

Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects

Michele Vallisneri

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109, USA
(Received 15 March 2007; published 4 February 2008)

The Fisher-matrix formalism is used routinely in the literature on gravitational-wave detection to characterize the parameter-estimation performance of gravitational-wave measurements, given parameterized models of the waveforms, and assuming detector noise of known colored Gaussian distribution. Unfortunately, the Fisher matrix can be a poor predictor of the amount of information obtained from typical observations, especially for waveforms with several parameters and relatively low expected signal-to-noise ratios (SNR), or for waveforms depending weakly on one or more parameters, when their priors are not taken into proper consideration. In this paper I discuss these pitfalls; show how they occur, even for relatively strong signals, with a commonly used template family for binary-inspiral waveforms; and describe practical recipes to recognize them and cope with them. Specifically, I answer the following questions: (i) What is the significance of (quasi-)singular Fisher matrices, and how must we deal with them? (ii) When is it necessary to take into account prior probability distributions for the source parameters? (iii) When is the signal-to-noise ratio high enough to believe the Fisher-matrix result? In addition, I provide general expressions for the higher-order, beyond-Fisher-matrix terms in the $1/\text{SNR}$ expansions for the expected parameter accuracies.

DOI: [10.1103/PhysRevD.77.042001](https://doi.org/10.1103/PhysRevD.77.042001)

PACS numbers: 04.80.Nn, 02.50.Tt, 95.55.Ym

I. INTRODUCTION

Over the last two decades, the prevailing attitude in the gravitational-wave (GW) source-modeling community has been one of *predata positioning*: in the absence of confirmed detections, the emphasis has been on exploring which astrophysical systems, and which of their properties, would become accessible to GW observations with the sensitivities afforded by planned (or desired) future experiments, with the purpose of committing theoretical effort to the most promising sources, and of directing public advocacy to the most promising detectors. In this positioning and in this exploration, the *expected accuracy* of GW source parameters, as determined from the signals yet to be observed, is often employed as a proxy for the amount of physical information that could be gained from detection campaigns. However, predicting the parameter-estimation performance of future observations is a complex matter, even with the benefit of accurate theoretical descriptions of the expected waveforms and of faithful characterizations of the noise and response of detectors; in practice, the typical source modeler has had much less to go with. The main problem is that there are few analytical tools that can be applied generally to the problem, before resorting to relatively cumbersome numerical simulations that involve multiple explicit realizations of signal-plus-noise data sets.

In the source-modeling community, the analytical tool of choice has been the Fisher information matrix $F_{ij}[h] = (h_i, h_j)$: here $h_i(t)$ is the partial derivative of the gravitational waveform $h(t)$ of interest with respect to the i th source parameter θ_i , and “ (\cdot, \cdot) ” is a signal product weighted by the expected power spectral density of detec-

tor noise, as described in Sec. II B. Now, it is usually claimed that the inverse Fisher matrix $F_{ij}^{-1}[h_0]$ represents the covariance matrix of parameter errors in the parameter-estimation problem for the true signal $h_0(t)$. This statement can be interpreted in three slightly different ways (all correct), which we examine in detail in Sec. II, and preview here:

- (1) The inverse Fisher matrix $F_{ij}^{-1}[h_0]$ is a lower bound (generally known as the *Cramér-Rao bound*) for the error covariance of any *unbiased estimator* of the true source parameters. Thus, it is a *frequentist* error (see Sec. II A): for any *experiment* characterized by the true signal $h_0(t)$ and a certain realization $n(t)$ of detector noise, the parameter estimator $\hat{\theta}$ is a vector function of the total detector output $s = n + h_0$, and $F_{ij}^{-1}[h_0]$ is a lower bound on the covariance (i.e., the fluctuations) of $\hat{\theta}$ in an imaginary infinite sequence of experiments with different realizations of noise. The Cramér-Rao bound is discussed in Sec. II C.
- (2) The inverse Fisher matrix $F_{ij}^{-1}[h_0]$ is the frequentist error covariance for the *maximum-likelihood* (ML) parameter estimator $\hat{\theta}^{\text{ML}}$, assuming Gaussian noise, in the limit of strong signals (i.e., high signal-to-noise ratio SNR) or, equivalently, in the limit in which the waveforms can be considered as linear functions of source parameters. We shall refer to this limit as the *linearized-signal approximation*, or LSA. This well-known result is rederived in Sec. II D.
- (3) The inverse Fisher matrix $F_{ij}^{-1}[h_0]$ represents the covariance (i.e., the multidimensional spread around the mode) of the *posterior probability distribution*

$p(\theta_0|s)$ for the true source parameters θ_0 , as inferred (in *Bayesian* fashion) from a *single* experiment with true signal h_0 , assuming Gaussian noise, in the high-SNR limit (or in the LSA), and in the case where any *prior probabilities* for the parameters are constant over the parameter range of interest. Properly speaking, the inverse Fisher matrix is a measure of uncertainty rather than error, since in any experiment the mode will be displaced from the true parameters by an unknown amount due to noise.¹ See Sec. IIE for a rederivation of this result.

As pointed out by Jaynes [1], while the numerical identity of these three different errorlike quantities has given rise to much confusion, it arises almost trivially from the fact that in a neighborhood of its maximum, the signal likelihood $p(s|\theta_0)$ is approximated by a normal probability distribution with covariance F_{ij}^{-1} . In this paper, I argue that the Cramér-Rao bound is seldom useful in the work of GW analysts (Sec. IIC), and while the high-SNR/LSA frequentist and Bayesian results are legitimate, they raise the question of whether the signals of interest are strong (or linear) enough to warrant the limit, and of what happens if they are not. In addition, if we possess significant information about the prior distributions (or even the allowed ranges) of source parameters, it is really only in the Bayesian framework that we can fold this information reliably into the Fisher result (Sec. IID).

Thus, I recommend the Bayesian viewpoint as the most fruitful way of thinking about the Fisher-matrix result (although I will also derive parallel results from the frequentist viewpoint). Of course, the study of Bayesian inference for GW parameter-estimation problems need not stop at the leading-order (Fisher-matrix) expression for the posterior likelihood: Markov Chain Monte Carlo (MCMC) algorithms [2] can provide very reliable results, immune from any considerations about signal strength, but they require a significant investment of time to implement them, and of computational resources to run them, since they necessarily involve explicit realizations of the noise. More rigorous Bayesian bounds (such as the Weiss-Weinstein and Ziv-Zakai bounds examined by Nicholson and Vecchio [3]) can also be derived, but they require a careful appraisal of the nonlocal structure of the likelihood function.

By contrast, the Fisher-matrix formalism is singularly economical, and it seems clear that it will always be the first recourse of the GW data analyst. To use it reliably, however, we must understand the limits of its applicability. The purpose of this paper is to explore these limits. I do so by providing practical solutions to three issues that were already raised in the seminal treatments of GW detection

by Finn [4] and by Cutler and Flanagan [5], but that seem to have been almost ignored after that:

- (1) What is the significance of the singular or ill-conditioned Fisher matrices that often appear in estimation problems with several source parameters, and how do we deal with them? Can we still believe the Fisher result in those cases? (See Sec. IV.)
- (2) When is it necessary to take into account the prior probability distributions for the parameters, even if specified trivially by their allowed ranges? (See Sec. V.)
- (3) When is the high-SNR/LSA approximation warranted? (As anticipated above, the high-SNR limit is *equivalent* to the LSA, as we shall show in Secs. IID and IIE.) That is, how strong a signal will we need to measure if we are to believe the Fisher-matrix result for its uncertainty? (See Sec. VI.)

Last, I discuss the extension of the LSA beyond the leading order, in both the frequentist and Bayesian parameter-estimation frameworks (Sec. VII), in a form that the adventurous GW analyst can use to test the reliability of the Fisher result (but higher-order derivatives and many-indexed expressions start to mount rapidly, even at the next-to-leading order). By contrast, I do not address the reduction in parameter-estimation accuracy due to the presence of secondary maxima in the likelihood function, as noticed [6] and carefully modeled [7] by Balasubramanian and colleagues in their extensive Monte Carlo simulations of ML estimation for inspiraling binaries using Newtonian and first post-Newtonian waveforms.

My treatment follows Refs. [4,5], as well as the classic texts on the statistical analysis of noisy data (e.g., Refs. [8–10]). I am indebted to Jaynes and Bretthorst [1,11] for their enlightening, if occasionally blunt, perspective on frequentist and Bayesian parameter estimation. The reader already familiar with the standing of the Fisher-matrix formalism in the frequentist and Bayesian frameworks can skip Secs. IIA (a refresher on the difference between the frequentist and Bayesian viewpoints) and IIC, IID, and IIE (a pedagogical derivation of the three approaches to the inverse-Fisher-matrix result that were introduced at the beginning of this section), and move directly to discussion of the three *issues* in Secs. IV, V, and VI, and to the higher-order formalism in Sec. VII, referring back to Sec. II as needed to establish notation. Whenever my discussion requires a practical example, I consider signals from inspiraling binaries of two black holes, both of mass $10M_\odot$, as described by the restricted post-Newtonian approximation for adiabatic, circular inspirals (see Sec. III); in my examples, I assume detection and parameter estimation are performed on Initial-LIGO [12] data, and I adopt the LIGO noise curve of Table IV in Ref. [13]. Throughout, I use geometric units; I assume the Einstein summation conven-

¹In the high-SNR/LSA limit with negligible priors, the posterior probability mode, seen as a frequentist statistic, coincides with the ML estimator; thus its fluctuations are again described by the inverse Fisher matrix.

tion for repeated indices; and I do not distinguish between covariant and contravariant indices, except in Sec. VII.

II. THREE ROADS TO THE FISHER MATRIX

In this section I discuss the “three roads” to the inverse Fisher matrix as a measure of uncertainty for GW observations: the Cramér-Rao bound (Sec. II C), the high-SNR/LSA limit for the frequentist covariance of the ML estimator (Sec. II D), and the high-SNR/LSA limit for the single-experiment covariance of the Bayesian posterior distribution (Sec. II E). Sections II A and II B are refreshers about frequentist and Bayesian parameter estimation, and about the analytical expression for the likelihood of GW signals in Gaussian noise.

A. A refresher on the frequentist and Bayesian frameworks

The *frequentist* (or orthodox) approach to parameter estimation for GW signals can be summed up as follows:

- (1) We are given the detector data s and we take it to consist of the true signal $h_0 = h(\theta_0)$ (where θ_0 is the vector of the *true system parameters*) plus additive noise n .
- (2) We select a *point estimator* $\hat{\theta}(s)$: that is, a vector function of detector data that (it is hoped) approximates the true values of source parameters, except for the statistical error due to the presence of noise. One important example of point estimator is the ML estimator $\hat{\theta}^{\text{ML}}$, which maximizes the *likelihood* $p(s|\theta)$ of observing the measured data s given a value θ of the true parameters. For additive noise, this likelihood coincides with the probability of a noise realization $n = s - h(\theta)$, and for Gaussian noise it is given below in Sec. II B.
- (3) We characterize statistical error as the fluctuations of $\hat{\theta}(s)$, computed over a very long series of independent *experiments* where the source parameters are kept fixed, while detector noise n is sampled from its assumed probability distribution (often called the *sampling* distribution).

The estimator $\hat{\theta}$ is usually chosen according to one or more criteria of optimality: for instance, *unbiasedness* requires that $\langle \hat{\theta}(s) \rangle_n$ (the average of the estimator over the noise probability distribution) be equal to θ_0 .

A rather different approach is that of *Bayesian* inference:

- (1) We do not assume a true value of the system parameters, but we posit their *prior probability distribution* $p(\theta)$.
- (2) Given the data s , we do not compute estimators, but rather the full *posterior probability distribution* $p(\theta|s)$, using *Bayes’ theorem* $p(\theta|s) = p(s|\theta) \times p(\theta)/p(s)$, where $p(s) = \int p(s|\theta)p(\theta)d\theta$.
- (3) We characterize statistical error *in a single experiment* by the spread of the posterior distribution $p(\theta|s)$.

The differences between the frequentist and Bayesian approaches are not only mathematical, but also epistemic: as their name indicates, “frequentists” view probabilities essentially as the relative frequencies of outcomes in repeated experiments, while “Bayesians” view them as subjective² indices of certainty for alternative propositions. For an introduction to the contrasting views, I refer the reader to the excellent treatise (very partial to the Bayesian world view) by Jaynes [1], and to Ref. [5] for a more GW-detection-oriented discussion.

Once actual detections are made, the Bayesian approach of computing posterior probability distributions for the signal parameters *given the observed data* seems more powerful than the frequentist usage of somewhat arbitrary point estimators; the latter will always result in throwing away useful information, unless the chosen estimators are *sufficient statistics* (i.e., unless the likelihood depends on the data only through the estimators). As for statistical error, it seems preferable to characterize it from the data *we have* (actually, from the posterior distributions that we infer from that data), rather than from the data *we could have obtained* (i.e., from the sampling distribution of estimators in a hypothetical *ensemble* of experiments).

As Cutler and Flanagan [5] point out, however, it is in the current *predata* regime that we seek to compute expected parameter accuracies; in the absence of actual confirmed-detection data sets, it seems acceptable to consider *ensembles* of possible parameter-estimation experiments, and to use frequentist statistical error as an inverse measure of potential physical insight. The best solution, bridging the two approaches, would undoubtedly be to examine the frequentist distribution of some definite measure of Bayesian statistical error; unfortunately, such a hybrid study is generally unfeasible, given the considerable computational requirements of even single-data set Bayesian analyses.

B. Likelihood for GW signals in Gaussian noise

Under the assumption of stationary and Gaussian detector noise, the likelihood $\log p(s|\theta)$ can be obtained very simply from a noise-weighted inner product of the detector output and of the signal $h(\theta)$ (see for instance Eq. (2.3) in Ref. [5]):

$$p(s|\theta) \propto e^{-(s-h(\theta), s-h(\theta))/2}; \quad (1)$$

the weighting is performed with respect to the expected power spectral density of detector noise by defining the noise-weighted inner product of two real-valued signals as

$$(h, g) = 4\text{Re} \int_0^{+\infty} \frac{\tilde{h}(f)^* \tilde{g}(f)}{S_n(f)} df, \quad (2)$$

²Only in the sense that subjects with different prior assumptions could come to different conclusions after seeing the same data; indeed, Bayesian statistics describes how prior assumptions become *deterministically* modified by the observation of data.

where $\tilde{h}(f)$ and $\tilde{g}(f)$ are the Fourier transforms of $h(t)$ and $g(t)$, “ $*$ ” denotes complex conjugation, and $S_n(f)$ is the one-sided power spectral density of the noise. From the definition of $S_n(f)$ as $\langle \tilde{n}^*(f)\tilde{n}(f') \rangle_n = \frac{1}{2}S_n(|f|)\delta(f-f')$, we get the useful property

$$\langle (h, n)(n, g) \rangle_n = (h, g), \quad (3)$$

where again “ $\langle \cdot \rangle_n$ ” denotes averaging over the probability distribution of the noise.

C. First road: Derivation and critique of the Cramér-Rao bound

The derivation in this section is inspired by the treatment of Ref. [1] (p. 518), and it is given for simplicity in the case of one source parameter. We wish to pose a bound on the frequentist estimator variance

$$\text{var } \hat{\theta} = \langle (\hat{\theta}(s) - \langle \hat{\theta}(s) \rangle_n)^2 \rangle_n: \quad (4)$$

to do this, we consider the ensemble product

$$\langle u(s), v(s) \rangle_n = \int u(s)v(s)p(s|\theta_0)ds, \quad (5)$$

where $p(s|\theta_0)$ is the likelihood of observing the detector output s given the true source parameter θ_0 , or equivalently the likelihood of observing the noise realization $n = s - h_0$. Setting $v(s) = \hat{\theta}(s) - \langle \hat{\theta}(s) \rangle_n$, we obtain a bound on $\langle v, v \rangle_n \equiv \text{var } \hat{\theta}$ from the Schwarz inequality:

$$\text{var } \hat{\theta} \equiv \langle v, v \rangle_n \geq \frac{\langle u, v \rangle_n^2}{\langle u, u \rangle_n}. \quad (6)$$

This inequality is true for any function $u(s)$ of the data, and it becomes an equality when $u(s) \propto v(s)$. Since we wish to derive a bound that applies generally to all estimators, we should not have (or try) to provide too much detail about $\hat{\theta}$ (and therefore $v(s)$). A simple assumption to make on $\hat{\theta}$ is that it is an unbiased estimator:

$$\langle \hat{\theta}(s) \rangle_n = \theta_0 \Rightarrow \partial_{\theta_0} \langle \hat{\theta}(s) \rangle_n = 1. \quad (7)$$

How does this help us? It turns out that we can write a function $d(s)$ whose ensemble product with any other function $w(s)$ yields the derivative $\partial_{\theta_0} \langle w(s) \rangle_n$; this function is just $d(s) = \partial_{\theta_0} \log p(s|\theta_0)$, because

$$\begin{aligned} & \int w(s)[\partial_{\theta_0} \log p(s|\theta_0)]p(s|\theta_0)ds \\ &= \int w(s)\partial_{\theta_0} p(s|\theta_0)ds \\ &= \partial_{\theta_0} \int w(s)p(s|\theta_0)ds = \partial_{\theta_0} \langle w(s) \rangle_n, \end{aligned} \quad (8)$$

assuming of course³ that we can exchange integration and differentiation with respect to θ_0 . For any s , $d(s)$ encodes

³This assumption fails for some (mildly) pathological likelihood functions, which can provide counterexamples to the Cramér-Rao bound.

the *local relative change* in the likelihood function as θ_0 is changed. It follows that $\langle d(s), v(s) \rangle_n = \partial_{\theta_0} \langle \hat{\theta}(s) \rangle_n = 1$, so from Eq. (6) we get⁴

$$\text{var } \hat{\theta} \geq \frac{1}{\langle d(s), d(s) \rangle} \equiv \frac{1}{\langle \partial_{\theta_0} \log p(s|\theta_0), \partial_{\theta_0} \log p(s|\theta_0) \rangle_n}, \quad (9)$$

which is the unbiased-estimator version of the Cramér-Rao bound. If the estimator is biased, we can still use the Schwarz inequality by providing the derivative of the bias $b(\theta_0)$ with respect to θ_0 :

$$\langle \hat{\theta}(s) \rangle_n = \theta_0 + b(\theta_0) \Rightarrow \partial_{\theta_0} \langle \hat{\theta}(s) \rangle_n = 1 + \partial_{\theta_0} b(\theta_0), \quad (10)$$

and therefore

$$\text{var } \hat{\theta} \geq \frac{(1 + \partial_{\theta_0} b)^2}{\langle \partial_{\theta_0} \log p(s|\theta_0), \partial_{\theta_0} \log p(s|\theta_0) \rangle_n}. \quad (11)$$

Generalizing to a multidimensional expression is straightforward, if verbose (see, e.g., Ref. [10]):

$$\text{covar}_n(\hat{\theta}_i, \hat{\theta}_l) \geq (\delta_{im} + \partial_m b_i(\theta_0))F_{mj}^{-1}(\delta_{jl} + \partial_j b_l(\theta_0)), \quad (12)$$

where the *Fisher information matrix* is defined by

$$\begin{aligned} F_{il} &= \langle (\partial_i \log p(s|\theta_0)), (\partial_l \log p(s|\theta_0)) \rangle_n \\ &= -\langle \partial_i \partial_l \log p(s|\theta_0) \rangle_n. \end{aligned} \quad (13)$$

The second equality is established by taking the gradient of $\int (\partial_i \log p(s|\theta_0))p(s|\theta_0)ds$, and remembering that $\partial_i \int p(s|\theta_0)ds = \partial_i 1 = 0$. With the help of Eqs. (1) and (3), we can compute the Fisher matrix for GW signals in additive Gaussian noise, which is the familiar expression $F_{ij} = (\partial_i h, \partial_j h)$.

The full expression (12) for the Cramér-Rao bound, which includes the effects of bias, has interesting consequences, for it implies that biased estimators can actually *outperform*⁵ unbiased estimators, since the $\partial_m b_i(\theta_0)$ can be negative. Unfortunately, we have no handle on these derivatives without explicitly choosing a particular estimator (which goes against the idea of having a generic *bound*), so

⁴To obtain Eq. (9), we need to notice also that for any $w(s)$, $\langle d(s), \langle w(s) \rangle_n \rangle_n = 0$, since $\langle w(s) \rangle_n$ does not depend on s (but only on θ_0), and the integral of Eq. (8) reduces to $\langle w(s) \rangle_n \times \int \partial_{\theta_0} p(s|\theta_0)ds = \langle w(s) \rangle_n \partial_{\theta_0} 1 = 0$.

⁵This is true even if we evaluate the performance of estimators on the basis of their quadratic error

$$\begin{aligned} \langle (\hat{\theta}_i - \theta_{0i})(\hat{\theta}_l - \theta_{0l}) \rangle_n &\geq b_i(\theta_0)b_l(\theta_0) \\ &\quad + (\delta_{im} + \partial_m b_i(\theta_0))F_{mj}^{-1}(\delta_{jl} + \partial_j b_l(\theta_0)) \end{aligned}$$

rather than on the basis of their variance.

the Cramér-Rao bound can only give us a definite result for the subclass of unbiased estimators.

As pointed out by Cutler and Flanagan [5] (Appendix A 5), it follows that the bound cannot be used to place absolute limits on the accuracy of estimators (i.e., lower bounds on frequentist error)—limits that would exclude or severely limit the possibility of inferring the physical properties of sources from their emitted GWs. Even if the lower bound for unbiased estimators is very discouraging, there is always a chance that a biased estimator could do much better, so we cannot use the bound to prove “no go” theorems.

Going back to Eq. (6), we note that the bound is satisfied as an equality when

$$\begin{aligned} u(s) \propto v(s) \Rightarrow d(s) &\equiv \partial_{\theta_0} \log p(s|\theta_0) \\ &= q(\theta_0)[\hat{\theta}(s) - \langle \hat{\theta}(s) \rangle_n]. \end{aligned} \quad (14)$$

By integrating, we obtain a relation between the likelihood and the estimator:

$$p(s|\theta_0) = \frac{m(s)}{Z(\theta_0)} e^{-l(\theta_0)\hat{\theta}(s)}, \quad (15)$$

the estimation problems (i.e., the pairings of given likelihoods and chosen estimators) for which this relation holds true are said to belong to the *exponential family*, and these problems are the only ones for which the Cramér-Rao bound is satisfied exactly as an equality. Equation (15) generalizes trivially to multidimensional problems by replacing the exponential with $\exp\{-l_k(\theta_0)\hat{\theta}_k(s)\}$. Unfortunately, for a given $p(s|\theta_0)$ there is no guarantee that any unbiased estimator exists that satisfies Eq. (15) and that therefore can actually *achieve* the bound; all we can say in general about the performance of *unbiased* estimators is that they will underperform the Cramér-Rao bias, but we do not know how badly. As discussed above, the bound tells us nothing in general about *biased* estimators.

It follows that the bound cannot be used to establish guaranteed levels of accuracy (i.e., upper bounds on frequentist error), which would prove the possibility of inferring the physical properties of sources from their GWs. We can only do so if we can identify a specific estimator that achieves the bound. In the next section we shall see that the ML estimator⁶ does so in the high-SNR limit, where waveforms can be approximated accurately as *linear* functions of their parameters within the region of parameter space where $p(s|\theta)$ is not negligible (so the high-SNR limit coincides with the limit in which the LSA is accurate).

We conclude that the Cramér-Rao bound is seldom useful to the GW analyst as a proper bound, whether to

⁶Indeed, Eq. (15) implies that if both an *efficient* (i.e., bound-achieving) unbiased estimator and the ML estimator exist, they must coincide. To show this, we notice that if the ML estimator exists, the log-derivative $\partial_i \log p(s|\theta) = -\partial_i l_k(\theta)(\hat{\theta}_k - \theta_k)$ must be zero at $\theta = \hat{\theta}^{\text{ML}}$, from which it follows that $\hat{\theta}_k = \hat{\theta}_k^{\text{ML}}$.

make positive or negative expected-accuracy statements; where it is useful, it reduces to the high-SNR/LSA result for the ML estimator.

D. Second road: Derivation and critique of the frequentist high-SNR/LSA result

We denote the true signal as h_0 (so $s = h_0 + n$), and expand the generic waveform $h(\theta)$ around h_0 , normalizing signals by the *optimal signal-to-noise ratio* of the true signal, $A = \sqrt{\langle h_0, h_0 \rangle}$ (also known in this context as *signal strength*):

$$\begin{aligned} h(\theta) &= h_0 + \theta_k h_k + \theta_j \theta_k h_{jk}/2 + \dots \\ &= A(\bar{h}_0 + \theta_k \bar{h}_k + \theta_j \theta_k \bar{h}_{jk}/2 + \dots); \end{aligned} \quad (16)$$

here we are translating source parameters as needed to have $h(0) = h_0$, defining $h_i = \partial_i h|_{\theta=0}$, $h_{ij} = \partial_{ij} h|_{\theta=0}$ (and so on), and $\bar{h}_0 = h_0/A$, $\bar{h}_k = h_k/A$ (and so on).⁷ The likelihood is then given by⁸

$$\begin{aligned} p(s|\theta) &\propto e^{-(s-h(\theta), s-h(\theta))/2} \\ &= \exp\{-(n, n)/2 \\ &\quad - A^2[\theta_j \theta_k (\bar{h}_j, \bar{h}_k) + \theta_j \theta_k \theta_l (\bar{h}_j, \bar{h}_{kl}) + \dots]/2 \\ &\quad + A[\theta_j (n, \bar{h}_j) + \theta_j \theta_k (n, \bar{h}_{jk})/2 \\ &\quad + \theta_j \theta_k \theta_l (n, \bar{h}_{jkl})/3! + \dots]\}. \end{aligned} \quad (17)$$

The ML equations $\partial_j p(s|\theta^{\text{ML}}) = 0$ are given by

$$\begin{aligned} 0 &= \frac{1}{A} [(n, \bar{h}_j) + \hat{\theta}_k (n, \bar{h}_{jk}) + \hat{\theta}_k \hat{\theta}_l (n, \bar{h}_{jkl})/2 + \dots] \\ &\quad - [\hat{\theta}_k (\bar{h}_j, \bar{h}_k) + \hat{\theta}_k \hat{\theta}_l ((\bar{h}_j, \bar{h}_{kl})/2 + (\bar{h}_k, \bar{h}_{jl})) + \dots], \end{aligned} \quad (18)$$

where we have divided everything by A^2 , and we omit the “ML” superscript for conciseness. A careful study of Eq. (18) shows that it can be solved in perturbative fashion by writing $\hat{\theta}_j^{\text{ML}}$ as a series in $1/A$,

$$\hat{\theta}_j^{\text{ML}} = \hat{\theta}_j^{(1)}/A + \hat{\theta}_j^{(2)}/A^2 + \hat{\theta}_j^{(3)}/A^3 + \dots, \quad (19)$$

and by collecting the terms of the same order in Eq. (18),

⁷The statistical uncertainty in the estimated signal strength can still be handled in this notation by taking one of the \bar{h}_k to lie along \bar{h}_0 ; the corresponding θ_k represents a fractional correction to the true A .

⁸Formally, it is troubling to truncate the series expression for the exponent at any order beyond quadratic, since the integral of the truncated likelihood may become infinite; the important thing to keep in mind, however, is that the series need converge only within a limited parameter range determined self-consistently by the truncated-likelihood estimator, by compact parameter ranges, or (in the Bayesian case) by parameter priors. Similar considerations apply to the derivation of the higher-order corrections given in Sec. VII.

$$\begin{aligned}
O(1/A): (n, \bar{h}_j) - \hat{\theta}_k^{(1)}(\bar{h}_j, \bar{h}_k) &= 0, \\
O(1/A^2): \hat{\theta}_k^{(1)}(n, \bar{h}_{jk}) - \hat{\theta}_k^{(1)}\hat{\theta}_l^{(1)}((\bar{h}_j, \bar{h}_{kl})/2 \\
&\quad + (\bar{h}_k, \bar{h}_{jl})) - \hat{\theta}_k^{(2)}(\bar{h}_j, \bar{h}_k) &= 0, \\
O(1/A^3): \dots &
\end{aligned} \tag{20}$$

thus the ML solution $\hat{\theta}_j^{\text{ML}}$ is given by

$$\begin{aligned}
\hat{\theta}_j^{\text{ML}} &= \frac{1}{A}(\bar{h}_j, \bar{h}_k)^{-1}(\bar{h}_k, n) + \frac{1}{A^2}(\bar{h}_j, \bar{h}_i)^{-1}\{(n, \bar{h}_{ik}) \\
&\quad \times (\bar{h}_k, \bar{h}_i)^{-1}(\bar{h}_i, n) - ((\bar{h}_i, \bar{h}_{kl})/2 + (\bar{h}_k, \bar{h}_{il})) \\
&\quad \times (\bar{h}_k, \bar{h}_m)^{-1}(\bar{h}_m, n)(\bar{h}_l, \bar{h}_n)^{-1}(\bar{h}_n, n)\} \\
&\quad + \frac{1}{A^3}\{\dots\} + \dots
\end{aligned} \tag{21}$$

Thus we see that the limit of large A (i.e., high SNR) coincides with the *linearized-signal approximation* (LSA) where only the first derivatives of the signals are included. In the LSA, the likelihood is just

$$\begin{aligned}
p(s|\theta) &\propto \exp\{-(n, n)/2 - \theta_j\theta_k(h_j, h_k)/2 + \theta_j(h_j, n)\} \\
&= \exp\{-(n, n)/2 - A^2\theta_j\theta_k(\bar{h}_j, \bar{h}_k)/2 + A\theta_j(\bar{h}_j, n)\} \\
\text{(LSA),} &
\end{aligned} \tag{22}$$

and the ML estimator is given by

$$\hat{\theta}_j^{\text{ML}} = (1/A)(\bar{h}_j, \bar{h}_k)^{-1}(\bar{h}_k, n) \quad \text{(LSA).} \tag{23}$$

Since $\langle(\bar{h}_k, n)\rangle_n = 0$, we see also that the ML estimator is unbiased. The variance of $\hat{\theta}^{\text{ML}}$ is then obtained by averaging $\hat{\theta}_j^{\text{ML}}\hat{\theta}_k^{\text{ML}}$ over noise realizations,

$$\begin{aligned}
\langle\hat{\theta}_j^{\text{ML}}\hat{\theta}_k^{\text{ML}}\rangle_n &= \frac{1}{A^2}(\bar{h}_j, \bar{h}_i)^{-1}\langle(\bar{h}_i, n)(n, \bar{h}_m)\rangle_n(\bar{h}_m, \bar{h}_k)^{-1} \\
&= \frac{1}{A^2}(\bar{h}_j, \bar{h}_i)^{-1}(\bar{h}_i, \bar{h}_m)(\bar{h}_m, \bar{h}_k)^{-1} \\
&= \frac{1}{A^2}(\bar{h}_j, \bar{h}_k)^{-1} \quad \text{(LSA),}
\end{aligned} \tag{24}$$

and it coincides with the mean quadratic error in the frequentist sense. In Eq. (24), the second equality follows from Eq. (3). The interpretation of the limit is that, for strong signals, the typical $\hat{\theta}_j^{\text{ML}} - \theta_{0j}$ becomes small enough that the log-likelihood is accurately described by the product of detector data and a linearized signal.

Equation (24) is the standard Fisher-information-matrix result, and it implies that in the high-SNR/LSA limit the ML estimator achieves the Cramér-Rao bound. As we shall see in Sec. VII, the next-order correction to the variance scales as $1/A^4$, not $1/A^3$. This is because all $O(1/A^3)$ terms contain odd numbers of n , whose products vanish under the ensemble average. The fact itself that there is a next-order correction shows that for generic A the ML estimator does not achieve the bound.

The fact that the Cramér-Rao bound is achieved in the high-SNR/LSA limit, but not beyond it, can also be seen in the light of Eq. (15), which encodes a standard form for the estimation problems in the exponential family. To express the LSA likelihood in this form, we can set $m(s) = e^{-(s-h_0, s-h_0)/2}$ and $Z(\theta) = e^{\theta_j\theta_k(h_j, h_k)/2}$; it remains to establish that

$$-l_j(\theta)\hat{\theta}_j^{\text{ML}}(s) = \theta_j(h_j, s - h_0), \tag{25}$$

which is satisfied by $l_j(\theta) = -(h_j, h_k)\theta_k$ [see Eq. (23)]. Now, if additional terms are added to Eq. (22), beginning with terms cubic in the θ_i , $\hat{\theta}_j^{\text{ML}}(s)$ comes to be a nonlinear function of the signal, such that no $-l_j(\theta)$ can multiply it in the right way to reconstruct the likelihood. It then follows that the estimation problem moves outside the exponential family, and the Cramér-Rao bound cannot be achieved.

It is possible (but perhaps not desirable, as we shall see shortly) to modify the ML estimator to take into account prior knowledge about the expected distribution of sources. The resulting maximum-posterior estimator $\hat{\theta}^{\text{MP}}$ is defined as the mode of the posterior probability $p(\theta|s) = p(s|\theta)p(\theta)/p(s)$,

$$\hat{\theta}^{\text{MP}} = \text{maxloc}_\theta p(\theta|s) = \text{maxloc}_\theta p(s|\theta)p(\theta). \tag{26}$$

This is a biased estimator: in the high-SNR/LSA limit, and with a Gaussian prior $p(\theta) \propto \exp\{-P_{ij}(\theta_i - \theta_i^P) \times (\theta_j - \theta_j^P)/2\}$ centered at θ^P (the only prior that can be easily handled analytically), we find

$$\begin{aligned}
b_i^{\text{MP}} &= \langle\theta_i^{\text{MP}}\rangle_n = [(\bar{h}_i, \bar{h}_j) + P_{ij}/A^2]^{-1}(P_{jk}/A^2)\theta_k^P \\
&\quad \text{(LSA/Gaussian prior);}
\end{aligned} \tag{27}$$

thus $\hat{\theta}^{\text{MP}}$ becomes unbiased for $A \rightarrow \infty$ (indeed, in that limit $\hat{\theta}^{\text{MP}}$ tends to $\hat{\theta}^{\text{ML}}$). For the frequentist variance we find

$$\begin{aligned}
\langle\hat{\theta}_i^{\text{MP}}\hat{\theta}_j^{\text{MP}}\rangle_n - \langle\hat{\theta}_i^{\text{MP}}\rangle_n\langle\hat{\theta}_j^{\text{MP}}\rangle_n &= \langle\hat{\theta}_i^{\text{MP}}\hat{\theta}_j^{\text{MP}}\rangle_n - b_i^{\text{MP}}b_j^{\text{MP}} \\
&= \frac{1}{A^2}[(\bar{h}_i, \bar{h}_k) + P_{ik}/A^2]^{-1}(\bar{h}_k, \bar{h}_l) \\
&\quad \times [(\bar{h}_l, \bar{h}_j) + P_{lj}/A^2]^{-1} \quad \text{(LSA w/prior),}
\end{aligned} \tag{28}$$

which coincides⁹ with the generalized Cramér-Rao bound of Eq. (12), proving that the estimation problem defined by

⁹Note that the Fisher matrix that must be substituted into Eq. (12) is still $-\langle\partial_j\partial_k p(s|\theta)\rangle_n = (h_j, h_k)$, and not $-\langle\partial_j\partial_k [p(s|\theta)p(\theta)]\rangle_n = (h_j|h_k) + P_{jk}$. The prior distribution does not concern the Cramér-Rao bound, which is computed from the likelihood alone for a *fixed known* value of the true source parameters. Instead, we happen to be using an estimator that takes into account prior information, which enters into the Cramér-Rao bound via the derivative of the bias.

the LSA likelihood and $\hat{\theta}^{\text{MP}}$ belongs to the exponential family.

The reason why $\hat{\theta}^{\text{MP}}$ is not too useful to characterize future parameter-estimation performance is that we expect a reasonable measure of error to converge to the effective width of the prior in the limit of vanishing signal strength. Instead, in the absence of any information from the experiment, $\hat{\theta}^{\text{MP}}$ becomes stuck at the mode of the prior, and its variance [in Eq. (28)] tends to zero. This behavior occurs for any nonuniform prior.¹⁰

E. Third-road derivation of the Bayesian high-SNR/LSA result

We now wish to show that in any single experiment, if the high-SNR/LSA limit is warranted (and if the parameter priors are uniform over the parameter region of interest), the inverse Fisher-information matrix yields the variance of the Bayesian posterior probability distribution. To do so, we rewrite Eq. (17) in terms of *normalized parameters* $\bar{\theta}_i = A\theta_i$:

$$\begin{aligned} p(s|\theta) \propto & \exp\{-(n, n)/2 + [(n, \bar{h}_j)\bar{\theta}_j + \frac{1}{A}(n, \bar{h}_{jk})\bar{\theta}_j\bar{\theta}_k]/2 \\ & + \frac{1}{A^2}(n, \bar{h}_{jkl})\bar{\theta}_j\bar{\theta}_k\bar{\theta}_l/3! + O(1/A^3)] \\ & - [(\bar{h}_j, \bar{h}_k)\bar{\theta}_j\bar{\theta}_k + \frac{1}{A}(\bar{h}_j, \bar{h}_{kl})\bar{\theta}_j\bar{\theta}_k\bar{\theta}_l \\ & + \frac{1}{A^2}(\bar{h}_{jk}, \bar{h}_{lm})\bar{\theta}_j\bar{\theta}_k\bar{\theta}_l\bar{\theta}_m/4 \\ & + \frac{2}{A^2}(\bar{h}_j, \bar{h}_{klm})\bar{\theta}_j\bar{\theta}_k\bar{\theta}_l\bar{\theta}_m/3! + O(1/A^3)]/2\}. \end{aligned} \quad (29)$$

We can build the variance from the posterior mean

$$\langle \bar{\theta}_i \rangle_p \equiv \int \bar{\theta}_i p(s|\theta) d\theta / \int p(s|\theta) d\theta \quad (30)$$

and the quadratic moment

$$\langle \bar{\theta}_i \bar{\theta}_j \rangle_p = \int \bar{\theta}_i \bar{\theta}_j p(s|\theta) d\theta / \int p(s|\theta) d\theta, \quad (31)$$

where “ $\langle \cdot \rangle_p$ ” denotes integration over $p(s|\theta)$. The idea is to proceed in perturbative fashion, writing the moments as series in $\epsilon = 1/A$: taking $\langle \bar{\theta}_i \rangle_p$ as an example,

$$\langle \bar{\theta}_i \rangle_p = \sum_{n=0}^{\infty} \frac{\epsilon^n}{n!} \langle \bar{\theta}_i \rangle_p^{(n)} \Rightarrow \langle \bar{\theta}_i \rangle_p^{(n)} = \left. \frac{\partial^n \langle \bar{\theta}_i \rangle_p}{\partial \epsilon^n} \right|_{\epsilon=0}. \quad (32)$$

Since ϵ appears at both the numerator and denominator of Eq. (30), we write

¹⁰For uniform priors (e.g., rectangular distributions corresponding to the allowed parameter ranges), $\hat{\theta}^{\text{MP}}$ actually becomes undefined in the $A \rightarrow 0$ limit.

$$\langle \bar{\theta}_i \rangle_p = \frac{\int \bar{\theta}_i p(0) d\theta + \epsilon \int \bar{\theta}_i \frac{\partial p(0)}{\partial \epsilon} d\theta + \frac{\epsilon^2}{2} \int \bar{\theta}_i \frac{\partial^2 p(0)}{\partial \epsilon^2} d\theta + \dots}{\int p(0) d\theta + \epsilon \int \frac{\partial p(0)}{\partial \epsilon} d\theta + \frac{\epsilon^2}{2} \int \frac{\partial^2 p(0)}{\partial \epsilon^2} d\theta + \dots} \quad (33)$$

(where the argument of p implies that the (n) th derivative is evaluated at $\epsilon = 0$), and therefore

$$\begin{aligned} \langle \bar{\theta}_i \rangle_p^{(0)} &= \int \bar{\theta}_i p(0) d\theta / \int p(0) d\theta, \\ \langle \bar{\theta}_i \rangle_p^{(1)} &= \left[\int \bar{\theta}_i \frac{\partial p(0)}{\partial \epsilon} d\theta \right. \\ &\quad \left. - \langle \bar{\theta}_i \rangle_p^{(0)} \int \frac{\partial p(0)}{\partial \epsilon} d\theta \right] / \int p(0) d\theta; \quad \dots \end{aligned} \quad (34)$$

similar expressions hold for $\langle \bar{\theta}_i \bar{\theta}_j \rangle_p$, and a general expression for the (n) th-order contribution is given in Sec. VII B. The $\epsilon \rightarrow 0$ limit coincides with the limit of large signal strengths, or of vanishing derivatives higher than the first, since in that case Eq. (29) truncates to Eq. (22). In this limit,

$$\langle \bar{\theta}_i \rangle_p = \langle \bar{\theta}_i \rangle_p^{(0)} = (\bar{h}_i, \bar{h}_j)^{-1}(n, \bar{h}_j) \quad (\text{LSA}) \quad (35)$$

and

$$\begin{aligned} \langle \Delta \bar{\theta}_i \Delta \bar{\theta}_j \rangle_p &= \langle (\bar{\theta}_i - \langle \bar{\theta}_i \rangle_p^{(0)}) (\bar{\theta}_j - \langle \bar{\theta}_j \rangle_p^{(0)}) \rangle_p^{(0)} \\ &= (\bar{h}_i, \bar{h}_j)^{-1} \quad (\text{LSA}), \end{aligned} \quad (36)$$

and therefore

$$\langle \theta_i \theta_j \rangle_p = \frac{1}{A^2} (\bar{h}_i, \bar{h}_j)^{-1} = (h_i, h_j)^{-1} \quad (\text{LSA}), \quad (37)$$

as can be seen by rewriting the exponential of Eq. (22) as

$$p(s|\bar{\theta}) \propto \exp\{-(\bar{h}_i, \bar{h}_j)(\bar{\theta}_i - \langle \bar{\theta}_i \rangle_p)(\bar{\theta}_j - \langle \bar{\theta}_j \rangle_p)/2\}, \quad (38)$$

where we have omitted factors independent from $\bar{\theta}$ that cancel out in the normalization of $p(s|\bar{\theta})$.

Reinstating A in Eq. (35) we see that in the high-SNR/LSA limit the mean of the posterior distribution coincides with the ML estimator, as is reasonable, since the average of a normal distribution coincides with its mode. The two, however, differ when higher-order terms are included, as we shall see in Sec. VII. From Eq. (36) we see also that, to leading order, the variance of the posterior distribution is experiment-independent, and it coincides with the variance of the ML estimator (remember, however, that the two have very different interpretations).¹¹

¹¹If we define the quadratic error of the posterior distribution as $\langle \bar{\theta}_i \bar{\theta}_j \rangle_p$ (which is appropriate given that the true signal is at $\theta = 0$), we must increment $(\bar{h}_i, \bar{h}_j)^{-1}$ by the experiment-dependent quantity $\langle \bar{\theta}_i \rangle_p \langle \bar{\theta}_j \rangle_p = (\bar{h}_i, \bar{h}_j)^{-1}(n, \bar{h}_i)(\bar{h}_m, n)(\bar{h}_m, \bar{h}_j)^{-1}$. Interestingly, the frequentist average of the Bayesian error $\langle \langle \bar{\theta}_i \bar{\theta}_j \rangle_p \rangle_n$ is $2(\bar{h}_i, \bar{h}_j)^{-1}$, twice the frequentist variance of $\hat{\theta}^{\text{ML}}$.

With the addition of a Gaussian prior $p(\theta) \propto e^{-P_{ij}\theta_i\theta_j/2}$ centered at $\theta = 0$, Eqs. (35) and (36) change only slightly¹²:

$$\begin{aligned}\langle \bar{\theta}_i \rangle_p &= [(\bar{h}_i, \bar{h}_j) + P_{ij}/A^2]^{-1}(n, \bar{h}_j), \\ \langle \Delta \bar{\theta}_i \Delta \bar{\theta}_j \rangle_p &= [(\bar{h}_i, \bar{h}_j) + P_{ij}/A^2]^{-1} \\ &\quad (\text{LSA/Gaussian prior}).\end{aligned}\quad (39)$$

Note that $p(\theta) \propto e^{-(1/A^2)P_{ij}\bar{\theta}_i\bar{\theta}_j/2}$ is formally an $O(1/A^2)$ contribution to the likelihood exponential that would enter the $1/A$ expansion beginning at that order. However, if P_{ij} is large enough to matter at the signal strengths of interest, it probably makes sense to bundle it with the zeroth-order terms as we did here. In contrast with Eq. (28) for the frequentist variance of $\hat{\theta}^{\text{MP}}$, we see that in the limit of vanishing signal strength the variance of the posterior goes to the variance P_{ij} of the prior.

III. STANDARD COMPACT-BINARY SIGNAL MODEL

Throughout the rest of this paper, our fiducial model for compact-binary signals will be simple stationary-phase-approximated (SPA) waveforms including phasing terms from the spin-orbit and spin-spin interactions of parallel or antiparallel component spins. Parameter estimation with these waveforms was studied by Poisson and Will [14]. In this paper we adopt second-order post-Newtonian¹³ (2PN) Fourier-domain waveforms as written by Arun and colleagues [15]:

$$\begin{aligned}\tilde{h}(M_c, \eta, \beta, \sigma, \phi_0, t_0; f) &\propto f^{-7/6} \exp\{i\psi(M_c, \eta, \beta, \sigma; f) \\ &\quad + \phi_0 + 2\pi f t_0\},\end{aligned}\quad (40)$$

¹²With the Gaussian prior, the quadratic error $\langle \langle \bar{\theta}_i \bar{\theta}_j \rangle_p \rangle_n$ becomes $[(\bar{h}_i, \bar{h}_j) + P_{ij}/A^2]^{-1}(\bar{h}_i, \bar{h}_m)[(\bar{h}_m, \bar{h}_j) + P_{mj}/A^2]^{-1}$.

¹³Waveform phasing expressions accurate to 3.5PN order are also provided in Ref. [15]. We do not use these in this article for the sake of simplicity, since they would not change the qualitative picture of parameter estimation presented here. For the reader's reference, however, the higher-than-2PN corrections to Eq. (41), including the errata to Ref. [15], are

$$\begin{aligned}&\pi \left(\frac{38\,645}{252} - \frac{65}{9} \eta^3 \right) v^5 \log v + \left[\left(\frac{11\,583\,231\,236\,531}{4\,694\,215\,680} - \frac{640}{3} \pi^2 \right. \right. \\ &\quad \left. \left. - \frac{6848}{21} \gamma - \frac{6848}{21} \log(4) \right) + \left(-\frac{15\,335\,597\,827}{3\,048\,192} + \frac{2255}{12} \pi^2 \right. \right. \\ &\quad \left. \left. - \frac{1760}{3} \theta + \frac{12\,320}{9} \lambda \right) \eta + \frac{76\,055}{1728} \eta^2 - \frac{127\,825}{1296} \eta^3 \right] v^6 \\ &\quad - \frac{6848}{21} v^6 \log v + \pi \left(\frac{77\,096\,675}{254\,016} + \frac{378\,515}{1512} \eta \right. \\ &\quad \left. - \frac{74\,045}{756} \eta^2 \right) v^7,\end{aligned}$$

where $\gamma = 0.57721 \dots$ is Euler's constant, and $\lambda = -1987/3080$ and $\theta = -11\,831/9240$ are recently determined constants in the PN expansion [16].

with

$$\begin{aligned}\psi(M_c, \eta, \beta, \sigma; f) &= \frac{3}{128\eta v^5} \left\{ 1 + \frac{20}{9} \left(\frac{743}{336} + \frac{11}{4} \eta \right) v^2 + (4\beta - 16\pi) v^3 \right. \\ &\quad \left. + 10 \left(\frac{3\,058\,673}{1\,016\,064} + \frac{5429}{1008} \eta + \frac{617}{144} \eta^2 - \sigma \right) v^4 \right\},\end{aligned}\quad (41)$$

where $v = (\pi M f)^{1/3}$, $M = m_1 + m_2$ is the total mass of the binary, $\eta = m_1 m_2 / M^2$ is the symmetric mass ratio, $M_c = M \eta^{3/5}$ is the chirp mass. The spin-orbit parameter β and the spin-spin parameter σ [14,17] are given by

$$\begin{aligned}\beta &= \sum_{i=1}^2 \frac{\hat{\mathbf{L}} \cdot \mathbf{S}_i}{12m_i^2} \left[113 \left(\frac{m_i}{M} \right)^2 + 75\eta \right] \\ &= \frac{113 \hat{\mathbf{L}} \cdot [\mathbf{S}_1 + \mathbf{S}_2] + 75 \hat{\mathbf{L}} \cdot [(m_2/m_1)\mathbf{S}_1 + (m_1/m_2)\mathbf{S}_2]}{12M^2},\end{aligned}\quad (42)$$

and

$$\sigma = \frac{721(\hat{\mathbf{L}} \cdot \mathbf{S}_1)(\hat{\mathbf{L}} \cdot \mathbf{S}_2) - 247\mathbf{S}_1 \cdot \mathbf{S}_2}{48m_1 m_2 M^2},\quad (43)$$

with \mathbf{S}_1 and \mathbf{S}_2 the spins of the binary components. We truncate waveforms at the (Keplerian) last stable circular orbit ($v = 1/\sqrt{6}$).

For simplicity, in this article we do not discuss the estimation of the amplitude parameter \mathcal{A} that would multiply the right-hand side of Eq. (40). [From Eqs. (2) and (40) it follows that $(\partial_{\mathcal{A}} h, \partial_i h) = 0$ for $i \neq \mathcal{A}$, so the amplitude \mathcal{A} effectively decouples from all other parameters in the Fisher matrix.] However, all discussions to follow can accommodate the addition of \mathcal{A} with trivial modifications.

IV. THE SINGULAR CASE OF THE DISAPPEARING PARAMETER

In Sec. II we have examined the interpretation of the inverse Fisher matrix as a frequentist or Bayesian measure of error or uncertainty. In this section, we discuss what happens when the Fisher matrix is singular, or almost so, so that the attempts to invert it numerically yield warnings that it is badly conditioned. It is pedagogical to begin this discussion by considering the case where the matrix is exactly singular (Sec. IV A), and then to widen our scope to approximate singularity (Sec. IV B). The conclusion is that a singular Fisher matrix is almost always a symptom that the high-SNR/LSA limit is not to be trusted, that prior probabilities play an important role, or both.

A. Singular Fisher matrix

A singular Fisher matrix implies that the corresponding LSA likelihood (22) is a *singular normal* distribution [18], which is constant along the directions of the Fisher-matrix

eigenvectors with null eigenvalues¹⁴ (henceforth, somewhat improperly, we shall call these *null eigenvectors*), so the ML equation has no solutions, and the even moments of the distribution are infinite, even for parameters that do not appear in the null eigenvectors. Thus, the frequentist variance of the ML estimator and the Bayesian variance of the posterior distribution are (formally) infinite for all parameters.

How to deal with this? If the signal is *really* linear, so that the LSA expressions are exact, it is possible to discard the combinations of parameters that correspond to the null eigenvectors, and characterize the variance of the remaining parameters. Let us see how, in the frequentist and Bayesian frameworks. In what follows, we denote the total number of source parameters by N , and the number of non-null eigenvectors by R .

In the frequentist ML framework, we write (\bar{h}_i, \bar{h}_j) in the singular-value (SV) decomposition¹⁵ [20] as $\sum_{\lambda^{(k)} \neq 0} \theta_i^{(k)} \lambda^{(k)} \theta_j^{(k)}$ (with $(k) = 1, \dots, R$), or $\Theta \Sigma \Theta^T$ in matrix notation (with Θ an $N \times R$ matrix with orthonormal columns, and Σ a diagonal matrix formed from the R nonzero eigenvalues). We can then refactor the ML equation as

$$(\Theta \Sigma \Theta^T) \hat{\theta} = A^{-1} \hat{\mathbf{n}} \Rightarrow (\Theta^T \hat{\theta}) = A^{-1} (\Sigma^{-1} \Theta^T \hat{\mathbf{n}}) \Rightarrow \hat{c}^{(k)} = A^{-1} (\lambda^{(k)})^{-1} n^{(k)}, \quad (44)$$

where $\hat{c}^{(k)}$ and $n^{(k)}$ denote the coefficients of the decompositions of $\hat{\theta}$ and (\bar{h}_i, n) with respect to the normalized non-null eigenvectors of (\bar{h}_i, \bar{h}_j) . Since the ensemble average $\langle n^{(k)} n^{(l)} \rangle_n$ is just $\lambda^{(k)} \delta^{(k)(l)}$ (where δ is Kronecker's delta), the frequentist covariance of the ML estimators $\hat{c}^{(k)}$ is the diagonal matrix $A^{-2} (\lambda^{(k)})^{-1} \delta^{(k)(l)}$.

In the Bayesian framework, the quantities of interest are the moments of the $c^{(k)}$ over infinite ranges of the $c^{(k)}$ and of the coefficients $C^{(K)}$ (with $(K) = 1, \dots, N - R$) corresponding to the null eigenvectors, which are not included in the SV decomposition. Formally, these moments are ratios of two infinities, because the LSA likelihood is not a function of the C [not even through the $n^{(K)} \equiv \theta_i^{(K)} (\bar{h}_i, n)$ terms, which are zero since $(\theta_i^{(K)} \bar{h}_i, \theta_j^{(K)} \bar{h}_j) = 0$], but they

may be evaluated as improper integrals, *in the limit* of the ranges for the $c^{(K)}$ extending to infinity:

$$\begin{aligned} \langle \Delta c^{(k)} \Delta c^{(l)} \rangle_p &= \frac{\int c^{(k)} c^{(l)} p(s|c) dc dC}{\int p(s|c) dc dC} \\ &= \lim_{\Delta C^{(k)} \rightarrow \infty} \frac{(\int_{-\Delta C^{(k)}}^{+\Delta C^{(k)}} dC) \int c^{(k)} c^{(l)} p(s|c) dc}{(\int_{-\Delta C^{(k)}}^{+\Delta C^{(k)}} dC) \int p(s|c) dc} \\ &= A^{-2} (\lambda^{(k)})^{-1} \delta^{(k)(l)}. \end{aligned} \quad (45)$$

We can then work back to the frequentist components of the covariance matrix (or the Bayesian posterior moments) that involve any $\hat{\theta}_i$ that *do not appear* in the null eigenvectors. All other $\hat{\theta}_i$, however, are *completely* indeterminate.¹⁶ In the frequentist framework, it may be possible to work back to *interval* estimates of their values by combining a ML estimate of the $\hat{c}^{(k)}$ with finite allowed ranges for some of the θ_i ; however, this would constitute a form of prior distribution for the θ_i , which is not entirely compatible with the ML estimator (what happens if the solution of the ML equation falls outside the allowed range?). In the Bayesian framework, salvation may come from the prior probability distributions that make the posterior integrable.¹⁷ Unless the priors are also normal, though, the resulting moments cannot be expressed simply as analytical expressions of the Fisher matrix.

The most benign outcome occurs when the null eigenvectors correspond individually to one or more of the original parameters, or when the subspace spanned by null eigenvectors corresponds to a subset of the original parameters. The null-eigenvector combinations of parameters may also have clear physical interpretations: for instance, for a monochromatic, continuous sinusoid of frequency f , the absolute time offset t_0 and the initial phase ϕ_0 are essentially the same parameter, so the Fisher matrix has a null eigenvector along the parameter combination $ft_0 - \phi_0$, which can be discarded, while $ft_0 + \phi_0$ remains well determined. A similar case is the degeneracy between luminosity distance and a certain function of the sky-position angles in the analysis of short GW chirps with a single ground-based detector.¹⁸ Other

¹⁴A reasonable objection to computing the eigensystem of the Fisher matrix is that it leads to taking linear combinations of parameters that may have different units. It is possible to avoid this problem by looking at the Fisher matrix more abstractly as a linear operator, and talking of its range and null space [19]; or more pragmatically, by dividing all parameters by their typical range; or perhaps by taking their logarithm (since we are working with errors, units can be forgotten as additive constants), which in the linearized theory is equivalent to dividing by the true parameters. We are going to largely ignore this issue, treating the parameters as pure numbers resulting from adopting a God-given system of units.

¹⁵For square matrices, the SV decomposition is essentially equivalent to an eigenvector-diagonal-matrix decomposition where we drop the rows and columns corresponding to the null eigenvalues and eigenvectors.

¹⁶In a truly linear system, this is true no matter how small the eigenvector component in that parameter direction; clearly, this raises a problem of accuracy in the numerical computation of eigenvectors.

¹⁷Even a single prior in the form of a rectangle function will regularize the integration over all the null-eigenvector coordinates that include that parameter. For normal priors, whether the posterior becomes integrable depends on the eigenstructure of $A^2(\bar{h}_i, \bar{h}_j) + P_{ij}$.

¹⁸Although neither of these examples is a linear model described exactly by the LSA, the degeneracy persists in the exact likelihood, so its Fisher-matrix diagnosis is correct. For such "perfect" degeneracies to occur, the two parameters must appear in all waveform expressions only as a sum or product; this would imply that their units can be sensibly summed, or that their combination has direct physical meaning.

combinations of parameters can be more ambiguous and troubling—what is the meaning of estimating a parameter equal to a mass plus a spin? In those cases, our best hope is again that the degeneracy will be cured by prior probabilities, or by higher-order corrections in the $1/A$ expansion, in which cases the Fisher-matrix formalism is certainly insufficient.

B. Ill-conditioned Fisher matrix

All nonsingular matrices have well-defined inverses, although these *may* be difficult to compute. The notion of *ill conditioning* from the theory of linear systems of equations [20] can be invoked here to provide a bound (valid under reasonable conditions) on the perturbation of the inverse of a perturbed matrix,

$$\frac{\|(M + \delta M)^{-1} - M^{-1}\|}{\|M^{-1}\|} \leq \kappa(M) \frac{\|\delta M\|}{\|M\|} + O(\|\delta M\|^2); \quad (46)$$

here “ $\|\cdot\|$ ” is a matrix norm (e.g., the 2-norm $\|M\|_2 = \sup_{\mathbf{x}} \|M\mathbf{x}\|_2 / \|\mathbf{x}\|_2$ derived from the vector 2-norm $\|\mathbf{x}\|_2 = (\sum_i x_i^2)^{1/2}$), and $\kappa(M) = \|M\| \times \|M^{-1}\|$ is the *condition number*. Since $\|M\|_2$ is equal to M 's largest eigenvalue, $\kappa_2(M)$ is given by the ratio of its largest-to-smallest-modulus eigenvalues. From a numerical-analysis perspective, as Finn [4] points out, the gist of Eq. (46) is that, roughly speaking, matrix inversion can amplify roundoff error by a factor κ , leading to the loss of up to $\log_{10}\kappa$ digits of precision. The same amplification will apply to any inaccuracies in our knowledge of M . Taken at face value, this means that the Fisher-matrix results of Eqs. (24) and (36) may be inaccurate at a 100% level if the components of the Fisher matrix are not known to a fractional accuracy better than $\kappa^{-1}(F)$.

Of course, Eq. (46) is only an upper bound, and this doomsday scenario need not be realized in practice. One way to check whether the matrix-inversion sensitivity is a concern is to add small random perturbations, Monte Carlo-style, to the Fisher-matrix elements, and then verify the change in the covariance matrix. Such an experiment for our standard SPA model (with $m_1 = m_2 = 10M_\odot$ and no spins) shows that perturbing the 12th significant digits of the F_{ij} components is already enough to engender 100% changes in the diagonal elements of $(F^{-1})_{ij}$ (i.e., the predicted parameter variances). This behavior is ~ 100 times less severe than what is predicted by Eq. (46), but it still tells a rather cautionary tale about numerical sensitivity in the inversion of that particular Fisher matrix.¹⁹ These problems can be cured, somewhat

¹⁹Augmenting the Fisher matrix with normal priors for η , β , and σ , as described in Sec. VA, can somewhat cure this instability to inversion, although the result is SNR-dependent: for SNR = 10, the errors in $(F_{ij} + P_{ij})^{-1}$ become intolerable for fractional perturbations in F_{ij} of order 10^{-7} rather than 10^{-12} , but in the high-SNR limit the threshold reverts to the latter.

trivially, by adopting higher-precision arithmetics, and by computing the Fisher matrix to better accuracy. It may also be possible to improve the condition number by changing the units of source parameters, which may reduce the magnitude gap between the largest- and smallest-modulus eigenvalues.

More to the point, it is the consequences of the Fisher-matrix condition number on the substance (rather than the numerical accuracy) of Eqs. (23) and (35) that should attract our attention. A large condition number implies one or more small Fisher-matrix eigenvalues, and consequently large statistical fluctuations for the combinations of source parameters corresponding to the small-eigenvalue eigenvectors, at least according to the LSA. The interpretation is that large parameter changes in the direction of the small-eigenvalue eigenvectors are needed to produce changes in the waveform comparable to typical noise fluctuations. Under this condition, we have to worry whether the LSA can really describe the likelihood over the entire parameter ranges of interest: of course, these depend on the SNR available at detection (at leading order, their extent is inversely proportional to signal strength). In Sec. VI we describe a numerical criterion to decide when the SNR is high enough to believe the LSA. We also have to worry whether prior probability distributions for the parameters (perhaps in the simple form of allowed ranges) already restrict the estimated (for frequentists) or probable (for Bayesians) values of parameters beyond what is predicted by the Fisher-matrix variance. In the next section we discuss a simple test to decide whether priors should be included.

V. THE BURDEN OF PRIOR COMMITMENTS

As Cutler and Flanagan point out [5] (p. 2691), “it is not necessary for *a priori* information to be very detailed or restrictive in order that it have a significant effect on parameter-extraction accuracy. All that is necessary is that it be more restrictive than the information contained in the waveform, for some of the parameters [...] what is more surprising is that due to the effects of correlations, the root-mean-square (rms) errors obtained for the other parameters may also be overestimated by large factors.” Roughly speaking, this happens because as we move in parameter space, the change in the signal can be partially absorbed by changing correlated parameters together; thus, limiting the range available to one parameter also limits the range over which a correlated parameter can run while not significantly modifying the signal. In this section we seek a practical recipe to determine, in the context of a parameter-estimation problem specified by a family of waveforms and a fiducial SNR, whether it is necessary to take priors into consideration when evaluating projected parameter accuracies.

Since prior probabilities can only be discussed consistently in the framework of Bayesian parameter estimation,

in this section we will restrict ourselves to that context. The Gaussian priors examined at the end of Sec. II E are rarely appropriate in actual practice, but they do provide a quick test to see if the priorless Fisher result can be taken as it stands, or whether a more careful analysis is needed that includes the effects of priors. In Sec. V A we try out this quick test on the SPA model of Sec. III. For simplicity, we shall consider the effects of priors as logically independent from the sufficiency of the LSA, although the two problems clearly come into play together in real situations.

A. Testing for the influence of priors (normal true-parameter-centered priors)

We shall discuss our quick test by way of an example. The standard SPA model of Sec. III has six parameters: M_c , η , β , σ , ϕ_0 , and t_0 (plus A , which we disregard). We work at 2PN with $\text{SNR} = A = 10$, with true parameters $m_1 = m_2 = 10M_\odot$ (corresponding to $M_c = 8.71M_\odot$, $\eta = 0.25$), and $\beta = \sigma = \phi_0 = t_0 = 0$. We wish to examine the effect of priors for three related parameter-estimation problems involving different subsets of parameters: a 4-parameter problem (4pp) where we disregard spin parameters (i.e., where we assume we know *a priori* that the true binary has no spin); a 5pp where spin-orbit coupling [as represented by β in Eq. (41)] is important, but spin-spin interactions can be neglected; and a 6pp where we include also spin-spin interactions [as represented by σ in Eq. (41)]. As we shall see, priors become increasingly important as the number of parameters increases.

In each problem, we compute the expected covariance matrix of the posterior distribution as the inverse of (a submatrix of) the Fisher matrix, neglecting any non-LSA effects. We represent priors as normal distributions centered around null parameter displacements (i.e., the true parameter value), with standard deviations of 0.25 for η and, following Poisson and Will [14] 8.5 for β and 5 for σ (in Ref. [21], Berti and colleagues derive and adopt approximate priors for β and σ with standard deviations $\Delta\beta = 9.4$ and $\Delta\sigma = 2.5$). This representation is very crude, but it is the only one that leads to a simple analytical result [Eq. (39)] for the posterior covariance, and it should give at least a qualitative idea of the effect of imposing rectangular priors covering the allowed parameter ranges. Results are shown in the upper section of Table I, and are as follows.

The first line of Table I shows the 4pp no-prior 1σ values for the single-parameter rms errors (i.e., the square roots of the diagonal elements in the covariance matrix). Among these, ΔM_c and $\Delta\eta$ seem reasonable, but we get hung up on the value of $\Delta\phi_0$. Can the 1σ error region be larger than the physically meaningful range for this angle? On general grounds, we should worry that the LSA cannot know that the waveforms are exactly periodic (and therefore non-linear) in the angular parameters, so it blithely extrapolates small-angle effects to infinite ranges. However, as pointed

out by Cutler [22], this extrapolation is roughly correct for a simple complex phase such as ϕ_0 [see Eq. (40)], for which the main correlated-parameter effect is to absorb the global phase shifts due to changes in the other parameters.²⁰ A large $\Delta\phi_0$ indicates that this absorption can happen through several cycles of phasing. We conclude that ϕ_0 is essentially undetermined, but we have no reason to distrust the errors for ΔM_c , $\Delta\eta$, and Δt_0 .

Applying a prior to η does not change the picture significantly, but priors do matter once we add the spin parameters, which are very poorly determined at this SNR. In the 5pp, we find unphysically large errors for both β and η , which are cured only by imposing priors on *both* parameters. In the 6pp, we find that a prior is needed also for σ ; adding it engenders measurable changes in ΔM_c and $\Delta\eta$. As a rule of thumb, we should expect such effects whenever the signal derivatives show significant correlations, and when the magnitudes of the priors, measured crudely as the squared inverses $(\theta_i^{\max} - \theta_i^{\min})^{-2}$ of the effective parameter ranges induced by the priors, are comparable to the corresponding diagonal Fisher-matrix elements F_{ii} .

B. Testing for the influence of priors (exact priors)

We can perform an even better test by evaluating the effects of *exact* priors while still working in the LSA. Doing this requires some numerics, which are, however, very manageable on a workstation-class system. The idea is to integrate $\langle \Delta\theta_i \Delta\theta_j \rangle_p$ as a Monte Carlo sum, which can be accomplished as follows. First, we need to fix a reference experiment by drawing the random variable $n_j \equiv (n, h_j)$ from its ensemble distribution, which in Gaussian noise is a normal distribution with mean zero and covariance matrix $F_{ij} \equiv (h_i, h_j)$ [see Eq. (3)]. To do so, we generate a zero-mean, unit-variance, normal N -tuple, and multiply it by $\sqrt{F_{ij}^{-1}}$ (where the square root is taken in the linear-operator sense and exists for nonsingular Fisher matrices). Note that we cannot work with SNR-invariant expressions (e.g., normalized parameter errors $\bar{\theta}_i$), since the priors set a scale for the strength of the signal.

We can now draw samples distributed with the LSA likelihood $p(s|\theta) \propto \exp\{-(h_j, h_k)\theta_j\theta_l/2 + n_j\theta_j\}$. To do so, we generate zero-mean, unit-variance, normal N -tuples, multiply them by $(F_{ij}^{-1})^{1/2}$, and offset them by $F_{ij}^{-1} n^j$. We include the effects of priors (therefore obtaining a population $\{\theta^{(i)}\}$ distributed according to the LSA *posterior* probability) by going through the samples, and discarding each of them with a probability $1 - p(\theta)/\max_\theta p$ (for rectangular priors the probability of discarding is al-

²⁰In particular, ϕ_0 is strongly coupled to t_0 , which produces frequency-dependent phase shifts through the exponential $\exp(2\pi i f t_0)$. Adopting the new phase parameter $\phi'_0 = \phi_0 + 2\pi f_0 t_0$, where f_0 is the *dominant* frequency at which $f^{-7/3}/S_n(f)$ is maximum, largely removes this coupling [22].

TABLE I. Fisher-matrix rms errors in the 4-, 5-, and 6-parameter-estimation problems for a $(10 + 10)M_{\odot}$ binary with $\beta = \sigma = 0$ and $\text{SNR} = 10$, evaluated under different combinations of normal true-parameter-centered (NTC) priors (upper section of table) and of the exact priors of Sec. VB (lower section). The underlined errors are larger than the physical range for the parameter.

	$\Delta M_c/M_c$	$\Delta \eta$	$\Delta \beta$	$\Delta \sigma$	$\Delta \phi_0$	$\Delta t_0(\text{ms})$
4pp, no priors	2.9×10^{-2}	8.3×10^{-2}			7.3	3.0
4pp, NTC prior on η	2.7×10^{-2}	7.9×10^{-2}			7.0	2.8
5pp, no priors	<u>1.1</u>	<u>5.1×10^1</u>	<u>1.2×10^3</u>		3.3×10^2	6.9
5pp, NTC prior on β	3.0×10^{-2}	<u>3.8×10^{-1}</u>	8.5		7.7	3.0
5pp, NTC prior on η, β	3.0×10^{-2}	2.1×10^{-1}	4.9		7.7	3.0
6pp, NTC priors on η, β	4.3×10^{-2}	2.5×10^{-1}	8.4	<u>2.5×10^1</u>	4.3×10^1	5.3
6pp, NTC priors on η, β, σ	3.0×10^{-2}	2.1×10^{-1}	5.1	4.9	1.1×10^1	3.1
4pp, exact priors on η	1.8×10^{-2}	5.0×10^{-2}			4.4	1.9
5pp, exact priors on η, β	2.9×10^{-2}	7.1×10^{-2}	2.4		7.5	2.9
6pp, exact priors on η, β, σ	2.9×10^{-2}	7.1×10^{-2}	2.6	2.9	9.0	3.0

ways 0 or 1). The covariance matrix of the posterior distribution can then be computed from the surviving samples.

We repeat this procedure for many different experiments (i.e., n_j 's), and take a frequentist average of the covariance-matrix components (or study their frequentist distribution). Again, the Bayesian interpretation of this entire procedure is as follows: we place the true signal at $\theta = 0$; we draw from the possible noise realizations according to their ensemble probability; and we compute the variance of the posterior distribution for each noise realization. If the priors are very restrictive compared to Fisher-matrix-only errors, we may find that we are discarding a very large percentage of the samples. To avoid this, we can incorporate a normal approximation to the priors in the probability distribution used to generate the samples (i.e., by multiplying normal variates by $\sqrt{[(h_i, h_j) + P_{ij}]^{-1}}$, and offsetting them by $[(h_i, h_j) + P_{ij}]^{-1}n^j$), and then sieve the resulting samples with respect to $\propto p(\theta)e^{P_{ij}\theta_i\theta_j/2}$ instead of $p(\theta)$. It is also possible to use rejection sampling [23], as we did for the results reported in this section, or the Metropolis algorithm [24] with the likelihood or likelihood-plus-NTC-prior as proposal distribution, and the full posterior as the target distribution.

Applying the procedure outlined above to our $(10 + 10)M_{\odot}$ system yields the results listed in the lower section of Table I. We adopt exact priors given by rectangular probability distributions covering the intervals $[0, \infty]$ for M_c , $[0, 0.25]$ for η , $[-8.5, 8.5]$ for β , and $[-5, 5]$ for σ . Each quoted error is a frequentist average of 200 independent Monte Carlo estimates, each computed for a different realization of noise from an initial sampling of 10^6 parameter sets, reduced to $5 \times 10^4 - 2 \times 10^5$ samples after rejection sampling, depending on the estimation problem.

The expected errors are significantly reduced compared to the NTC-prior estimates. These reductions stem mainly from the greater tightness of the rectangular priors, and are especially significant for η , for which the symmetric NTC prior is indeed very crude. The lesson is that we can use

$[(h_i, h_j) + P_{ij}]^{-1}$ (i.e., the quick test) to decide whether priors are important, but we need something more sophisticated, such as the procedure described in this section, to gauge their effects accurately. Of course, this gain in accuracy may be only virtual if the LSA is not warranted for our problem. Deciding that question is the object of the next section.

VI. THE UNBEARABLE LIGHTNESS OF SIGNAL TO NOISE

As we have seen in Sec. II, the high-SNR and LSA limits coincide because larger signal strengths correspond to smaller statistical errors, which in turn imply that the linearized-signal expression (22) for the likelihood is more accurate. The equivalence of the two limits is manifest in the $1/A$ expansions such as Eqs. (20) and (29). Indeed, Finn [4] cautions that “it is important that the probability contours of interest (e.g. 90%) do not involve [errors] so large that the linearization of [the likelihood] is a poor approximation.”

In practice, given a family of waveforms and the true parameter values, we need to ask how high a SNR is needed for the limits to yield accurate expected errors. One approach involves comparing the Fisher-matrix results with errors computed at the next order in the $1/A$ expansions: in Appendix A 5 of Ref. [5], Cutler and Flanagan provide next-order formulas for the frequentist variance (although they do not apply them to the Fisher-matrix estimates in the same article). In Sec. VII we provide the mathematical tools to do so in our notation; we must, however, warn the reader that the calculation is rather cumbersome, except for simple waveforms, and the verdict is still not definitive: the smallness of a term in a series does not guarantee that the series is converging.

A simpler approach involves working with the ratio $r(\theta, A)$ of the LSA likelihood to the exact likelihood to build a consistency criterion for the Fisher-matrix formalism. In this section we shall see that under reasonable conditions, the ratio r is given in logarithm by

$$|\log r(\theta, A)| = (\theta_j h_j - \Delta h(\theta), \theta_k h_k - \Delta h(\theta))/2, \quad (47)$$

where $\Delta h(\theta) = h(\theta) - h(0)$, A is the signal strength, and θ is the error (in a sense to be made precise shortly). Since $h(\theta) = h(0) + \theta_i h_i + \dots$, the product in Eq. (47) represents the noise-weighted norm of the higher-than-linear contributions to $h(\theta)$, expanded around the true source parameters. The idea of the criterion is to choose an isoprobability surface (say, the 1σ surface), as predicted by the Fisher matrix, and then explore it to verify that the mismatch between the LSA and exact likelihoods is smaller than a fiducial value (say, $|\log r| < 0.1$), so that we can actually believe the LSA in predicting the 1σ surface to begin with.

We stress that this is just a criterion of *consistency*. Even if the Fisher-matrix result is internally consistent, it may still be inaccurate; conversely, the structure of the ambiguity function across parameter space could conspire in such a way as to make the LSA results correct, although we have no reason to expect that in general. In the rest of this section, we explain how the criterion comes about in the frequentist (Sec. VIA) and Bayesian (Sec. VIB) frameworks, and we show a concrete example of the criterion in use (Sec. VIC).

A. Frequentist justification of the maximum-mismatch criterion

As in Secs. II and V, we assume that the detector output is $s = A\bar{h}_0 + n$. In the LSA, the ML estimator $\hat{\theta}^{\text{ML}}$ is a normal variable with mean zero and covariance matrix $(1/A^2)(\bar{h}_j, \bar{h}_k)^{-1}$ [Eq. (3)]. For a given signal amplitude A , let $\hat{\theta}^{\text{ML}}$ take on the specific value $\theta^{1\sigma}$ on its 1σ surface. From Eqs. (1), (22), and (47), the mismatch ratio r is given by

$$r(\theta^{1\sigma}) = \exp\left\{-\frac{(s - A(\bar{h}_0 + \theta_j^{1\sigma}\bar{h}_j), s - A(\bar{h}_0 + \theta_k^{1\sigma}\bar{h}_k))/2}{\exp\left\{-\frac{(s - A\bar{h}(\theta^{1\sigma}), s - A\bar{h}(\theta^{1\sigma}))/2}{2}\right\}}\right\}, \quad (48)$$

writing s out, we eliminate all instances of \bar{h}_0 :

$$\begin{aligned} \log r(\theta^{1\sigma}) &= -A^2\theta_j^{1\sigma}\theta_k^{1\sigma}(\bar{h}_j, \bar{h}_k)/2 \\ &\quad + A^2(\Delta\bar{h}(\theta^{1\sigma}), \Delta\bar{h}(\theta^{1\sigma}))/2 \\ &\quad + A\theta_j^{1\sigma}(\bar{h}_j, n) - A(\Delta\bar{h}(\theta^{1\sigma}), n). \end{aligned} \quad (49)$$

The first two terms in the exponent can be computed given $\theta^{1\sigma}$; not so the two products involving n . To obtain the first, we note that if $\hat{\theta}^{\text{ML}} = \theta^{1\sigma}$, then the noise must be such that $\theta^{1\sigma} = (1/A)(\bar{h}_j, \bar{h}_k)^{-1}(\bar{h}_k, n)$ [Eq. (23)], so $(\bar{h}_j, n) = A(\bar{h}_j, \bar{h}_k)\theta_k^{1\sigma}$. To obtain the second, we change our perspective slightly, and average $\log r(\theta^{1\sigma})$ over all noise realizations n compatible with $\theta^{1\sigma}$. This is how. Let $x_j \equiv (\bar{h}_j, n)$, $y \equiv (\Delta\bar{h}(\theta^{1\sigma}), n)$: separately, x_j and y are normal random variables with mean zero and covariances equal to (\bar{h}_j, \bar{h}_k) and $(\Delta\bar{h}(\theta^{1\sigma}), \Delta\bar{h}(\theta^{1\sigma}))$, respectively; taken together, they are jointly normal variables with covariance $(\bar{h}_j, \Delta\bar{h}(\theta^{1\sigma}))$ [Eq. (3) again]. We now know enough to build $p(x, y)$,

from which we can derive the conditional probability $p(y|x)$ and compute the conditional mean of y , which (after the algebra of Appendix A) turns out to be $A\theta_j^{1\sigma}(\bar{h}_j, \Delta\bar{h}(\theta^{1\sigma}))$. Altogether, we find

$$\begin{aligned} \langle \log r(\theta^{1\sigma}) \rangle_{n(\theta^{1\sigma})} &= A^2\theta_j^{1\sigma}\theta_k^{1\sigma}(\bar{h}_j, \bar{h}_k)/2 + A^2(\Delta\bar{h}(\theta^{1\sigma}), \Delta\bar{h}(\theta^{1\sigma}))/2 \\ &\quad - A^2\theta_j^{1\sigma}(\bar{h}_j, \Delta\bar{h}(\theta^{1\sigma})) \\ &= (\theta_j^{1\sigma}h_j - \Delta h(\theta^{1\sigma}), \theta_k^{1\sigma}h_k - \Delta h(\theta^{1\sigma}))/2, \end{aligned} \quad (50)$$

just as anticipated in Eq. (47). The signal strength A enters Eq. (50) explicitly, but also implicitly through the parameter width of the Fisher-matrix 1σ surface. Thus Eq. (47) can be solved for the A that corresponds to $\theta^{1\sigma}$ small enough to yield r as close to unity as desired. Since to leading order $\theta_j^{1\sigma} - \Delta h(\theta^{1\sigma}) = A\bar{h}_{jk}\theta_j^{1\sigma}\theta_k^{1\sigma}$, and since to leading order $\theta^{1\sigma}$ scales as $1/A$, we expect $\log r$ to scale as $1/A^2$ for large enough A .

In summary, the maximum-mismatch criterion is justified from a frequentist viewpoint as a constraint on the ratio r at points on a constant-LSA-probability surface, averaged over all realizations of noise compatible with finding the ML estimator at those points.

B. Bayesian justification of the maximum-mismatch criterion

The justification of the maximum-mismatch criterion from a Bayesian viewpoint requires another slight change of perspective. Again we assume $s = A\bar{h}_0 + n$; this time, however, we expand the waveform not with respect to the true parameters ($\theta = 0$), but to the observed location $\hat{\theta}^{\text{ML}}(n) \equiv \theta_0^n$ of maximum LSA likelihood for a given experiment. In the absence of priors, it is with respect to this location that the uncertainty of the posterior would be judged in a single experiment. Thus we write $h(\theta) \equiv h(\theta_0^n + \theta^n) \simeq h_0^n + \theta_j^n h_j^n$, where $h_0^n \equiv h(\theta_0^n)$: the “ n ” superscripts serve to remind us that the parameter displacements θ^n (and the waveform derivatives h_j^n) are evaluated from (at) θ_0^n . We also write Δh_0^n for $h_0 - h_0^n$.

The 1σ surface over which we are going to evaluate the ratio r will be a surface of equiprobable true-signal locations, given the observed location θ_0^n of maximum LSA likelihood. In the LSA, the distribution over experiments of the true-signal location with respect to θ_0^n is again normal with covariance matrix $(h_i^n, h_j^n)^{-1}$. Thus we have $\theta^n = -\theta_0^n$, and the mismatch ratio r is given by

$$\begin{aligned} r(\theta_0^n) &= \exp\left\{-\frac{(s - A(\bar{h}_0^n + \theta_j^n\bar{h}_j^n), s - A(\bar{h}_0^n + \theta_k^n\bar{h}_k^n))/2}{\exp\left\{-\frac{(s - A\bar{h}_0^n, s - A\bar{h}_0^n)}{2}\right\}}\right\}, \end{aligned} \quad (51)$$

writing s out, the denominator reduces to $\exp\{-(n, n)/2\}$, and h_0 enters the numerator only through Δh_0^n :

$$\begin{aligned} \log r(\theta_0^n) &= -A^2\theta_j^n\theta_k^n(\bar{h}_j^n, \bar{h}_k^n)/2 - A^2(\Delta\bar{h}_0^n, \Delta\bar{h}_0^n)/2 \\ &\quad + A^2\theta_j^n(\bar{h}_j^n, \Delta\bar{h}_0^n) + A\theta_j^n(\bar{h}_j^n, n) - A(\Delta\bar{h}_0^n, n). \end{aligned} \quad (52)$$

Now, since the LSA likelihood can be written as $p(s|\theta^n) \propto \exp[-(n + \Delta h_0^n - \theta_j^n h_j^n, n + \Delta h_0^n - \theta_j^n h_j^n)/2]$, the ML equation $\partial p/\partial \theta_i^n = 0$ at $\theta_i^n = 0$ implies $(h_j^n, n) = -(h_i^n, \Delta h_0^n)$. We handle the last term of the equation by evaluating the conditional mean of $y^n \equiv -A(\Delta \bar{h}_0^n, n)$ given $x^n = (\bar{h}_j^n, n) = -A(\bar{h}_j^n, \Delta \bar{h}_0^n)$, producing $-A(\bar{h}_j^n, n) \times (\bar{h}_j^n, \bar{h}_k^n)^{-1}(\bar{h}_k^n, \Delta \bar{h}_0^n) = A(\bar{h}_j^n, n)(h_j^n, h_k^n)^{-1}(h_k^n, n)$ (again, see Appendix A). We can then use Eq. (23) to replace $(h_j^n, h_k^n)^{-1}(h_k^n, n)$ with $\theta_{0j}^n = -\theta_j^n$ (working to leading order), so that the last two terms of Eq. (52) end up canceling out:

$$\begin{aligned} \log r(\theta_0) &= -A^2 \theta_j^n \theta_k^n (\bar{h}_j^n, \bar{h}_k^n)/2 + A^2 \theta_j^n (\bar{h}_j^n, \Delta \bar{h}_0^n) \\ &\quad - A^2 (\Delta \bar{h}_0^n, \Delta \bar{h}_0^n)/2 \\ &= -(\theta_j^n h_j^n - \Delta h(\theta^n), \theta_k^n h_k^n - \Delta h(\theta^n))/2. \end{aligned} \quad (53)$$

Again, this equation can be solved for the A that corresponds to 1σ true-signal locations θ^n small enough to yield r close to unity. Interestingly, the signs of the frequentist and Bayesian expressions (50) and (53) are opposite, indicating (at least *prima facie*) that the likelihood is overestimated in the frequentist case, underestimated in the Bayesian case. Given the conditions under which we have obtained Eqs. (50) and (53), it is perhaps best to consider only their absolute value as rough indicators of the appropriateness of the high-SNR/LSA limit.

In summary, the maximum-mismatch criterion is justified from a Bayesian viewpoint by fixing the location of maximum LSA likelihood, and then exploring a surface of equiprobable true-signal locations, evaluating for each the average of $\log r$ over all experiments (i.e., realizations of noise) compatible with having the true signal there.

C. Practical usage of the maximum-mismatch criterion

In both the frequentist and the Bayesian pictures, Eq. (47) yields the noise-averaged logarithm of the likelihood mismatch, $|\log r|$, as a function of the signal strength A and of a direction in parameter space that identifies a point on the 1σ surface, given by the solutions of the LSA equation $A^2(\bar{h}_j, \bar{h}_k)\theta_j\theta_k = 1$, and interpreted as equiprobable locations for the ML estimator given the true signal $\theta = 0$ (in the frequentist picture), or for the true signal given the mode of the likelihood at $\theta = 0$ (in the Bayesian picture). We use Eq. (47) by fixing the signal strength to what is reasonably expected in observations, perhaps close to the minimum detection SNR, although the astronomical distribution and intrinsic strengths of sources may prompt other choices (e.g., the supermassive-black-hole binaries to be observed by LISA have typical SNRs in the hundreds); and then by evaluating $|\log r|$ as a function of direction in parameter space.

Figure 1 shows an example of this procedure for a very simple and benign one-dimensional estimation problem (a sinusoid of known amplitude and frequency in Gaussian stationary noise), where the only parameter left to estimate

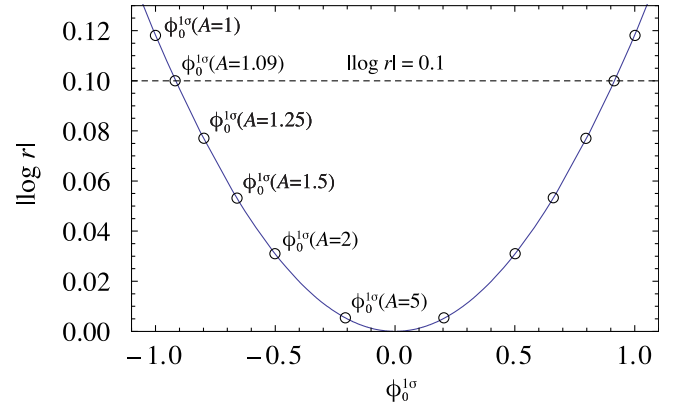


FIG. 1 (color online). Consistency criterion for the simple waveform model $h(\phi_0) = A \cos(2\pi f t + \phi_0)$ in Gaussian stationary noise, with fixed (known) A and f . The curve plots $|\log r|$ as a function of the 1σ error $\phi_0^{1\sigma} = \pm 1/A$, with specific values of A called out by the circles. For the threshold $|\log r| = 0.1$, consistency is achieved for $A \geq 1.09$. (To generate this graph, the integration time was set to $1000/f$, and the variance of noise adjusted so that $(\bar{h}(\phi_0), \bar{h}(\phi_0)) = 1$.)

is the initial phase ϕ_0 ($= 0$ for the true signal). For each value of $\text{SNR} \equiv A$, the expected 1σ surface consists of just the two points $\phi_0^{1\sigma} = \pm 1/A$. Figure 1 shows $|\log r|$ as a function of $\phi_0^{1\sigma}$, and therefore of A . If we set a threshold of $|\log r| = 0.1$ (the dashed line) to claim the high-SNR/LSA limit as consistent, we see that the consistency criterion is not satisfied for $A = 1$, where $|\log r| \approx 0.12$, but it begins to be satisfied for $A \geq 1.09$. Once again, for a given SNR, $|\log r|$ at $\phi_0^{1\sigma}$ is an index of the closeness of the LSA and exact likelihoods at typical values of the errors, and averaged among compatible noise realizations.

The principle is the same for multiparameter estimation problems, where we have the additional task of sampling the entire 1σ surface in a manner consistent with the LSA distribution at 1σ . One way to do so is to obtain the eigenvalues $\lambda^{(i)}$ and eigenvectors $\theta_j^{(i)}$ of (\bar{h}_j, \bar{h}_k) , and then sample the parameter values $\theta = \sum_{(i)=1}^N \tilde{c}^{(i)} \theta_j^{(i)} / (A\sqrt{\lambda^{(i)}})$, with $\tilde{c}^{(i)}$ distributed uniformly on the N -dimensional unit sphere.²¹ We then obtain the cumulative distribution func-

²¹To see why this is the right thing to do, consider the integration of a function against the LSA distribution, and make a change of variables (with unit Jacobian) to eigenvalue components, and a second to rescaled components $\tilde{c}^{(i)} = A\sqrt{\lambda^{(i)}} c^{(i)}$:

$$\begin{aligned} &\int (\dots) \exp\{-A^2 \theta_i \theta_j (\bar{h}_i, \bar{h}_j)/2\} d\theta \\ &= \int (\dots) \exp\{-A^2 \sum_{(i)} \lambda^{(i)} [c^{(i)}]^2/2\} dc \propto \int (\dots) e^{-\sum_{(i)} [\tilde{c}^{(i)}]^2/2} d\tilde{c}; \end{aligned}$$

we see that the source parameters that correspond to \tilde{c} lying on a sphere of fixed radius must lie on an isoprobability surface. To reassemble θ from the \tilde{c} , we need to *divide* the eigenvectors by $A\sqrt{\lambda^{(i)}}$.

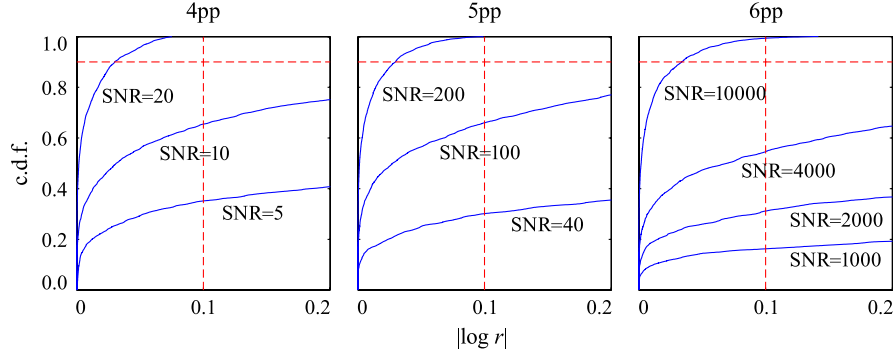


FIG. 2 (color online). Cumulative distribution function for $|\log r|$ on the 1σ surface at various SNRs for our reference SPA model with $m_1 = m_2 = 10M_\odot$. The SNR required to have 90% of the 1σ points at $|\log r| = 0.1$ (dashed lines) increases considerably (in fact, to unrealistic values) as we move from the 4pp to 5pp and 6pp. This figure was produced without imposing any priors on the source parameters.

tion for the values of $|\log r|$, which we plot in Fig. 2 for our reference model. If we consider the high-SNR/LSA limit to be sufficiently realized when $|\log r| < 0.1$ over 90% of the 1σ surface, we conclude that the Fisher-matrix formalism (with no priors) is self-consistent for SNRs between 10 and 20 in the 4-parameter problem, between 100 and 200 in the 5pp, and between 4000 and 10 000 in the 6pp.

The eigenvector directions that push the required SNR toward higher values are usually those associated with the smallest-magnitude eigenvalues. To confirm this, and to get some clues about the beyond-LSA structure of the likelihood, we can fix the maximum acceptable value of $|\log r|$ (say, again to 0.1) and then solve Eq. (47) for A as a function of direction in parameter space. We do so for the 4pp in Fig. 3, where we show all two-dimensional parameter subspaces along pairs of eigenvectors (strictly speaking, we are not sampling a single 1σ surface, but considering the set of such surfaces for all SNRs, and determining on which of them $|\log r| = 0.1$, as a function of parameter angle).

In the Bayesian framework, it is also possible to combine the maximum-mismatch criterion with the normal prior test of Sec. VA, by investigating the values of $|\log r|$ on the 1σ surface given by the solutions of the LSA-cum-prior equation $[(h_j, h_k) + P_{jk}]\theta_j\theta_k = 1$. This test can help decide whether the LSA is warranted once priors are factored in: it can be shown that Eq. (53) continues to hold with NTC priors, although its interpretation is not as clean, because their mode moves around the 1σ surface as we explore it. The maximum-mismatch criterion may indicate that, at the signal strengths of interest, the LSA becomes consistent only with the priors; in that case, the reliable predictor of source parameter accuracy would not be $[(h_j, h_k) + P_{jk}]^{-1}$ (given the crudeness of NTC priors), but rather the result of a LSA Monte Carlo procedure such as that described in Sec. VB.

It is very hard to make a general statement about the errors in the expected accuracies when the LSA Fisher-

matrix result is not self-consistent. Such errors are strongly SNR-dependent, and it is usually necessary to include parameter priors into consideration. As anecdotal evidence, I offer that for our reference model at $\text{SNR} = 10$, a full-blown Monte Carlo sampling of the posterior distribution, involving an explicit time-domain realization of noise and adopting the priors of Sec. VB, reports posterior variances that differ from the last three rows of Table I by a few tens percent for the 4pp, and by factors of a few (for M_c and η only, since $\Delta\beta$ and $\Delta\sigma$ are dominated by the priors) for the 5pp and 6pp.

In conclusion, I submit that graphs like those of Fig. 2 can be useful to assess the consistency of the “straight” Fisher-matrix formalism, are easy to produce with little additional machinery, and should be included in all articles that use the formalism to predict the future parameter-estimation performance of GW observations. If a single number must be quoted, it could be the SNR at which 90% of the 1σ surface yields $|\log r| < 0.1$.

VII. BEYOND THE LINEARIZED-SIGNAL APPROXIMATION

In this section we develop mathematical tools to derive higher-than-LSA expressions for the frequentist mean and variance of the ML estimator over an ensemble of noise realizations (Sec. VII A), and for the Bayesian mean and variance of the posterior distribution (without priors) in a single experiment (Sec. VII B). These expressions provide corrections to the Fisher-matrix result, and can therefore be used to check its accuracy, as suggested by Cutler and Flanagan [5], who derive a general expression for the $1/A^4$ correction to the frequentist variance. A formal treatment of the $1/A$ expansion for the frequentist moments can be found in Barndorff-Nielsen and Cox [25] and in Zanolin, Naftali, and Makris [26], who use the expansion to determine conditions for the ML estimate to become unbiased and attain the Cramér-Rao bound [27].

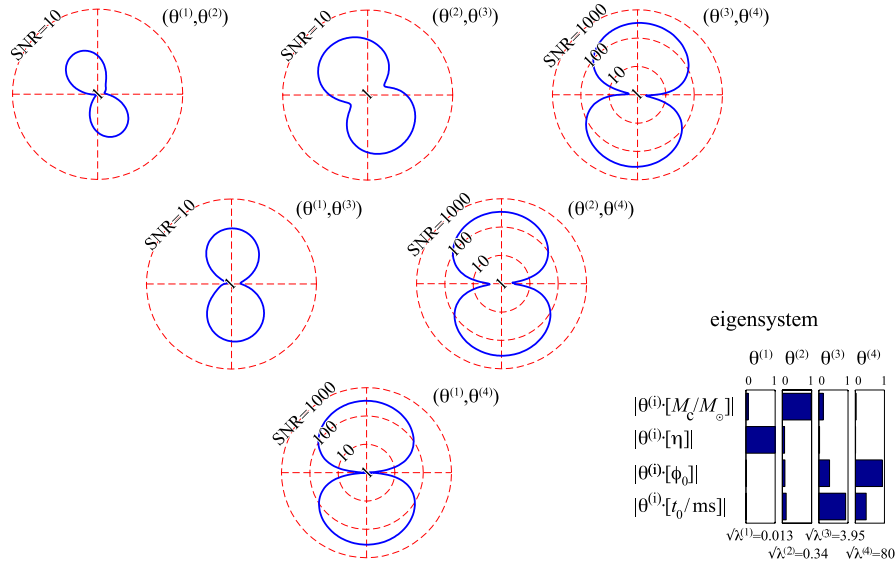


FIG. 3 (color online). SNR values at which $|\log r| = 0.1$ in the 4pp for our reference SPA model with $m_1 = m_2 = 10M_\odot$. The six subplots display sections of parameter space corresponding to all distinct pairs of eigenvectors of the Fisher matrix; the polar radius of the curves shows the required SNR (plotted logarithmically from $\log \text{SNR} = 0$), while the polar angle is computed between pairs of renormalized eigenvector coefficients $\tilde{c}^{(i)}$. The graph at the bottom right shows the composition of the four Fisher-matrix eigenvectors in terms of the source parameters, as well as their respective eigenvalues. This figure was produced without imposing any priors on the source parameters.

However, computing higher-order corrections involves a considerable amount of tensor algebra that calls for the use of specialized software, such as MATHTENSOR [28]; they also involve higher-than-first derivatives of the waveforms and products of several inverse Fisher matrices, which may raise concerns about the numerical accuracy of the computations. Throughout this section, we distinguish between covariant and contravariant indices (as in $n_i = (\bar{h}_i, n)$ and θ^i , respectively); in fact, we find it convenient to use the inverse normalized Fisher matrix $(\bar{h}^i, \bar{h}^j)^{-1}$ to raise indices, therefore hiding its repeated appearance in tensor expressions.

A. In the frequentist framework

Using the $1/A$ expansion of Eqs. (18) and (19), the perturbative ML equations can be written in general as

$$\begin{aligned}
 H_{ij}\theta_{(1)}^j &= N_i, & H_{ij}\theta_{(2)}^j &= N_{ij}\theta_{(1)}^j - H_{ijk}\theta_{(2)}^{jk}, \\
 H_{ij}\theta_{(3)}^j &= N_{ij}\theta_{(2)}^j + N_{ijk}\theta_{(2)}^{jk} - H_{ijk}\theta_{(3)}^{jk} - H_{ijkl}\theta_{(3)}^{jkl}, \\
 H_{ij}\theta_{(4)}^j &= N_{ij}\theta_{(3)}^j + N_{ijk}\theta_{(3)}^{jk} + N_{ijkl}\theta_{(3)}^{jkl} - H_{ijk}\theta_{(4)}^{jk} \\
 &\quad - H_{ijkl}\theta_{(4)}^{jkl} - H_{ijklm}\theta_{(3)}^{klm}, \quad \dots
 \end{aligned} \tag{54}$$

with $N_i = (n, \bar{h}_i)/0!$, $N_{ij} = (\bar{h}_{ij}, n)/1!$, $N_{ijk} = (\bar{h}_{ijk}, n)/2!$ (and so on), and

$$\begin{aligned}
 H_{ij} &= \frac{\bar{h}_{i,j}}{0!1!}, & H_{ijk} &= \frac{\bar{h}_{i,jk}}{0!2!} + \frac{\bar{h}_{ij,k}}{1!1!}, \\
 H_{ijkl} &= \frac{\bar{h}_{i,jkl}}{0!3!} + \frac{\bar{h}_{ij,kl}}{1!2!} + \frac{\bar{h}_{ijk,l}}{2!1!}, \quad \dots
 \end{aligned} \tag{55}$$

where $\bar{h}_{i,j} = (\bar{h}_i, \bar{h}_j)$, $\bar{h}_{i,jk} = (\bar{h}_i, \bar{h}_{jk})$ (and so on), and where the two factorials at each denominator are those (respectively) of the number of indices before the comma minus one, and of the number of indices after the comma. Also, the $\theta_{(n)}^j$ of Eq. (54) are the unknown $1/A^n$ contributions to $\hat{\theta}^{\text{ML}}$ (as in Eq. (19), dropping hats for simplicity), while the multi-index parameter objects such as $\theta_{(2)}^{jk}$ are given by

$$\begin{aligned}
 \theta_{(2)}^{jk} &= \theta_{(1)}^j \theta_{(1)}^k, & \theta_{(3)}^{jk} &= \theta_{(1)}^j \theta_{(2)}^k + \theta_{(2)}^j \theta_{(1)}^k, \\
 \theta_{(4)}^{jk} &= \theta_{(1)}^j \theta_{(3)}^k + \theta_{(2)}^j \theta_{(2)}^k + \theta_{(3)}^j \theta_{(1)}^k, & \theta_{(3)}^{jkl} &= \theta_{(1)}^j \theta_{(1)}^k \theta_{(1)}^l, \\
 \theta_{(4)}^{jkl} &= \theta_{(1)}^j \theta_{(1)}^k \theta_{(2)}^l + \theta_{(1)}^j \theta_{(2)}^k \theta_{(1)}^l + \theta_{(2)}^j \theta_{(1)}^k \theta_{(1)}^l, \\
 \theta_{(4)}^{klm} &= \theta_{(1)}^k \theta_{(1)}^l \theta_{(1)}^m,
 \end{aligned} \tag{56}$$

and so on. In general, the object $\theta_{(n)}^{j_1 \dots j_m}$ will consist of as many addends as there are partitions of n into m integers, including all permutations of each partition. For instance, the $n = 5, m = 3$ object $\theta_{(5)}^{jkl}$ would have terms for each of the partitions $1 + 1 + 3, 1 + 3 + 1, 3 + 1 + 1, 1 + 2 + 2, 2 + 1 + 2, 2 + 2 + 1$.

The solution of each equation in Eq. (54) is trivial given the solutions of all equations of lower order. Since the inverse matrix $(H^{-1})^{ij} \equiv (\bar{h}_i, \bar{h}_j)^{-1} = A^2 F_{ij}^{-1}$ appears multiple times in the solutions (because H_{ij} multiplies the unknown $\theta_{(n)}^j$ in each equation), it is convenient to adopt a compact notation that hides the $(H^{-1})^{ij}$ by raising every index into which they are contracted. We then find

$$\begin{aligned}\theta_{(1)}^i &= N^i, \\ \theta_{(2)}^i &= N^i_j \theta_{(1)}^j - H^i_{jk} \theta_{(2)}^{jk} = N^i_j N^j - H^i_{jk} N^j N^k, \\ \theta_{(3)}^i &= N^i_j \theta_{(2)}^j + N^i_{jk} \theta_{(1)}^j \theta_{(1)}^k \\ &\quad - H^i_{jk} (\theta_{(1)}^j \theta_{(2)}^k + \theta_{(2)}^j \theta_{(1)}^k) - H^i_{jkl} \theta_{(1)}^j \theta_{(1)}^k \theta_{(1)}^l = \dots\end{aligned}\quad (57)$$

The frequentist mean and covariance of the θ^i can be built from these solutions, remembering the Wick-product rule [29] for the ensemble average of products of Gaussian variables:

$$\begin{aligned}\langle \theta^i \theta^j \rangle_n - \langle \theta^i \rangle_n \langle \theta^j \rangle_n &= \left[\frac{1}{A^2} \langle \theta_{(1)}^i \theta_{(1)}^j \rangle_n + \frac{1}{A^3} \langle \theta_{(1)}^i \theta_{(2)}^j + \theta_{(2)}^i \theta_{(1)}^j \rangle_n + \frac{1}{A^4} \langle \theta_{(1)}^i \theta_{(3)}^j + \theta_{(2)}^i \theta_{(2)}^j + \theta_{(3)}^i \theta_{(1)}^j \rangle_n + \dots \right] \\ &\quad - \left[\frac{1}{A^2} \langle \theta_{(1)}^i \rangle_n \langle \theta_{(1)}^j \rangle_n + \frac{1}{A^3} (\langle \theta_{(1)}^i \rangle_n \langle \theta_{(2)}^j \rangle_n + \langle \theta_{(2)}^i \rangle_n \langle \theta_{(1)}^j \rangle_n) \right. \\ &\quad \left. + \frac{1}{A^4} (\langle \theta_{(1)}^i \rangle_n \langle \theta_{(3)}^j \rangle_n + \langle \theta_{(2)}^i \rangle_n \langle \theta_{(2)}^j \rangle_n + \langle \theta_{(3)}^i \rangle_n \langle \theta_{(1)}^j \rangle_n) + \dots \right],\end{aligned}\quad (60)$$

where all the underlined terms vanish because they are proportional to ensemble products of an odd number of n terms. The surviving contributions are given by

$$\begin{aligned}\langle \theta_{(1)}^i \theta_{(1)}^j \rangle_n &= \langle N^i N^j \rangle_n = \bar{h}^{i,j}, \\ \langle \theta_{(1)}^i \theta_{(3)}^j \rangle_n &= \langle N^i N^j_k N^k N^l N^l \rangle_n - H^k_{lm} \langle N^i N^j_k N^l N^m \rangle_n + \langle N^i N^j_{kl} N^k N^l \rangle_n - H^j_{klm} \langle N^i N^k N^l N^m \rangle_n \\ &\quad - H^j_{kl} (\langle N^i N^k N^l N^m \rangle_n - H^l_{mq} \langle N^i N^k N^m N^q \rangle_n + \langle N^i N^k_m N^m N^l \rangle_n - H^k_{mq} \langle N^i N^m N^q N^l \rangle_n), \\ \langle \theta_{(2)}^i \theta_{(2)}^j \rangle_n &= \langle N^i_k N^k N^j_m N^m \rangle_n - H^i_{kl} \langle N^k N^l N^j_m N^m \rangle_n - H^j_{mq} \langle N^m N^q N^i_k N^k \rangle_n + H^i_{kl} H^j_{mq} \langle N^k N^l N^m N^q \rangle_n,\end{aligned}\quad (61)$$

and of course $\langle \theta_{(3)}^i \theta_{(1)}^j \rangle_n = \langle \theta_{(1)}^i \theta_{(3)}^j \rangle_n$. The first of these equations reproduces the standard Fisher-matrix result. The four- N products in Eq. (61) follow from Eq. (58). For instance, the last two products are given by

$$\begin{aligned}\langle N^k N^l N^m N^q \rangle_n &= \langle N^k N^l \rangle_n \langle N^m N^q \rangle_n + \langle N^k N^m \rangle_n \langle N^l N^q \rangle_n \\ &\quad + \langle N^k N^q \rangle_n \langle N^l N^m \rangle_n \\ &= \bar{h}^{k,l} \bar{h}^{m,q} + \bar{h}^{k,m} \bar{h}^{l,q} + \bar{h}^{k,q} \bar{h}^{l,m}, \\ \langle N^m N^q N^i_k N^k \rangle_n &= \bar{h}^{m,q} \bar{h}^{i,k}_k + \bar{h}^{m,i}_k \bar{h}^{q,k}_k + \bar{h}^{m,k} \bar{h}^{q,i}_k.\end{aligned}\quad (62)$$

These expressions can be substituted into those of Eq. (61), and those into Eq. (60), yielding the frequentist variance to order $1/A^4$. Unfortunately, this requires computing second- and third-order waveform derivatives (the latter for H_{ijklm}).

$$\langle (a, n) \rangle_n = 0,$$

$$\langle (a, n)(b, n) \rangle_n = (a, b),$$

$$\langle (a, n)(b, n)(c, n) \rangle_n = 0,$$

$$\begin{aligned}\langle (a, n)(b, n)(c, n)(d, n) \rangle_n &= (a, b)(c, d) + (a, c)(b, d) \\ &\quad + (a, d)(b, c)\end{aligned}$$

$$\dots \quad (58)$$

(for any signals a, b, c , and d), where all the products with an odd number of factors vanish, while the products with an even number of factors are given by the sum of terms corresponding to all distinct pairings of signals into inner products. Thus we find that all the $\langle \theta_{\text{odd}(k)}^i \rangle_n$ vanish, while the first nonzero correction to $\langle \theta^i \rangle_n$ is

$$\langle \theta_{(2)}^i \rangle_n = \langle N^i_{,j} N^j \rangle_n - H^i_{jk} \langle N^j N^k \rangle_n = \bar{h}^{ij}_{,j} - H^i_{jk} \bar{h}^{j,k}. \quad (59)$$

As for the covariance,

B. In the Bayesian framework

To generalize Eq. (34), we write

$$I^{(n)} = \int \bar{\theta}^i \frac{\partial^n p(0)}{\partial \epsilon^n} d\theta, \quad \mathcal{N}^{(n)} = \int \frac{\partial^n p(0)}{\partial \epsilon^n} d\theta, \quad (63)$$

and find the recurrence relation

$$\langle \bar{\theta}^i \rangle_p^{(n)} = \left(I^{(n)} - \sum_{j=1}^n \left[\binom{n}{j} \langle \bar{\theta}^i \rangle_p^{(n-j)} \times \mathcal{N}^{(j)} \right] \right) / \mathcal{N}^{(0)}, \quad (64)$$

which we may prove by expanding the identity $I^{(0)} + \epsilon I^{(1)} + \frac{\epsilon^2}{2} I^{(2)} + \dots = [\langle \bar{\theta}^i \rangle_p^{(0)} + \epsilon \langle \bar{\theta}^i \rangle_p^{(1)} + \dots] \times [\mathcal{N}^{(0)} + \epsilon \mathcal{N}^{(1)} + \dots]$ on both sides as a series of ϵ , leading to

$$I^{(n)} = \sum_{j=0}^n \binom{n}{j} \mathcal{N}^{(j)} \langle \bar{\theta}^i \rangle_p^{(n-j)}, \quad (65)$$

whence Eq. (64). To obtain all needed derivatives with

respect to ϵ , we rewrite Eq. (29) as

$$p(s|\theta) \propto \exp\{-(n, n)/2 + [N_i \bar{\theta}^i + \epsilon N_{ij} \bar{\theta}^i \bar{\theta}^j + \epsilon^2 N_{ijk} \bar{\theta}^i \bar{\theta}^j \bar{\theta}^k + \dots] - [H_{jk} \bar{\theta}^j \bar{\theta}^k + \epsilon H'_{jkl} \bar{\theta}^j \bar{\theta}^k \bar{\theta}^l + \epsilon^2 H'_{jklm} \bar{\theta}^j \bar{\theta}^k \bar{\theta}^l \bar{\theta}^m + \dots]/2\}, \quad (66)$$

where $N_i = (n, \bar{h}_i)/1! = n_i$, $N_{ij} = (\bar{h}_{ij}, n)/2! = n_{ij}/2$, $N_{ijk} = (\bar{h}_{ijk}, n)/3!$ (and so on), and also the $H'_{j_1 \dots j_n}$ have slightly different denominators than the $H_{j_1 \dots j_n}$ of Eq. (55):

$$H'_{ijk} = \frac{\bar{h}_{i,jk}}{1!2!} + \frac{\bar{h}_{i,j,k}}{2!1!}, \quad (67)$$

$$H'_{ijkl} = \frac{\bar{h}_{i,jkl}}{1!3!} + \frac{\bar{h}_{ij,kl}}{2!2!} + \frac{\bar{h}_{ijk,l}}{3!1!}; \quad \dots$$

namely, the denominator is $m!l!$ for the product $h_{j_1 \dots j_m, j_1 \dots j_l} \equiv (\bar{h}_{j_1 \dots j_m}, \bar{h}_{j_1 \dots j_l})$. Expanding as a series of ϵ yields

$$p(s|\theta) \propto e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2 + n_j \bar{\theta}^j} \times \left\{ 1 + \epsilon \left(N_{jk} \bar{\theta}^j \bar{\theta}^k - \frac{1}{2} H'_{jkl} \bar{\theta}^j \bar{\theta}^k \bar{\theta}^l \right) + \frac{\epsilon^2}{2} \left[\left(N_{jk} \bar{\theta}^j \bar{\theta}^k - \frac{1}{2} H'_{jkl} \bar{\theta}^j \bar{\theta}^k \bar{\theta}^l \right)^2 + 2 \left(N_{jkl} \bar{\theta}^j \bar{\theta}^k \bar{\theta}^l - \frac{1}{2} H'_{jklm} \bar{\theta}^j \bar{\theta}^k \bar{\theta}^l \bar{\theta}^m \right) \right] + \dots \right\}, \quad (68)$$

so that the $\mathcal{I}^{(n)}$ and $\mathcal{N}^{(n)}$ are given by expressions akin to

$$\mathcal{I}^{(1)} / \mathcal{N}^{(0)} = \int \bar{\theta}^i \left(N_{jk} \bar{\theta}^j \bar{\theta}^k - \frac{1}{2} H'_{jkl} \bar{\theta}^j \bar{\theta}^k \bar{\theta}^l \right) \times e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2 + n_j \bar{\theta}^j} d\theta / \int e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2 + n_j \bar{\theta}^j} d\theta. \quad (69)$$

Now, the integrals of the general form

$$\langle \bar{\theta}^i \dots \bar{\theta}^m \rangle_p^{(0)} = \int \bar{\theta}^i \dots \bar{\theta}^m e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2 + n_j \bar{\theta}^j} d\theta / \int e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2 + n_j \bar{\theta}^j} d\theta, \quad (70)$$

can be computed with the Wick identity²² [29]

²²Another way to organize this computation is to offset the integration variable $\bar{\theta}^j$ to $\bar{\theta}^j - (H^{jk})^{-1} n_k = \bar{\theta}^j - n^j$ in Eq. (70), obtaining

$$\langle \bar{\theta}^i \dots \bar{\theta}^m \rangle_p^{(0)} = \int (\bar{\theta}^i + n^i) \dots (\bar{\theta}^m + n^m) e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2} d\theta / \int e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2} d\theta;$$

we can then expand the product in the integrand, bring the n^i outside the integral, and apply Wick's theorem [Eq. (58)] to obtain each addend of the form

$$n^i \dots n^{i_{m-1}} \int \bar{\theta}^i \dots \bar{\theta}^i e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2} d\theta / \int e^{-H_{jk} \bar{\theta}^j \bar{\theta}^k / 2} d\theta;$$

all integrals with odd l are zero, while the integrals with even l are given by the sum of all possible pairings of indices into products of $(H^{\dots})^{-1}$.

$$\langle F(\bar{\theta}) \rangle_p^{(0)} = F\left(\frac{\partial}{\partial n}\right) \exp\{n_i (H^{ij})^{-1} n_j / 2\}; \quad (71)$$

in particular (again using $(H^{ij})^{-1}$ to raise indices),

$$\langle \bar{\theta}^i \rangle_p^{(0)} = n^i, \quad \langle \bar{\theta}^i \bar{\theta}^j \rangle_p^{(0)} = (H^{ij})^{-1} + n^i n^j, \\ \langle \bar{\theta}^i \bar{\theta}^j \bar{\theta}^k \rangle_p^{(0)} = (H^{ij})^{-1} n^k + (H^{ik})^{-1} n^j + n^i (H^{jk})^{-1} + n^i n^j n^k, \quad \dots \quad (72)$$

Unfortunately, the $1/A^4$ (i.e., ϵ^2) corrections to the variance turn out to be rather unwieldy, and belong in a symbolic-manipulation software package rather than on these pages. We content ourselves with the $1/A^2$ correction to the posterior mean (remember that the normalized parameters $\bar{\theta}$ carry an A),

$$\langle \bar{\theta}^i \rangle_p = n^i + \epsilon [n^i_k n^k - (\frac{1}{2} \bar{h}^i_{kl} + \bar{h}^i_{k,l}) (n^k n^l + \bar{h}^{kl})] + O(\epsilon^2), \quad (73)$$

and the $1/A^3$ correction to the variance,

$$\langle \bar{\theta}^i \bar{\theta}^j \rangle_p - \langle \bar{\theta}^i \rangle_p \langle \bar{\theta}^j \rangle_p = \bar{h}^{ij} + \epsilon [n^{ij} + \frac{1}{2} n^i n^j n_k (n^k n^l - \bar{h}^{kl}) - n^k (\bar{h}^{i,j}_k + \bar{h}^{j,i}_k + \bar{h}^{ij}_k) - n^i n^j n^k (\frac{1}{2} \bar{h}^l_{k,i} + \bar{h}^l_{kl} + \frac{1}{2} \bar{h}^l_{k,lm} n^l n^m)] + O(\epsilon^2). \quad (74)$$

Thus we see that the $1/A^3$ contribution to the variance does not vanish in any single experiment (unless $n^i = 0$). It does vanish, however, under frequentist average, since it involves products of odd numbers of noises.

VIII. CONCLUSION

In this article I tried to provide, as it were, a user's manual for the Fisher information matrix. It seems clear that the Fisher-matrix formalism will continue to be featured prominently in research dealing with the parameter-estimation prospects of future GW observations, because of its compactness and accessibility, and because of the difficulty of computing higher-order corrections and running full-blown simulations. Yet the three questions posed in the introduction loom over the credibility of Fisher-matrix results, which is all the more worrisome when these results are used to justify choices in science policy or experiment design.

The recipes provided in this paper to answer the initial questions can help assert (or falsify) the accuracy of the formalism for specific signal models. In particular:

- (1) As discussed in Sec. IV, ill-conditioned or singular Fisher matrices point to the need for increased numerical accuracy, and occasionally to a case for discarding a parameter or combination of parameters, but more often to suspicions about the appropriateness of the high-SNR/LSA limit. Section IVA

describes how to use the singular-value decomposition of the Fisher matrix to discard truly degenerate linear combinations of parameters; Sec. IV B describes how to roughly assess the sensitivity of the Fisher-matrix inverse to numerical error by means of the Fisher-matrix condition number, and more carefully by a simple Monte Carlo test.

- (2) The necessity of including prior distributions for the source parameters, perhaps in as simple a form as uniform distributions over the physically allowed ranges, can be roughly assessed by verifying whether Fisher-matrix results change with the addition of simple Gaussian priors, as shown in Sec. V A; more accurate estimates of the effect of priors can be obtained by integrating the variance of an exact-prior-LSA-likelihood posterior with the simple Monte Carlo algorithm of Sec. V B.
- (3) The detected-signal strength (i.e., the SNR) necessary for Fisher-matrix results to be internally consistent can be evaluated with the likelihood-mismatch criterion that follows from Eqs. (50) and (53) of Sec. VI, or (at the price of some algebra) by computing the higher-order corrections presented in Sec. VII.

If the Fisher-matrix formalism remains inconsistent at the SNRs of interest, even with the help of priors, there is little recourse but to embark in explicit Monte Carlo simulations of frequentist [6] or Bayesian [2] parameter estimation. Such simulations can consistently include sophisticated priors, and explore the secondary maxima of the posterior (or likelihood, in the frequentist case). They are the gold standard of this trade, but as such they are expensive in human effort and CPU resources. The recipes given in this paper can help establish when they are truly needed.

ACKNOWLEDGMENTS

I would like to thank Yanbei Chen, Curt Cutler, Yi Pan, and Michele Zanolin for useful discussions; for reviewing this manuscript, I am grateful to John Armstrong, Emanuele Berti, Steve Drasco, Frank Estabrook, Sam Finn, Éanna Flanagan, and especially Alessandra Buonanno. My work was supported by the LISA Mission Science Office and by the Human Resources Development Fund at the Jet Propulsion Laboratory, California Institute of Technology, where it was performed under contract with the National Aeronautics and Space Administration.

APPENDIX: LEMMA FOR THE CONDITIONAL AVERAGE OF JOINTLY NORMAL RANDOM VARIABLES

Assume the vector x_j and the scalar y are jointly normal random variables with mean zero and covariance matrix

$$\mathcal{C} = \begin{pmatrix} F_{ij} & H_i \\ H_j & G \end{pmatrix}. \quad (\text{A1})$$

From the standard Frobenius-Schur formula for the inverse of a block matrix [30],

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}^{-1} = \begin{pmatrix} A^{-1} + A^{-1}BS_A^{-1}CA^{-1} & -A^{-1}BS_A^{-1} \\ -S_A^{-1}CA^{-1} & S_A^{-1} \end{pmatrix} \quad (\text{A2})$$

(with $S_A = D - CA^{-1}B$ the *Schur complement* of A), we find

$$\mathcal{C}^{-1} = \begin{pmatrix} F_{ij}^{-1} + S_A^{-1}(F_{ik}^{-1}H_k)(F_{jl}^{-1}H_l) & -S_A^{-1}(F_{jk}^{-1}H_k) \\ -S_A^{-1}(F_{ik}^{-1}H_k) & S_A^{-1} \end{pmatrix}, \quad (\text{A3})$$

since in our case F_{ij}^{-1} is symmetric and S_A is the scalar $G - (F_{ij}^{-1}H_iH_j)$. Now, the joint distribution of x_j and y is given by

$$p(x, y) \propto \exp\left\{-\frac{1}{2}(x_i y) \cdot \mathcal{C}^{-1} \cdot \begin{pmatrix} x_i \\ y \end{pmatrix}\right\}, \quad (\text{A4})$$

while the conditional distribution of y given x_j is $p(y|x) = p(x, y)/p(x) = p(x, y)/[\int p(x, y)dy]$. Since, however, $p(x)$ can be a function only of x , by the properties of Gaussian integrals it must be that $p(x) \propto \exp(\dots)_{ij}x_ix_j$. It follows that $p(y|x)$ must be of the form

$$p(y|x) \propto \exp\left\{-\frac{1}{2}S_A^{-1}y^2 - 2S_A^{-1}(x_iF_{ij}^{-1}H_j)y + (\dots)_{ij}x_ix_j\right\}/2, \quad (\text{A5})$$

from which, by inspection, we conclude that

$$\langle y \rangle_{x_i} = \int y p(y|x) dy = x_i F_{ij}^{-1} H_j \quad (\text{A6})$$

and that

$$\text{var}_{x_i, y} = \int (y - \langle y \rangle_{x_i})^2 p(y|x) dx = S_A = G - F_{ij}^{-1} H_i H_j. \quad (\text{A7})$$

[1] E. T. Jaynes and G. L. Bretthorst, *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, England, 2003).

[2] See, e.g., N. Christensen and R. Meyer, Phys. Rev. D **58**, 082001 (1998); N. Christensen, R. J. Dupuis, G. Woan, and R. Meyer, Phys. Rev. D **70**, 022001 (2004).

- [3] D. Nicholson and A. Vecchio, *Phys. Rev. D* **57**, 4588 (1998).
- [4] L. S. Finn, *Phys. Rev. D* **46**, 5236 (1992).
- [5] C. Cutler and É.E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [6] R. Balasubramanian, B.S. Sathyaprakash, and S.V. Dhurandhar, *Phys. Rev. D* **53**, 3033 (1996).
- [7] R. Balasubramanian and S.V. Dhurandhar, *Phys. Rev. D* **57**, 3408 (1998).
- [8] L. A. Wainstein and L. D. Zubakov, *Extraction of Signals From Noise* (Prentice-Hall, Englewood Cliffs, NJ, 1962).
- [9] A. V. Oppenheim, A. S. Willsky, and I. T. Young, *Signals and Systems* (Prentice-Hall, Englewood Cliffs, NJ, 1983).
- [10] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory* (Prentice-Hall, Englewood Cliffs, NJ, 1993).
- [11] G.L. Bretthorst, *Bayesian Spectrum Analysis and Parameter Estimation* (Springer-Verlag, Berlin, 1988).
- [12] A. Abramovici *et al.*, *Science* **256**, 325 (1992).
- [13] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, *Phys. Rev. D* **63**, 044023 (2001); **66**, 027502 (2002).
- [14] E. Poisson and C. M. Will, *Phys. Rev. D* **52**, 848 (1995).
- [15] K.G. Arun, B.R. Iyer, B.S. Sathyaprakash, and P.A. Sundararajan, *Phys. Rev. D* **71**, 084008 (2005); **72**, 069903(E) (2005).
- [16] L. Blanchet, T. Damour, G. Esposito-Farèse, and B.R. Iyer, *Phys. Rev. Lett.* **93**, 091101 (2004).
- [17] L. E. Kidder, C. M. Will, and A. G. Wiseman, *Phys. Rev. D* **47**, R4183 (1993).
- [18] See, e.g., A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions* (Chapman and Hall/CRC, Boca Raton, FL, 2000).
- [19] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation* (SIAM, Philadelphia, PA, 2005).
- [20] G. Golub and C. van Loan, *Matrix Computations* (Johns Hopkins University Press, London, 1996), 3rd ed.
- [21] E. Berti, A. Buonanno, and C. M. Will, *Classical Quantum Gravity* **22**, S943 (2005).
- [22] C. Cutler (private communication).
- [23] W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, *Numerical Recipes in C* (Cambridge University Press, Cambridge, England, 1988).
- [24] N. Metropolis *et al.*, *J. Chem. Phys.* **21**, 1087 (1953); W.K. Hastings, *Biometrika* **57**, 97 (1970).
- [25] O.E. Barndorff-Nielsen and D.R. Cox, *Inference and Asymptotics* (Chapman and Hall, London, 1994).
- [26] M. Zanolin, E. Naftali, and N. C. Makris (unpublished).
- [27] E. Naftali and N. C. Makris, *J. Acoust. Soc. Am.* **110**, 1917 (2001); A. Thode, M. Zanolin, E. Naftali, I. Ingram, P. Ratilal, and N. C. Makris, *ibid.* **112**, 1890 (2002).
- [28] L. Parker and S. M. Christensen, *MathTensor: A System for Doing Tensor Analysis by Computer* (Addison-Wesley, Reading, MA, 1994).
- [29] J. Zinn-Justin, *Path Integrals in Quantum Mechanics* (Oxford University, New York, 2005).
- [30] E. Bodewig, *Matrix Calculus* (North-Holland, Amsterdam, 1959).