the other hand, the full likelihood analysis, the reduced likelihood analysis, and the

LIGO-Virgo stochastic analysis, the latter two of which use data segments different from the analysis segment to estimate the autocorrelated power, all show good coverage. These analyses are not identical, however, due to different conditioning of the data. In particular, the LIGO-Virgo analysis computes the cross-correlation using 50% overlapping Hann windows and estimates the autocorrelated power by averaging Welch estimators from adjacent data segments. In contrast, for our simple reduced likelihood analysis, we do not window the data (which is okay for white data) and use a single additional, adjacent, data segment for estimating the autocorrelated power.

Second, we perform a more realistic set of colored signal + noise injections, and recover the signal with the




FIG. 4. White signal + noise analysis. The pp plot compares recoveries of the amplitude by the full likelihood, reduced likelihood, and LIGO-Virgo stochastic analyses, showing that they both have good Bayesian coverage. The latter two analyses are not identical because of different choices made in conditioning the data. The dotted line shows the bias obtained when using the reduced likelihood analysis if the same data segment as the analysis segment is used to estimate the autocorrelated power in the two detectors.




FIG. 5. Colored signal + noise analysis. In the top panel, we show the injected noise power spectrum $P_n(f)$ (the same for both detectors) and the injected GWB power spectrum $P_h(f)$, along with the optimal estimator $\hat{P}_{h,\text{opt}}(f)$ and its uncertainty $\hat{\sigma}_{h,\text{opt}}(f)$ for one segment in one realization. We rescale the optimal estimator and its uncertainty by a factor of $\sqrt{2T\Delta f}$, where T is the observation time and Δf is the frequency resolution, so that they can be directly compared with the injected autopower and cross power. In the bottom panel, we show a pp plot generated by performing 300 strong-signal and 300 weak-signal injections and recoveries. We see that when the weak-signal approximation is satisfied, the LIGO-Virgo stochastic analysis has excellent Bayesian coverage. Outside of the weak-signal approximation, the coverage is less good, as expected.

LIGO-Virgo stochastic analysis. In these injections, the theoretical noise power spectra are kept the same, although the noise realization changes in every injection. The amplitude of $P_h(f)$ at a reference frequency of 1 Hz is fixed, but we draw the spectral index from a uniform prior probability distribution ranging from -3 to 3 . We use the same prior on the spectral index to perform the recovery. The realization of the GWB signal changes in every injection. We perform two sets of 300 injections. In the first set, the amplitude of $P_h(f)$ is set to be $1/8$ of the amplitude of $P_n(f)$ at the reference frequency in each segment, which is consistent with the weak-signal

approximation. In the second set, the amplitude of the signal is set to 2 times the amplitude of the noise at the reference frequency in each segment, which violates the weak-signal approximation.

In the top panel of Fig. 5, we show the spectrum associated with one injection in the weak-signal regime, along with the injected noise and GWB signal spectra. In the bottom panel, we show the results of the pp plot analysis for both strong- and weak-signal amplitudes. The recovery of both the amplitude and spectral index for weak signals lies within the 90% uncertainty, demonstrating good Bayesian coverage. On the other hand, the recovery of the amplitude and spectral index for strong signals lies outside the 90% region. This is expected, given that strong signals violate the weak-signal approximation, which was used to derive the approximate equivalence of the LIGO-Virgo stochastic analysis and the fully Bayesian approach.

V. DISCUSSION

In this paper, we have shown in what sense LIGO-Virgo's hybrid frequentist-Bayesian cross-correlation analysis for a GWB is equivalent to a fully Bayesian search. The main result was our argument that for a reduced signal + noise model consisting of weak signals and estimated autocorrelated power spectra, the frequency integrand of the cross-correlation statistic and its variance are sufficient statistics for the recovery of the GWB. This means that the posterior distributions of the recovered spectrum of the GWB (e.g., its amplitude and spectral index) will agree for the hybrid frequentist-Bayesian analysis and the fully Bayesian analysis in the context of this reduced signal + noise model. The results of our analyses on simulated data are consistent with those found in [30], which describes a fully Bayesian implementation of LIGO's stochastic search that can estimate the presence of a primordial GWB in the presence of an astrophysical foreground.

We note that a similar hybrid frequentist-Bayesian analysis is also being used for pulsar timing array searches for a GWB produced by inspiraling supermassive black hole binaries associated with galaxy mergers. Here a *noise-marginalized* version of the cross-correlation statistic [31,32] is used to avoid strong covariances that exist between individual estimates of pulsar red noise parameters and the amplitude of the GWB. Noise marginalization is performed by drawing values of the pulsars' red noise parameters from posterior distributions that were generated by an earlier Bayesian analysis which jointly estimates the pulsar's red noise parameters and that of a common red process (e.g., the autocorrelated power due to the GWB), which may exceed the noise at the lowest observed frequencies. This leads to a more accurate recovery of the amplitude of the GWB at relatively little computational cost to the cross-correlation statistic analysis.

A similar hierarchical approach, which first estimates the autocorrelated component of the background before looking for evidence of cross-correlations, would most likely also be needed for analyzing data from proposed *third generation* (3G) ground-based detectors, such as Cosmic Explorer [33] and Einstein Telescope [34]. The expected sensitivity of these 3G detectors is such that the weak-signal approximation in a single segment of data is no longer valid, at least for the rate of binary black hole mergers that advanced LIGO and Virgo are currently observing [35].² As such, simply estimating the autocorrelated power spectra without assigning some portion of it to the GWB would lead to biased estimates of the amplitude of the background, similar to what is seen in pulsar timing analyses [31].

It may be possible to extend the results of our paper to other signal + noise models, such as a nonstationary “popcornlike” background produced, e.g., by stellar-mass binary black hole mergers in the LIGO frequency band [29]. For such a background, one would need to use a *mixture* signal + noise prior to model the nonstationarity, which amounts to postulating the presence of a correlated GWB signal in a certain fraction ξ of the data segments and just noise for the remaining segments, with fraction $1 - \xi$ [29]. The detection of a GWB in this scenario would amount to a posterior distribution for ξ strongly peaked away from zero. By focusing on a GWB signal model that consists of just excess correlation in a certain fraction of data segments (as opposed to marginalizing over the parameters of potential binary black hole mergers [29]), one should be able to implement a computationally efficient and robust (albeit suboptimal) search for a nonstationary GWB.

ACKNOWLEDGMENTS

We thank Stephen Taylor for useful discussions about Bayesian analyses and noise-marginalized statistics in the context of pulsar timing arrays. We also thank Colm Talbot and Sylvia Biscoveanu for conversations at the early stages of this work, regarding fully Bayesian searches for stochastic backgrounds using LIGO-Virgo data, and Vuk Mandic for helpful feedback on a draft of this paper. J. D. R. thanks Tom Callister, Michael Coughlin, Joey Key, and Tyson Littenberg for many useful discussions, which took place several years ago and which eventually led to this project, albeit in a slightly different form than originally planned. J. D. R. also acknowledges support of start-up funds from Texas Tech University. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science

²Note, however, that it may be possible to subtract out the loud resolvable signals, bringing one back to the weak-signal regime [36,37], or else simultaneously fit for the presence of a GWB and resolvable signals [30].

Foundation Grants No. PHY-0757058 and No. PHY-0823459. Numerical computations were performed with the NumPy [38], SciPy [39], and BILBY [40] libraries, and plots with the Matplotlib [41] library.

APPENDIX A: USEFUL IDENTITY

A useful identity that appears in several calculations in the paper is

$$\sum_i \frac{(x_i - \mu)^2}{\sigma_i^2} = \sum_i \frac{x_i^2}{\sigma_i^2} - \frac{x_{\text{opt}}^2}{\sigma_{\text{opt}}^2} + \frac{(x_{\text{opt}} - \mu)^2}{\sigma_{\text{opt}}^2}, \quad (\text{A1})$$

where x_{opt} and σ_{opt}^2 are defined by

$$\frac{x_{\text{opt}}}{\sigma_{\text{opt}}^2} \equiv \sum_i \frac{x_i}{\sigma_i^2}, \quad \frac{1}{\sigma_{\text{opt}}^2} \equiv \sum_i \frac{1}{\sigma_i^2}. \quad (\text{A2})$$

One can think of x_i and σ_i , where $i = 1, 2, \dots, N$, as a set of N independent measurements and error bars of the quantity μ , whose value is to be determined from the measured data. It is a well-known result that x_{opt} defined above is the *minimal-variance unbiased estimator* of μ with variance σ_{opt}^2 . A proof of (A1) is the following:

$$\begin{aligned} \sum_i \frac{(x_i - \mu)^2}{\sigma_i^2} &= \sum_i \frac{x_i^2}{\sigma_i^2} - 2\mu \sum_i \frac{x_i}{\sigma_i^2} + \mu^2 \sum_i \frac{1}{\sigma_i^2} \\ &= \sum_i \frac{x_i^2}{\sigma_i^2} - 2\mu \frac{x_{\text{opt}}}{\sigma_{\text{opt}}^2} + \frac{\mu^2}{\sigma_{\text{opt}}^2} \\ &= \sum_i \frac{x_i^2}{\sigma_i^2} - \frac{x_{\text{opt}}^2}{\sigma_{\text{opt}}^2} + \frac{(x_{\text{opt}} - \mu)^2}{\sigma_{\text{opt}}^2}, \end{aligned} \quad (\text{A3})$$

where we completed the square in x_{opt} and μ to get the last equality.

APPENDIX B: UNCERTAINTIES IN POWER SPECTRUM ESTIMATES

In this Appendix, we provide a brief summary of uncertainties in power spectrum estimates. The final result of this analysis justifies the inclusion of the factor $(1 + 2/N_{\text{avg}})$ in the expressions for the reduced likelihood functions for both the white and colored signal + noise models. Our presentation follows that of an unpublished internal LIGO technical note by Warren Anderson (May 25, 2004).

To simplify the notation a bit, we will use \bar{P}_1, \bar{P}_2 to denote two power spectrum estimators, representing either the autocorrelations $\bar{\sigma}_1^2, \bar{\sigma}_2^2$ for the white signal + noise model or the autocorrelated power spectra $\bar{P}_{1l}(f_\ell), \bar{P}_{2l}(f_\ell)$ for the colored signal + noise model. The number of averages used in the construction of the power spectrum estimators is denoted by N_{avg} , which is proportional to the

number of data samples N for the white signal + noise models, or the number of frequency bins M averaged over for coarse graining in a Welch power spectrum estimate [22] for the colored signal + noise models.

Since \bar{P}_1 and \bar{P}_2 are unbiased estimators of P_1 and P_2 , we can write

$$\begin{aligned}\bar{P}_1 &= P_1 + \delta\bar{P}_1, \\ \bar{P}_2 &= P_2 + \delta\bar{P}_2,\end{aligned}\quad (\text{B1})$$

where

$$\langle\delta\bar{P}_1\rangle = \langle\delta\bar{P}_2\rangle = 0. \quad (\text{B2})$$

The variance and covariance of the power spectrum estimators are given by the quadratic expectation values

$$\begin{aligned}\langle(\delta\bar{P}_1)^2\rangle &= \langle\bar{P}_1^2\rangle - P_1^2 \equiv \text{var}(\bar{P}_1), \\ \langle(\delta\bar{P}_2)^2\rangle &= \langle\bar{P}_2^2\rangle - P_2^2 \equiv \text{var}(\bar{P}_2), \\ \langle\delta\bar{P}_1\delta\bar{P}_2\rangle &= \langle\bar{P}_1\bar{P}_2\rangle - P_1P_2 \equiv \text{cov}(\bar{P}_1\bar{P}_2).\end{aligned}\quad (\text{B3})$$

Explicitly evaluating $\langle\bar{P}_1^2\rangle$, etc., using the identity

$$\langle abcd\rangle = \langle ab\rangle\langle cd\rangle + \langle ac\rangle\langle bd\rangle + \langle ad\rangle\langle bc\rangle \quad (\text{B4})$$

for zero-mean, Gaussian random variables, leads to

$$\begin{aligned}\langle(\delta\bar{P}_1)^2\rangle &= P_1^2/N_{\text{avg}}, \\ \langle(\delta\bar{P}_2)^2\rangle &= P_2^2/N_{\text{avg}}, \\ \langle\delta\bar{P}_1\delta\bar{P}_2\rangle &= P_h^2/N_{\text{avg}}.\end{aligned}\quad (\text{B5})$$

Since the power spectra appear in the full likelihood function via the product of their inverses, $1/\sigma_1^2\sigma_2^2$ or $1/P_{1I}(f_\ell)P_{2I}(f_\ell)$, we need to calculate the expectation value $\langle 1/\bar{P}_1\bar{P}_2\rangle$. So making a Taylor series expansion

$$\frac{1}{\bar{P}} = \frac{1}{P + \delta\bar{P}} = \frac{1}{P} \left(1 - \frac{\delta\bar{P}}{P} + \frac{(\delta\bar{P})^2}{P^2} - \dots \right) \quad (\text{B6})$$

for both $1/\bar{P}_1$ and $1/\bar{P}_2$, it follows that

$$\begin{aligned}\frac{1}{\bar{P}_1\bar{P}_2} &= \frac{1}{P_1P_2} \left(1 - \frac{\delta\bar{P}_1}{P_1} + \frac{(\delta\bar{P}_1)^2}{P_1^2} - \dots \right) \left(1 - \frac{\delta\bar{P}_2}{P_2} + \frac{(\delta\bar{P}_2)^2}{P_2^2} - \dots \right) \\ &= \frac{1}{P_1P_2} \left(1 - \frac{\delta\bar{P}_1}{P_1} - \frac{\delta\bar{P}_2}{P_2} + \frac{(\delta\bar{P}_1)^2}{P_1^2} + \frac{(\delta\bar{P}_2)^2}{P_2^2} + \frac{\delta\bar{P}_1\delta\bar{P}_2}{P_1P_2} - \dots \right).\end{aligned}\quad (\text{B7})$$

Taking the expectation value of both sides of the above expression, we find

$$\left\langle \frac{1}{\bar{P}_1\bar{P}_2} \right\rangle = \frac{1}{P_1P_2} \left(1 + \frac{2}{N_{\text{avg}}} + \frac{1}{N_{\text{avg}}} \frac{P_h^2}{P_1P_2} - \dots \right) \simeq \frac{1}{P_1P_2} \left(1 + \frac{2}{N_{\text{avg}}} \right), \quad (\text{B8})$$

where we have ignored cubic and higher-order terms in $\delta\bar{P}_1$ and $\delta\bar{P}_2$, and assumed the weak-signal approximation, $P_h^2/P_1P_2 \ll 1$, to get the last approximate equality. This result shows that $1/\bar{P}_1\bar{P}_2$ is a *biased* estimator of $1/P_1P_2$. Nonetheless, this bias can be removed by simply moving the factor of $(1 + 2/N_{\text{avg}})$ to the left-hand side, so that

$$\frac{1}{\bar{P}_1\bar{P}_2(1 + 2/N_{\text{avg}})} \quad (\text{B9})$$

is a unbiased estimator of $1/P_1P_2$. This is the replacement we make for $1/P_1P_2$ in the reduced likelihood functions in the main text; cf. (23).

APPENDIX C: ALTERNATIVE DERIVATION OF A REDUCED LIKELIHOOD FUNCTION

In this Appendix, we give an alternative derivation of the likelihood function for a reduced signal + noise model, but under different assumptions than those given in the main

text. More specifically, we do not assume here that the GWB signal is weak compared to the detector noise, nor do we use estimators of the autocorrelated power calculated from data segments different from that being analyzed for the signal. Rather we assume the following: (i) the number of data points N (or coarse-grained averages M) for a given data segment I is sufficiently large that one can expand the likelihood function around the maximum-likelihood estimators of σ_1^2 , σ_2^2 , σ_h^2 [or $P_{1I}(f_\ell)$, $P_{2I}(f_\ell)$, $P_h(f_\ell)$] to second order without loss of information; and (ii) the data are informative for the autocorrelated power σ_1^2 , σ_2^2 [or $P_{1I}(f_\ell)$, $P_{2I}(f_\ell)$] allowing us to evaluate the second-order likelihood function at the values of σ_1^2 , σ_2^2 [or $P_{1I}(f_\ell)$, $P_{2I}(f_\ell)$] that maximize the likelihood for *fixed* values of σ_h^2 [or $P_h(f_\ell)$]. For concreteness, we give the derivation here in the context of the white signal + noise model for two coincident and coaligned detectors. But it can easily be extended to the case of colored data with non-stationary noise and a nontrivial overlap function. Our derivation follows that given in [42].

We start with the full likelihood function (22) for the white signal + noise model, which we rewrite here as

$$p(d|\sigma_{n_1}^2, \sigma_{n_2}^2, \sigma_h^2) = \exp\left[-\frac{N}{2}f(\sigma_1^2, \sigma_2^2, \sigma_h^2|d)\right], \quad (\text{C1})$$

where

$$f(\sigma_1^2, \sigma_2^2, \sigma_h^2|d) = 2\ln(2\pi) + \ln(\sigma_1^2\sigma_2^2 - (\sigma_h^2)^2) + \frac{1}{(\sigma_1^2\sigma_2^2 - (\sigma_h^2)^2)}(\hat{\sigma}_1^2\sigma_2^2 + \hat{\sigma}_2^2\sigma_1^2 - 2\hat{\sigma}_h^2\sigma_h^2). \quad (\text{C2})$$

In the above expressions, σ_1^2, σ_2^2 are the total autocorrelated variances in the two detectors [see (18)], and

$$\begin{aligned} \hat{\sigma}_1^2 &\equiv \frac{1}{N} \sum_i d_{1i}^2, \\ \hat{\sigma}_2^2 &\equiv \frac{1}{N} \sum_i d_{2i}^2, \\ \hat{\sigma}_h^2 &\equiv \frac{1}{N} \sum_i d_{1i}d_{2i}, \end{aligned} \quad (\text{C3})$$

are the maximum-likelihood estimators of $\sigma_1^2, \sigma_2^2, \sigma_h^2$.

Assuming that N is sufficiently large, we can Taylor expand f around its maximum-likelihood values ignoring terms higher than second order in the differences $\sigma_1^2 - \hat{\sigma}_1^2, \sigma_2^2 - \hat{\sigma}_2^2, \sigma_h^2 - \hat{\sigma}_h^2$. Doing so gives

$$f(\sigma_1^2, \sigma_2^2, \sigma_h^2|d) \simeq f(\hat{\sigma}_1^2, \hat{\sigma}_2^2, \hat{\sigma}_h^2|d) + \frac{1}{2} \sum_{i,j} \frac{\partial^2 f}{\partial x_i \partial x_j} \Big|_{\text{ML}} (x_i - \hat{x}_i)(x_j - \hat{x}_j), \quad (\text{C4})$$

where $x_i \equiv (\sigma_1^2, \sigma_2^2, \sigma_h^2)$ and

$$\begin{aligned} &\frac{\partial^2 f}{\partial x_i \partial x_j} \Big|_{\text{ML}} \\ &= \frac{1}{(\hat{\sigma}_1^2\hat{\sigma}_2^2 - (\hat{\sigma}_h^2)^2)} \begin{bmatrix} (\hat{\sigma}_2^2)^2 & (\hat{\sigma}_h^2)^2 & -2\hat{\sigma}_2^2\hat{\sigma}_h^2 \\ (\hat{\sigma}_h^2)^2 & (\hat{\sigma}_1^2)^2 & -2\hat{\sigma}_1^2\hat{\sigma}_h^2 \\ -2\hat{\sigma}_2^2\hat{\sigma}_h^2 & -2\hat{\sigma}_1^2\hat{\sigma}_h^2 & 2(\hat{\sigma}_1^2\hat{\sigma}_2^2 + (\hat{\sigma}_h^2)^2) \end{bmatrix}. \end{aligned} \quad (\text{C5})$$

Defining

$$\Gamma_{ij} \equiv \frac{N}{2} \frac{\partial^2 f}{\partial x_i \partial x_j} \Big|_{\text{ML}}, \quad (\text{C6})$$

it follows that the inverse matrix

$$C_{ij} \equiv (\Gamma^{-1})_{ij} = \frac{1}{N} \begin{bmatrix} 2(\hat{\sigma}_1^2)^2 & 2(\hat{\sigma}_h^2)^2 & 2\hat{\sigma}_1^2\hat{\sigma}_2^2 \\ 2(\hat{\sigma}_h^2)^2 & 2(\hat{\sigma}_2^2)^2 & 2\hat{\sigma}_2^2\hat{\sigma}_h^2 \\ 2\hat{\sigma}_1^2\hat{\sigma}_h^2 & 2\hat{\sigma}_2^2\hat{\sigma}_h^2 & \hat{\sigma}_1^2\hat{\sigma}_2^2 + (\hat{\sigma}_h^2)^2 \end{bmatrix} \quad (\text{C7})$$

is an estimator of the covariance matrix of the maximum-likelihood estimators $\hat{\sigma}_1^2, \hat{\sigma}_2^2, \hat{\sigma}_h^2$.

To proceed further, we construct a *reduced* likelihood function by assuming that the data are informative with respect to the detector autocorrelations σ_1^2, σ_2^2 . This means that we can evaluate (C4) at the values of σ_1^2, σ_2^2 that maximize (C4) for *fixed* values of σ_h^2 . So, simultaneously solving the two equations

$$\frac{\partial f}{\partial \sigma_1^2} = 0, \quad \frac{\partial f}{\partial \sigma_2^2} = 0, \quad (\text{C8})$$

for σ_1^2, σ_2^2 , we obtain

$$\begin{aligned} \sigma_1^2 &= \hat{\sigma}_1^2 + \frac{2\hat{\sigma}_1^2\hat{\sigma}_h^2}{(\hat{\sigma}_1^2\hat{\sigma}_2^2 + (\hat{\sigma}_h^2)^2)}(\sigma_h^2 - \hat{\sigma}_h^2), \\ \sigma_2^2 &= \hat{\sigma}_2^2 + \frac{2\hat{\sigma}_2^2\hat{\sigma}_h^2}{(\hat{\sigma}_1^2\hat{\sigma}_2^2 + (\hat{\sigma}_h^2)^2)}(\sigma_h^2 - \hat{\sigma}_h^2). \end{aligned} \quad (\text{C9})$$

Denoting the RHSs of these expressions by $\bar{\sigma}_1^2, \bar{\sigma}_2^2$, it follows that (C4) becomes

$$f(\sigma_1^2, \sigma_2^2, \sigma_h^2|d) \Big|_{\sigma_1^2=\bar{\sigma}_1^2, \sigma_2^2=\bar{\sigma}_2^2} \simeq f(\hat{\sigma}_1^2, \hat{\sigma}_2^2, \hat{\sigma}_h^2|d) + \frac{(\sigma_h^2 - \hat{\sigma}_h^2)^2}{(\hat{\sigma}_1^2\hat{\sigma}_2^2 + (\hat{\sigma}_h^2)^2)}. \quad (\text{C10})$$

The corresponding reduced likelihood function is

$$p(d|\bar{\sigma}_1^2, \bar{\sigma}_2^2, \sigma_h^2) \equiv p(d|\hat{\sigma}_1^2, \hat{\sigma}_2^2, \hat{\sigma}_h^2) \exp\left[-\frac{N}{2} \frac{(\sigma_h^2 - \hat{\sigma}_h^2)^2}{(\hat{\sigma}_1^2\hat{\sigma}_2^2 + (\hat{\sigma}_h^2)^2)}\right]. \quad (\text{C11})$$

Thus, we see that $\hat{\sigma}_h^2$ together with $\hat{\sigma}_1^2, \hat{\sigma}_2^2$ are sufficient statistics for σ_h^2 with variance

$$\text{var}[\hat{\sigma}_h^2] = (\hat{\sigma}_1^2\hat{\sigma}_2^2 + (\hat{\sigma}_h^2)^2)/N. \quad (\text{C12})$$

A couple of remarks are in order:

(i) Nowhere in the above derivation did we assume that the power in the GWB is small compared to the detector noise. Thus, the reduced likelihood function (C11) is valid for arbitrarily large GWB signals, which is relevant, for example, for searches for a GWB using pulsar timing arrays [15,21] or the proposed space-based interferometer LISA [17]. Equation (C12) contains an extra term, $(\hat{\sigma}_h^2)^2$, compared to (28), which takes into account the extra variance

associated with the GWB, over and above that which is already captured in the autocorrelations $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$.

(ii) Although we denoted the right-hand sides of (C9) by $\hat{\sigma}_1^2$, $\hat{\sigma}_2^2$, these expressions are not the same as those used in the main text (see the discussion in Sec. III B), which were estimators of σ_1^2 , σ_2^2 constructed from data segments *different* than that being analyzed for the signal. The use of different data segments to estimate the autocorrelations is necessary for LIGO-Virgo data, since the number of averages used to estimate power spectra is not sufficiently large to beat down the bias that arises from using the same data in both the numerator and denominator of the exponential in (C11). (Recall that for

the colored signal + noise model, the number of averages is proportional to the number of frequency bins M that are averaged together for coarse graining.) For LIGO-Virgo data, one is restricted to N_{avg} of order at most 50, due to the complexity of the detector noise (the need to take the coarse-grained frequency resolution to be $\delta f \sim 0.25$ Hz) and its broad nonstationarity on timescales of order minutes (ignoring shorter timescale instrumental glitches). By using different data segments to estimate the autocorrelated power, the bias goes away in the weak-signal limit, as shown in Fig. 4. But if the number of averages was sufficiently large, then one could use the expression in (C11) as is.

-
- [1] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 115012 (2015).
- [2] F. Acernese *et al.*, *Classical Quantum Gravity* **32**, 024001 (2015).
- [3] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **100**, 061101 (2019).
- [4] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **100**, 062001 (2019).
- [5] J. D. Romano and N. J. Cornish, *Living Rev. Relativity* **20**, 2 (2017).
- [6] V. Mandic, E. Thrane, S. Giampanis, and T. Regimbau, *Phys. Rev. Lett.* **109**, 171102 (2012).
- [7] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
- [8] T. Callister, A. S. Biscoveanu, N. Christensen, M. Isi, A. Matas, O. Minazzoli, T. Regimbau, M. Sakellariadou, J. Tasson, and E. Thrane, *Phys. Rev. X* **7**, 041058 (2017).
- [9] E. Thrane, N. Christensen, and R. Schofield, *Phys. Rev. D* **87**, 123009 (2013).
- [10] E. Thrane, N. Christensen, R. M. S. Schofield, and A. Effler, *Phys. Rev. D* **90**, 023013 (2014).
- [11] M. W. Coughlin *et al.*, *Classical Quantum Gravity* **33**, 224003 (2016).
- [12] M. W. Coughlin *et al.*, *Phys. Rev. D* **97**, 102007 (2018).
- [13] P. M. Meyers, K. Martinovic, N. Christensen, and M. Sakellariadou, *Phys. Rev. D* **102**, 102005 (2020).
- [14] Z. Arzoumanian *et al.* (NANOGrAV Collaboration), *Astrophys. J.* **859**, 47 (2018).
- [15] Z. Arzoumanian *et al.*, *Astrophys. J. Lett.* **905**, L34 (2020).
- [16] J. Baker *et al.*, [arXiv:1907.06482](https://arxiv.org/abs/1907.06482).
- [17] M. R. Adams and N. J. Cornish, *Phys. Rev. D* **89**, 022001 (2014).
- [18] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **69**, 122004 (2004).
- [19] A. Lazzarini, Bias from power spectrum measurement in parameter estimation for the stochastic gravitational wave background, <https://dcc.ligo.org/T040128/> (2004).
- [20] B. Allen and J. D. Romano, *Phys. Rev. D* **59**, 102001 (1999).
- [21] X. Siemens, J. Ellis, F. Jenet, and J. D. Romano, *Classical Quantum Gravity* **30**, 224015 (2013).
- [22] P. Welch, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).
- [23] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), [arXiv:2101.12130](https://arxiv.org/abs/2101.12130).
- [24] S. R. Cook, A. Gelman, and D. B. Rubin, *J. Comput. Graph. Stat.* **15**, 675 (2006).
- [25] N. Christensen, *Phys. Rev. D* **46**, 5250 (1992).
- [26] E. E. Flanagan, *Phys. Rev. D* **48**, 2389 (1993).
- [27] D. Davis *et al.* (LIGO Collaboration), [arXiv:2101.11673](https://arxiv.org/abs/2101.11673).
- [28] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **69**, 122004 (2004).
- [29] R. Smith and E. Thrane, *Phys. Rev. X* **8**, 021019 (2018).
- [30] S. Biscoveanu, C. Talbot, E. Thrane, and R. Smith, *Phys. Rev. Lett.* **125**, 241101 (2020).
- [31] S. J. Vigeland, K. Iso, S. R. Taylor, and J. A. Ellis, *Phys. Rev. D* **98**, 044003 (2018).
- [32] M. Anholm, S. Ballmer, J. D. E. Creighton, L. R. Price, and X. Siemens, *Phys. Rev. D* **79**, 084030 (2009).
- [33] D. Reitze *et al.*, *Bull. Am. Astron. Soc.* **51**, 035 (2019), <https://ui.adsabs.harvard.edu/abs/2019BAAS...51g..35R/abstract>.
- [34] M. Maggiore, C. V. D. Broeck, N. Bartolo, E. Belgacem, D. Bertacca, M. A. Bizouard, M. Branchesi, S. Clesse, S. Foffa, J. Garcia-Bellido *et al.*, *J. Cosmol. Astropart. Phys.* **03** (2020) 050.
- [35] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt *et al.*, *Astrophys. J.* **882**, L24 (2019).
- [36] T. Regimbau, M. Evans, N. Christensen, E. Katsavounidis, B. Sathyaprakash, and S. Vitale, *Phys. Rev. Lett.* **118**, 151105 (2017).
- [37] S. Sachdev, T. Regimbau, and B. Sathyaprakash, *Phys. Rev. D* **102**, 024051 (2020).
- [38] C. R. Harris *et al.*, *Nature (London)* **585** 357 (2020).
- [39] P. Virtanen *et al.* (SciPy 1.0 Contributors), *Nat. Methods* **17**, 261 (2020).
- [40] G. Ashton *et al.*, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [41] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [42] B. Allen, J. D. E. Creighton, E. E. Flanagan, and J. D. Romano, *Phys. Rev. D* **67**, 122002 (2003).