

Probing dynamical gravity with the polarization of continuous gravitational waves

Maximiliano Isi,^{1,*} Matthew Pitkin,² and Alan J. Weinstein¹

¹*LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA*

²*University of Glasgow, Glasgow G12 8QQ, Scotland, United Kingdom*

(Received 23 March 2017; published 15 August 2017)

The direct detection of gravitational waves provides the opportunity to measure fundamental aspects of gravity which have never been directly probed before, including the polarization of gravitational waves. In the context of searches for continuous waves from known pulsars, we present novel methods to detect signals of any polarization content, measure the modes present and place upper limits on the amplitude of nontensorial components. This will allow us to obtain new model-independent, dynamical constraints on deviations from general relativity. We test this framework on multiple potential sources using simulated data from three advanced-era detectors at design sensitivity. We find that signals of any polarization will become detectable and distinguishable for characteristic strains $h \gtrsim 3 \times 10^{-27} \sqrt{1 \text{ yr}/T}$, for an observation time T . We also find that our ability to detect nontensorial components depends only on the power present in those modes, irrespective of the strength of the tensorial strain.

DOI: [10.1103/PhysRevD.96.042001](https://doi.org/10.1103/PhysRevD.96.042001)

I. INTRODUCTION

The recent detection of gravitational waves (GWs) by the advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) heralds the beginning of the long-awaited era of GW astronomy [1,2]. One of the main goals of this field is to use GWs as a probe of fundamental physics in the highly dynamical and strong-field regimes of gravity, as predicted by the general theory of relativity (GR). The first few GW detections have already been used to place some of the most stringent constraints on deviations from GR in this domain, which is inaccessible to laboratory, Solar System or cosmological tests of gravity.

However, it has not been possible to use LIGO signals to learn about the polarization content of GWs [3], a measurement highly relevant when comparing GR to many of its alternatives [4,5]. The reason for this is that the relative orientation of the two LIGO detectors makes it nearly impossible to unequivocally characterize the polarizations of transient GW signals like the compact-binary coalescences (CBCs) observed so far. In fact, at least five noncoaligned quadrupolar detectors would be needed to break the degeneracies of all five nondegenerate polarizations allowed by generic metric theories of gravity [6,7].

Existing observations that are usually taken to constrain the amount of allowed non-GR polarizations can do so only in an indirect and strongly model-dependent manner. For example, measurements of the orbital decay of binary systems are sensitive to the total radiated GW power, but do not probe the waves directly (see e.g. [8,9], or [10,11]

for reviews). In the context of specific alternative theories (e.g. scalar-tensor) such observations can indeed constrain the power in extra polarizations; however, they provide no direct, model-independent information on the actual polarization content of the gravitational radiation. Thus, there may be multiple theories, with different polarization content, that still predict the correct observed GW emitted power. Because other traditional tests of GR (like Solar System tests) have no bearing on GWs, there currently exist no direct measurements of GW polarizations.

Unlike CBC transients, continuous gravitational waves (CWs) are, by definition, long-lasting narrow-band signals. Although they have not yet been observed [12–17], CWs are expected to be emitted by stable systems, like spinning neutron stars with an asymmetric moment of inertia [18]. If detected, such signals would allow for tests of gravity complementary to those achievable with transients, including the study of GW polarizations [19].

In [19], we showed that it is possible to search for CWs in a polarization-agnostic way and to disentangle the polarization content if a signal is present. However, the data analysis methods proposed were based on a frequentist approach to statistics and suffered from the associated limitations. In this paper, we reframe the ideas of [19] in a more sophisticated Bayesian framework that allows us to achieve the following novel goals:

- (1) *Model-independent detection*: determine whether a set of GW detector data, prepared for any given known pulsar and from one or multiple detectors, provides evidence for the presence of an astrophysical signal of any polarization content.
- (2) *Model selection*: in the presence of a signal, determine whether the data favor GR or a generic non-GR

*misi@ligo.caltech.edu

model, as well as comparing specific alternative theories among themselves and to GR; combine data for multiple sources into a single statement about the validity of GR.

- (3) *Inference*: if the data favor the presence of a GR signal, place constraints on specific alternative theories using the tools of Bayesian parameter estimation.

Furthermore, while [19] treated only the case of a single detector, we are now able to consider the generic case of a network of detectors.

We present Bayesian methods to achieve the three goals above in the context of searches targeted to known pulsars and present sensitivity estimates for the advanced detector era, including the first generic estimates of sensitivity to nontensorial CW polarizations ever published. In Sec. II, we review the basics of beyond-Einstein polarizations and the targeted pulsar CW search. In Sec. III, we phrase our problem in the language of model selection and explain the construction of hypotheses that will allow us to distinguish GR from non-GR signals. In Sec. IV, we specify the details of our analysis, and we explain our results in Sec. V. Finally, we summarize our findings and explain caveats in Sec. VI.

II. BACKGROUND

A. Polarizations

GWs can be decomposed into different polarizations, which arise from the linearly independent components of the three-dimensional, rank-two tensor representing the spatial metric perturbation [7]. A generic metric theory of gravity may thus allow any combination of up to six independent modes: plus (+), cross (\times), vector x (x), vector y (y), breathing (b) and longitudinal (l). The effect of each of these modes is represented in Fig. 1. The rotational properties of the fields underlying any given theory determines which polarizations the theory supports: + and \times correspond to tensor fields (helicity ± 2), x and y to vector fields (helicity ± 1), and b and l to scalar fields (helicity 0).

The components of the tensor and vector pairs are not separable, in the sense that a signal model that includes one element of the group must also include the other (e.g. it is not possible to have a model that allows plus + but not \times), because the distinction between + and \times , or x and y, is contingent on the frame of reference (e.g. relative orientation of source and detector).

Einstein's theory only allows the existence of the + and \times polarizations. On the other hand, scalar-tensor and massive-graviton theories may also predict the presence of some b and/or l component associated to the theory's extra scalar field [5]. On top of tensor and scalar modes, bimetric theories, like Rosen or Lightman-Lee theories, also predict vector modes [20]. Furthermore, less conventional theories

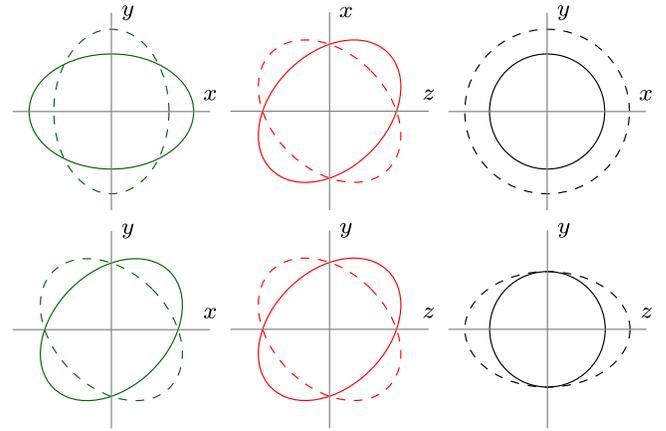


FIG. 1. Effect of different GW polarizations on a ring of free-falling test particles. Plus (+) and cross (\times) tensor modes (green); vector-x (x) and vector-y (y) modes (red); breathing (b) and longitudinal (l) scalar modes (black). In all of these diagrams, the wave propagates in the z direction. This decomposition into polarizations was first proposed for generic metric theories in [7].

might, in principle, predict the existence of vector or scalar modes *only*, while still possibly being in agreement with all other non-GW tests of GR (see [21] for an example). Although all these different theoretical frameworks serve as motivation for our study, our approach to the measurement of GW polarizations is phenomenological and, thus, theory-agnostic (Sec. III). It is important to underscore that the detection of a GW signal with a non-GR polarization, no matter how small, is sufficient to falsify GR (note the converse is not true, however).

Because different polarizations have geometrically distinct effects, GW detectors will react differently to each mode. This is manifested in the detector response function F_p for each polarization p , which encodes the effect of a linearly p -polarized GW with unit amplitude, $h_p = 1$. Ground-based GW detectors, like LIGO and Virgo are quadrupolar antennas that perform low-noise measurements of the strain associated with the differential motion of two orthogonal arms. Their detector response function can thus be written as [19,22–24]:

$$F_+ = \frac{1}{2} [(\mathbf{w}_x \cdot \mathbf{d}_x)^2 - (\mathbf{w}_x \cdot \mathbf{d}_y)^2 - (\mathbf{w}_y \cdot \mathbf{d}_x)^2 + (\mathbf{w}_y \cdot \mathbf{d}_y)^2], \quad (1)$$

$$F_\times = (\mathbf{w}_x \cdot \mathbf{d}_x)(\mathbf{w}_y \cdot \mathbf{d}_x) - (\mathbf{w}_x \cdot \mathbf{d}_y)(\mathbf{w}_y \cdot \mathbf{d}_y), \quad (2)$$

$$F_x = (\mathbf{w}_x \cdot \mathbf{d}_x)(\mathbf{w}_z \cdot \mathbf{d}_x) - (\mathbf{w}_x \cdot \mathbf{d}_y)(\mathbf{w}_z \cdot \mathbf{d}_y), \quad (3)$$

$$F_y = (\mathbf{w}_y \cdot \mathbf{d}_x)(\mathbf{w}_z \cdot \mathbf{d}_x) - (\mathbf{w}_y \cdot \mathbf{d}_y)(\mathbf{w}_z \cdot \mathbf{d}_y), \quad (4)$$

$$F_b = \frac{1}{2} [(\mathbf{w}_x \cdot \mathbf{d}_x)^2 - (\mathbf{w}_x \cdot \mathbf{d}_y)^2 + (\mathbf{w}_y \cdot \mathbf{d}_x)^2 - (\mathbf{w}_y \cdot \mathbf{d}_y)^2], \quad (5)$$

$$F_l = \frac{1}{2} [(\mathbf{w}_z \cdot \mathbf{d}_x)^2 - (\mathbf{w}_z \cdot \mathbf{d}_y)^2]. \quad (6)$$

Here, the spatial vectors \mathbf{d}_x , \mathbf{d}_y have unit norm and point along the detector arms such that $\mathbf{d}_z = \mathbf{d}_x \times \mathbf{d}_y$ is the local zenith; the direction of propagation of the wave from a source at known sky location (specified by right ascension α , and declination δ) is given by \mathbf{w}_z , and \mathbf{w}_x , \mathbf{w}_y are such that $\mathbf{w}_z = \mathbf{w}_x \times \mathbf{w}_y$. We choose \mathbf{w}_x to lie along the intersection of the equatorial plane of the source with the plane of the sky, and let the angle between \mathbf{w}_y and the celestial north be ψ , the polarization angle.

Because of their symmetries, the breathing and longitudinal modes are fully degenerate to networks of quadrupolar antennas (see e.g. Sec. VI of [20]). This means that no model-independent measurement with such a network can possibly distinguish between the two, so it is enough for us to consider just one of them explicitly. We will refer to the scalar modes jointly by the subscript ‘‘s’’.

The response of gravitational detectors to signals of a given polarization and direction of propagation can be represented, as in Fig. 2, by a spherical polar plot in which the radial coordinate corresponds to the sensitivity given by Eqs. (1)–(6). In the frame of a given detector, this can be written as [see e.g. Eqs. (13.98) in [24] with $\psi = -\pi/2$, to account for the different wave-frame definition]:

$$F_+(\vartheta, \varphi; \psi = 0) = -\frac{1}{2} (1 + \cos^2 \vartheta) \cos 2\varphi, \quad (7)$$

$$F_\times (\vartheta, \varphi; \psi = 0) = -\cos \vartheta \sin 2\varphi, \quad (8)$$

$$F_x (\vartheta, \varphi; \psi = 0) = -\sin \vartheta \sin 2\varphi, \quad (9)$$

$$F_y (\vartheta, \varphi; \psi = 0) = \sin \vartheta \cos \vartheta \cos 2\varphi, \quad (10)$$

$$F_{b/l} (\vartheta, \varphi; \psi = 0) = \mp \frac{1}{2} \sin^2 \vartheta \cos 2\varphi, \quad (11)$$

where ϑ and φ are the polar and azimuthal coordinates of the source with respect to the antenna at any given time (with detector arms along the x and y -axes), and we have fixed the wave frame so that $\psi = 0$. The representation of Fig. 2 makes it clear that quadrupolar detectors will generally be more sensitive to some polarizations than others, although this will vary with the sky location of the source. For example, for all but a few sky locations, quadrupolar antennas will respond significantly less to a breathing signal than a plus or cross signal.

For a given detector, polarization angle and sky location, the antenna patterns of Eqs. (1)–(6) become simple, distinct

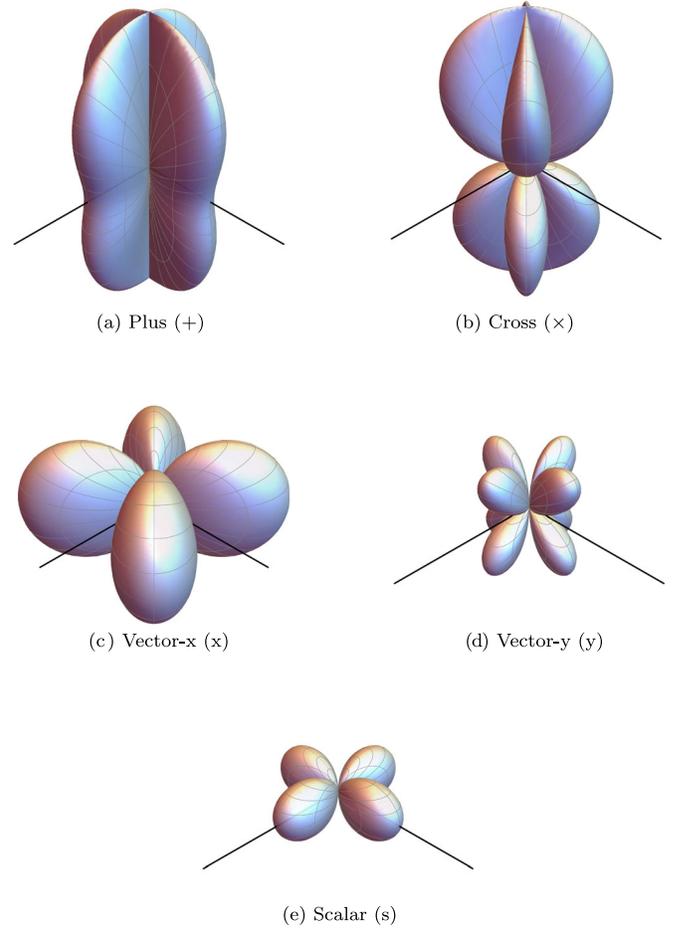


FIG. 2. Angular response of a quadrupolar detector to each GW polarization. The radial distance represents the response of a single quadrupolar antenna to a unit-amplitude gravitational signal of a tensor (top), vector (middle), or scalar (bottom) polarization, i.e. $|F_p|$ for each polarization p as given by Eqs. (7)–(11). The polar and azimuthal coordinates correspond to the source location with respect to the detector, which is to be imagined as placed with its vertex at the center of each plot and arms along the x and y -axes. The response is plotted to scale, such that the black lines representing the detector arms have unit length in all plots. The response to breathing and longitudinal modes is identical, so we only display it once and label it ‘‘scalar’’.

functions of time determined by the rotation of the Earth. This can be pictured by noting that, as the Earth spins on its axis, the angular location of the source with respect to detector will change, tracing an arc on the surfaces of Fig. 2 with varying radial distance. As we explain in Sec. III A 1, the F_p ’s of polarizations with different rotational properties can be distinguished even in the absence of information on the source orientation; for the minority of cases in which such information exists, it can be taken into account to better distinguish among specific signal models (see Sec. IV).

Because their characteristic period (a sidereal day) is much longer than the CBC timescale (order of minutes or

less), the F_p 's are treated as constants in transient searches; however, this simplification is not valid for CW searches, since their coherent-integration time can be of the order of months to years. As we have pointed out before, this can be used to distinguish the polarization content of a signal [19]. Assuming wave frequency and speed are the same for all modes, the only differences between CWs of different polarizations arise from the sidereal-period amplitude modulations caused by each antenna pattern.

B. Continuous waves

1. Signal

A CW is an almost-monochromatic gravitational perturbation with constant intrinsic amplitude and phase evolution $\phi(t)$. For arbitrary polarization content, such a GW will induce a strain in a quadrupolar detector which can be written as:

$$h(t) = \sum_p F_p(t) h_p(t), \quad (12)$$

where the sum is over the five independent polarizations, $p \in \{+, \times, x, y, s\}$; the F_p 's are those of Eqs. (1)–(5), and thus implicitly depend on the relative location and orientation of source and detector by means of ψ , α and δ ; the h_p term encodes the amplitude and phase of the wave before being projected onto the frame of the detector:

$$h_p(t) = a_p \cos(\phi(t) + \phi_p), \quad (13)$$

where a_p is a time-independent amplitude with a functional dependence on source parameters determined by each particular theory of gravity; $\phi(t)$ the phase evolution, a consequence of the dynamics of the source in that theory; and ϕ_p a phase offset for each polarization. The polarization amplitudes a_p and phases ϕ_p may take arbitrary values depending on the specific theory of gravity and emission mechanism.

In GR, there are several ways in which a neutron star could emit CWs, but the most likely is the presence of a nonaxisymmetry in the star's moment of inertia [25]. For this type of triaxial, nonprecessing source, GR predicts:

$$h_+(t) = h_0 \frac{1}{2} (1 + \cos^2 \iota) \cos \phi(t), \quad (14)$$

$$h_\times(t) = h_0 \cos \iota \sin \phi(t), \quad (15)$$

$$h_x = h_y = h_s = h_l = 0, \quad (16)$$

where ι is the inclination angle between the spin axis of the source and the observing line-of-sight, and h_0 is an overall amplitude given by:

$$h_0 = \frac{16\pi^2 G \epsilon I_{zz} f_{\text{rot}}^2}{c^4 r}, \quad (17)$$

where r is the source distance, f_{rot} its rotation frequency around the principal axis z , \mathbf{I} the moment-of-inertia tensor and $\epsilon \equiv (I_{xx} - I_{yy})/I_{zz}$ the equatorial ellipticity. For the triaxial case, the GW frequency f is twice the rotational value f_{rot} , so that we can write:

$$\phi(t) = 2\phi_{\text{rot}}(t) + \phi_{(\text{GW-EM})}, \quad (18)$$

where ϕ_{rot} is the rotational phase as measured via electromagnetic (EM) observations and $\phi_{(\text{GW-EM})}$ is a potential, constant phase offset between the GW and EM signals that can be absorbed into the definition of the ϕ_p 's in Eq. (13).

Note that other emission mechanisms may result in GW radiation at $f = f_{\text{rot}}$ [26], or even noninteger powers f_{rot} [27–29]. Furthermore, alternative theories of gravity may (and, in general, will) support signals at any harmonic. Although in this paper we only consider the case in which only the second rotational harmonic appears in the GW phase, the analysis can be easily generalized to also include contributions from the fundamental and other multiples of f_{rot} (see Sec. VI).

2. Targeted search

We would like to search a given set of data (from one or more detectors) for CW signals coming from a specific candidate pulsar which has already been observed and timed electromagnetically. Timing solutions are obtained through the pulsar timing package TEMPO2 [30]. We want to achieve this regardless of polarization content, and to reliably distinguish between the different modes present.

If we assume all polarizations share the same phase evolution, then detector response is the only factor distinguishing CW polarizations and, thus, all the relevant information is encoded in the sidereal-day-period amplitude modulation of the signal. This allows us to focus on a narrow frequency band around the expected GW frequency by processing the data following the complex-heterodyne method presented in [31,32]. This procedure is summarized below.

A signal like Eqs. (12) and (13) can be rewritten in the form:

$$h(t) = \Lambda(t) e^{i\phi(t)} + \Lambda^*(t) e^{-i\phi(t)}, \quad (19)$$

$$\Lambda(t) \equiv \frac{1}{2} \sum_{p=1}^5 a_p e^{i\phi_p} F_p(t_k; \psi, \alpha, \delta), \quad (20)$$

with $*$ indicating complex conjugation and $\phi(t)$ given by a Taylor expansion around $f = 2f_{\text{rot}}$:

$$\phi(t) = 2\pi(2f_{\text{rot}}\tau + \dot{f}_{\text{rot}}\tau^2 + \dots), \quad (21)$$

where τ is itself a function of time given by:

$$\tau(t) = t + \Delta_R + \Delta_E + \Delta_S + \Delta_{\text{binary}}. \quad (22)$$

In the above, τ is the time measured by a clock inertial with respect to the pulsar; t is the time as measured at a given detector; Δ_R is the Roemer delay; Δ_E is the Solar-System Einstein delay; Δ_S is the Solar-System Shapiro delay; Δ_{binary} is the delay originating from the motion of the pulsar in its binary (a term that vanishes for isolated sources) [32].

It is important to remember that, the F_p 's are functions of the source orientation and sky location relative to the detector, so we have made this dependence explicit in Eq. (20) by writing $F_p(t_k)$ as $F_p(t_k; \psi, \alpha, \delta)$. Also, recall that these functions have a characteristic period of a sidereal day ($\sim 10^{-5}$ Hz).

Because the phase evolution $\phi(t)$, including all corrections from Eq. (22), is known (with known uncertainties) from electromagnetic observations, we can digitally heterodyne the data by multiplying by $\exp[-i\phi(t)]$ so that the signal therein becomes

$$h'(t) \equiv h(t)e^{-i\phi(t)} = \Lambda(t) + \Lambda^*(t)e^{-i2\phi(t)}, \quad (23)$$

and the frequency modulation of the first term is removed, while that of the second term is doubled. A series of low-pass filters can then be used to remove the quickly varying term, which enables the down-sampling of the data by averaging over minute-long time bins. As a result, $\Lambda(t)$ is the only contribution from the original signal left in our data, and hence we can use Eq. (20) as the template for our search. Note that, although we started with real-valued data, after this process the data are now complex.

From Eq. (23), we see that, in the presence of a signal, the heterodyned and down-sampled noisy detector strain data B_k for the k^{th} time bin (which can also be labeled by the Earth-frame GPS time-of-arrival at the detector, t_k) are expected to be of the form:

$$B_{\text{expected}}(t_k) = \Lambda(t_k) + n(t_k), \quad (24)$$

where $n(t_k)$ is the heterodyned, filtered and downsampled noise in bin k , which carries no information about the GW signal. Note then that $B_k(t_k) - \Lambda(t_k)$ should be expected to have the statistical properties of noise, a fact that will be used below in defining likelihoods.

III. METHOD

A. Model selection

We use the tools of Bayesian model selection (also known as *second-level inference*) to determine whether the data contain a signal and, if so, whether that signal agrees

with the GR prediction or not. Our procedure is hierarchical and consists of the following stages:

- (1) *Detection*: select between signal and noise models;
- (2) *Test of GR*: if a signal is present, select between GR and non-GR models;
- (3) *Upper limits*: if GR is favored, place upper limits on nontensorial strain amplitudes, in the context of specific alternative polarization models.

This subsection covers only the first two items in this list, since the placement of upper limits belongs in the section on parameter estimation. We treat the case of a single data set in III A 1 and III A 2, and we show how to combine results from multiple analyses in III A 3; we offer some considerations about how to approach the problem of non-Gaussian noise in III A 4.

1. Hypotheses

For any given pulsar, we would first like to use *reduced* (i.e. heterodyned, filtered and downsampled) GW data to decide between the following two logically disjoint hypotheses:

- (1) *noise* (\mathcal{H}_N): no signal, the data are drawn from a Gaussian distribution of zero mean and some (possibly slowly varying) standard deviation;
- (2) *signal* (\mathcal{H}_S): the data contain noise drawn from a Gaussian distribution and a signal with the assumed phase evolution and *any* polarization content.

In order to perform model selection, we need to translate these hypotheses into the corresponding Bayesian models; this means setting a likelihood function derived from the expected noise properties and picking a multidimensional prior distribution over all parameters. It is important to underscore that a Bayesian model is *defined* by the choice of these two probability distributions.

For \mathcal{H}_N , the construction of the likelihood is straightforward. First, let σ be the standard deviation of the detector noise at or near the expected GW frequency; then, for each complex-valued data point B_k , Gaussianity implies:

$$p(B_k|\sigma, \mathcal{H}_N) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{|B_k|^2}{2\sigma^2}\right). \quad (25)$$

Here, and throughout this document, a lower-case p is used for probability densities, while an uppercase P is used for discrete probabilities.

If the data are split into N_S segments of lengths s_j ($j = 1, \dots, N_S$) over which the standard deviation σ_j is assumed to remain constant, we can analytically marginalize over this parameter to obtain a likelihood for the entire data set \mathbf{B} in the form of a Student's t -distribution [32,33]:

$$P(\mathbf{B}|\mathcal{H}_N) = \prod_{j=1}^{N_S} A_j \left(\sum_{k=\kappa_j}^{K_j} |B_k|^2 \right)^{-s_j}, \quad (26)$$

with $A_j = (s_j - 1)!/2\pi^{s_j}$, $\kappa_j = 1 + \sum_{n=1}^j s_{n-1}$, $K_j = \kappa_j + s_j - 1$ and $s_0 = 0$. Data streams from N_D detectors can be analyzed coherently by generalizing this to:

$$P(\mathbf{B}|\mathcal{H}_N) = \prod_{i=1}^{N_D} \prod_{j=1}^{N_{S_i}} A_{i,j} \left(\sum_{k=\kappa_{i,j}}^{K_{i,j}} |B_{i,k}|^2 \right)^{-s_{i,j}}, \quad (27)$$

where i indexes detectors, $B_{i,k} \equiv B_i(t_k)$ is the datum corresponding to the i^{th} detector at time t_k , and $A_{i,j}$, $\kappa_{i,j}$ and $K_{i,j}$ are defined analogously to s_j , κ_j above. The splitting of the data into segments of constant standard deviation may be achieved with a strategy similar to the Bayesian-blocks algorithm of [34], and explained in detail in [33].

Note that the likelihood $p(\mathbf{B}|\vec{\theta}, \mathcal{H})$ of some hypothesis \mathcal{H} , is the probability of observing the data \mathbf{B} assuming \mathcal{H} is true and given a specific choice of free parameters $\vec{\theta}$ from the model's parameter space Θ . However, in the case of the noise (“null”) hypothesis, as defined by the Student's t likelihood above, there are no free parameters. Consequently, $\Theta = \emptyset$ and $p(\mathbf{B}|\vec{\theta}, \mathcal{H}_N) = P(\mathbf{B}|\mathcal{H}_N)$.

The case of \mathcal{H}_S requires more careful attention. One could be tempted to use Eq. (24) to define a likelihood like Eq. (27) with the substitution $|B_k| \rightarrow |B_k - \Lambda_k|$, for $\Lambda_k \equiv \Lambda(t_k)$ including all polarizations like in Eq. (20); the priors would reflect uncertainties in measured source parameters and extend over reasonable ranges for a_p and ϕ_p . However, for most realistic prior choices, that would correspond to a hypothesis that assigns most of the prior probability to regions of parameter space for which $a_p \neq 0$ for all p , thus downweighting more conservative models (including GR) that we would like to prioritize. This is simply because the subspace in parameter space corresponding to any of these smaller subhypotheses (which, for example, fix one of the a_p 's to be zero) has infinitely less volume (i.e. it offers infinitesimally less support) than its complement; hence any practical choice of prior probability density will also assign this subspace infinitely less weight, and so the prior for the corresponding subhypothesis will be vanishingly small.

Formally, the inadequacy of the naive construction of \mathcal{H}_S as proposed in the previous paragraph is related to the logical independence of nested hypotheses. We refer to this important point multiple times in the following sections; in particular, we discuss it in the context of odds computations in the text surrounding Eq. (38). We refer readers not familiar with this line of reasoning to a similar discussion in [35], or, more generally, to Chap. 4 in [36] or Chap. 28 in [37].

Instead, we will construct \mathcal{H}_S from two logically disjoint component hypotheses:

- (1) *GR signal* (\mathcal{H}_{GR} or \mathcal{H}_t): the data contain Gaussian noise and a tensorial signal with the assumed $\phi(t)$;

- (2) *Non-GR signal* (\mathcal{H}_{nGR}): the data contain Gaussian noise and a signal with non-GR polarization content, but with the assumed $\phi(t)$.

The tensorial hypothesis is embodied most generally by a signal model such that

$$\Lambda_t(t) = \frac{1}{2} [a_+ e^{i\phi_+} F_+(t; \psi = 0) + a_\times e^{i\phi_\times} F_\times(t; \psi = 0)], \quad (28)$$

where a_+ , a_\times , ϕ_+ and ϕ_\times are free parameters, and we pick a specific polarization frame by setting $\psi = 0$ (we are allowed to do this because of a degeneracy between ψ and a_+ , a_\times explained in Appendix A). An alternative parametrization can be derived from the triaxial emission model of Eqs. (14)–(16), namely

$$\Lambda_{\text{GR}}(t) = \frac{1}{2} h_0 e^{i\phi_0} \left[\frac{1}{2} (1 + \cos^2 \iota) F_+(t; \psi) - i \cos \iota F_\times(t; \psi) \right], \quad (29)$$

where the free parameters are now h_0 , ϕ_0 , ι and ψ [in the notation of Eqs. (13) and (20), $\phi_+ = \phi_0$ and $\phi_\times = \phi_0 - \pi/2$]. This is the parametrization used in most traditional GR-only searches (see e.g. [17,33]).

The templates of Eqs. (28) and (29) span the same signal space; therefore, if we pick parameter priors properly related by their Jacobian, the respective hypotheses (\mathcal{H}_t and \mathcal{H}_{GR}) will be logically equivalent (i.e. $\mathcal{H}_t \equiv \mathcal{H}_{\text{GR}}$). However, we will sometimes want to restrict ψ or ι in Eq. (29) to incorporate measurements of the source orientation (see Table 3 in [38]), and compare those results to the unconstrained model of Eq. (28). In such cases, \mathcal{H}_t and \mathcal{H}_{GR} are no longer equivalent: the former corresponds to a free-tensor signal, while the latter now corresponds to a GR triaxial signal for some given source orientation [i.e. a signal with the functional dependence on ι and ψ of Eq. (29)]. Because of lack of any orientation information, this is a distinction without a difference for most pulsars. (See Appendix A for more details.)

The non-GR hypothesis, \mathcal{H}_{nGR} , can itself be seen as a composite hypothesis encompassing all the signal models that depart from GR in some way, i.e. models that include polarizations other than $+$ and \times . We denote such subhypotheses with a subscript listing the polarizations included in the signal. For example, “st” (meaning “scalar plus tensor”) corresponds to a model with unrestricted scalar and tensor contributions:

$$\Lambda_{\text{st}}(t) = \frac{1}{2} [a_+ e^{i\phi_+} F_+(t; \psi = 0) + a_\times e^{i\phi_\times} F_\times(t; \psi = 0) + a_s e^{i\phi_s} F_s(t; \psi = 0)]. \quad (30)$$

With this notation extended to the names of the relevant hypotheses, we may then write \mathcal{H}_{nGR} as the logical union (“or” junction, \vee)

$$\begin{aligned}\mathcal{H}_{\text{nGR}} &\equiv \mathcal{H}_{\text{s}} \vee \mathcal{H}_{\text{v}} \vee \mathcal{H}_{\text{st}} \vee \mathcal{H}_{\text{sv}} \vee \mathcal{H}_{\text{tv}} \vee \mathcal{H}_{\text{stv}} \\ &= \bigvee_{m \in \tilde{M}} \mathcal{H}_m,\end{aligned}\quad (31)$$

where, for convenience, we have defined the non-GR subscript set \tilde{M} :

$$\tilde{M} \equiv \{\text{s}, \text{v}, \text{st}, \text{sv}, \text{tv}, \text{stv}\}.\quad (32)$$

Just as before, we may equivalently use the triaxial parametrization, Eq. (29), for the tensor modes in the non-GR hypotheses by instead defining \tilde{M} as

$$\tilde{M} = \{\text{s}, \text{v}, \text{sv}, \text{GR} + \text{s}, \text{GR} + \text{v}, \text{GR} + \text{sv}\},\quad (33)$$

where, for example, $\text{GR} + \text{s}$ denotes a signal template like

$$\begin{aligned}\Lambda_{\text{GR}+\text{s}}(t) &= \frac{h_0}{2} e^{i\phi_0} \left[\frac{1}{2} (1 + \cos^2 t) F_+(t; \psi) - i \cos t F_\times(t; \psi) \right] \\ &\quad + \frac{1}{2} a_b e^{i\phi_b} F_b(t; \psi),\end{aligned}\quad (34)$$

and similarly for $\text{GR} + \text{v}$ and $\text{GR} + \text{sv}$, with the added vector modes. Again, the two definitions of \tilde{M} , Eqs. (32) and (33), are equivalent unless orientation information is incorporated in the way explained above.

By the same token, the signal hypothesis can be built from the logical union of \mathcal{H}_{GR} or \mathcal{H}_{t} , and \mathcal{H}_{nGR} :

$$\mathcal{H}_{\text{S}} \equiv \mathcal{H}_{\text{GR/t}} \vee \mathcal{H}_{\text{nGR}} = \bigvee_{m \in M} \mathcal{H}_m,\quad (35)$$

with M defined similarly to \tilde{M} , but also including the tensor-only hypothesis, \mathcal{H}_{GR} or \mathcal{H}_{t} :

$$M \equiv \tilde{M} \cup \{\text{GR/t}\}.\quad (36)$$

The validity of Eqs. (31) and (35) is contingent on the mutual logical independence of all the \mathcal{H}_m ’s. This requirement is satisfied by construction, since each of the \mathcal{H}_m ’s is defined to exclude regions of parameter space that would correspond to other hypotheses nested within it (e.g. $\mathcal{H}_{\text{GR}+\text{s}}$ is defined over all values of the scalar amplitude except $a_{\text{s}} = 0$, to avoid including \mathcal{H}_{GR}). In practice, however, it is not necessary to explicitly exclude these infinitesimal regions of parameter space, as will be explained in the following section.

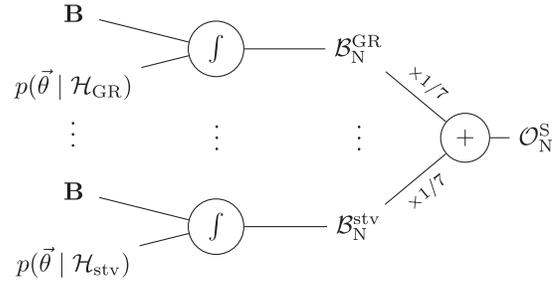


FIG. 3. Computation of $\mathcal{O}_{\text{N}}^{\text{S}}$. First, the Bayes factor \mathcal{B}_{N}^m is obtained from the data \mathbf{B} and corresponding priors $p(\theta|\mathcal{H}_m)$ for each model $m \in M$, by evaluating the integral of Eq. (40) using a nested sampling algorithm that samples over $\vec{\theta}$ (step indicated by integral sign); these values are then added and multiplied by $P(\mathcal{H}_m)/P(\mathcal{H}_{\text{N}})$ to obtain $\mathcal{O}_{\text{N}}^{\text{S}}$, as in Eq. (43). (Note that here we have set $P(\mathcal{H}_m)/P(\mathcal{H}_{\text{N}}) = 1/7$, as explained Sec. IV.) The computation of $\mathcal{O}_{\text{GR}}^{\text{nGR}}$ is analogous.

2. Odds

We can construct a Bayesian model for \mathcal{H}_{S} starting from its components: for each subhypothesis \mathcal{H}_m for $m \in M$, we use a likelihood function like Eq. (27) with the substitution $|B_{i,k}| \rightarrow |B_{i,k} - \Lambda_{m,i,k}|$, i.e.

$$P(\mathbf{B}|\vec{\theta}, \mathcal{H}_m) = \prod_{i=1}^{N_{\text{D}}} \prod_{j=1}^{N_{\text{S}_i}} A_{i,j} \left(\sum_{k=\kappa_{i,j}}^{K_{i,j}} |B_{i,k} - \Lambda_{m,i,k}|^2 \right)^{-s_{i,j}}\quad (37)$$

(where $\Lambda_{m,i,k}$ is the template corresponding to model m , for detector i and time-bin k), and suitable priors on the model parameters $\vec{\theta}_m \in \Theta_m$; then, we combine the posteriors with priors on the models themselves to obtain the posterior for \mathcal{H}_{S} . This last step allows us to incorporate our *a priori* beliefs about the validity of each of the components. This procedure is represented schematically in Fig. 3 and fleshed out below.

The choice of model priors can be made clearer by considering the posterior probability for the signal model. Given some set of detector data \mathbf{B} and underlying assumptions I (suppressed from the following expressions), the posterior probability for \mathcal{H}_{S} is

$$P(\mathcal{H}_{\text{S}}|\mathbf{B}) = \sum_{m \in M} P(\mathcal{H}_m|\mathbf{B})\quad (38)$$

by Eq. (35) and because the components are all logically independent [i.e. $\mathcal{H}_{m_1} \wedge \mathcal{H}_{m_2} = \text{False}$, hence $P(\mathcal{H}_{m_1} \wedge \mathcal{H}_{m_2}|\mathbf{B}) = 0$ for any $m_1, m_2 \in M$ such that $m_1 \neq m_2$]. Note that this is true even for hypotheses that may contain each other as special cases. For instance, even though the GR template can be obtained from $\text{GR} + \text{s}$ by setting the scalar amplitude to $a_{\text{s}} = 0$, the points in the $\text{GR} + \text{s}$ parameter space satisfying this condition define an infinitesimally-thin slice in parameter space that offers

no support to the prior distribution and is thus assigned no weight (see similar discussion in [35]).

We can expand each term on the RHS of Eq. (38) using Bayes' theorem:

$$P(\mathcal{H}_S|\mathbf{B}) = \sum_{m \in M} P(\mathcal{H}_m)P(\mathbf{B}|\mathcal{H}_m)/P(\mathbf{B}). \quad (39)$$

Each of the terms on the RHS is made up of three factors: a marginalized likelihood $P(\mathbf{B}|\mathcal{H}_m)$, a prior on the model $P(\mathcal{H}_m)$, and a normalization constant $P(\mathbf{B})$.

The marginalized likelihood (also known as *evidence*) is computed from the data:

$$P(\mathbf{B}|\mathcal{H}_m) = \int_{\Theta_m} p(\mathbf{B}|\vec{\theta}_m, \mathcal{H}_m)p(\vec{\theta}_m|\mathcal{H}_m)d\vec{\theta}_m, \quad (40)$$

where $p(\mathbf{B}|\vec{\theta}_m, \mathcal{H}_m)$ is itself the likelihood. The evaluation of the multidimensional integral of Eq. (40) is the most computationally intensive part of our analysis (see Sec. IV for details).

We are free to choose the model priors (discussed in Sec. IV), as long as we satisfy the constraint:

$$P(\mathcal{H}_N) + \sum_{m \in M} P(\mathcal{H}_m) = 1. \quad (41)$$

This is a statement about the exhaustiveness and disjointness of the hypotheses we are considering: we assume that reality will agree with *one and only one* of the hypotheses at hand. (As we will see in Sec. VI, this assumption might not hold; for example, the noise may not be Gaussian.) The particular choice of prior for each model will encode our expectations about the corresponding theory (before seeing the data), and thus allow for some degree of subjectivity.

Note that we cannot directly compute $P(\mathbf{B})$ in a straightforward manner and without assuming that our hypothesis set is indeed exhaustive (which is not the case for non-Gaussian detector noise, see Sec. III A 4). However, the need for this computation can be avoided by looking at relative probabilities, i.e. *odds*. The odds for \mathcal{H}_S vs \mathcal{H}_N is defined as:

$$\mathcal{O}_N^S \equiv \frac{P(\mathcal{H}_S|\mathbf{B})}{P(\mathcal{H}_N|\mathbf{B})}. \quad (42)$$

Using Bayes' theorem again and canceling the $P(\mathbf{B})$ factors, this simplifies to:

$$\mathcal{O}_N^S = \frac{\sum_{m \in M} P(\mathcal{H}_m)P(\mathbf{B}|\mathcal{H}_m)}{P(\mathcal{H}_N)P(\mathbf{B}|\mathcal{H}_N)} = \sum_{m \in M} \frac{P(\mathcal{H}_m)}{P(\mathcal{H}_N)} \mathcal{B}_N^m, \quad (43)$$

where, in the second equality, we have used the definition of the *Bayes factor*:

$$\mathcal{B}_j^i \equiv \frac{P(\mathbf{B}|\mathcal{H}_i)}{P(\mathbf{B}|\mathcal{H}_j)}, \quad (44)$$

for any two hypotheses $\mathcal{H}_i, \mathcal{H}_j$.

The odds in Eq. (43) can be used as a detection statistic to determine whether it is likely that the data contain a signal (of any polarization) or not. Once the presence of a signal has been established, a similar ratio can be constructed to assess agreement with GR:

$$\mathcal{O}_{GR}^{nGR} = \frac{P(\mathcal{H}_{nGR}|\mathbf{B})}{P(\mathcal{H}_{GR}|\mathbf{B})} = \sum_{m \neq GR} \frac{P(\mathcal{H}_m)}{P(\mathcal{H}_{GR})} \mathcal{B}_{GR}^m. \quad (45)$$

This ratio encodes the relative probability that there is a GR violation. Because it is now assumed that there is a signal in the data, $P(\mathcal{H}_N) = 0$ and the model priors must instead satisfy:

$$\sum_{m \in M} P(\mathcal{H}_m) = 1. \quad (46)$$

We can reduce the number of computations needed to obtain \mathcal{O}_N^S and \mathcal{O}_{GR}^{nGR} by using the fact that:

$$\mathcal{B}_j^i = \frac{P(\mathbf{B}|\mathcal{H}_i)}{P(\mathbf{B}|\mathcal{H}_j)} = \frac{P(\mathbf{B}|\mathcal{H}_i) P(\mathbf{B}|\mathcal{H}_N)}{P(\mathbf{B}|\mathcal{H}_N) P(\mathbf{B}|\mathcal{H}_j)} = \frac{\mathcal{B}_N^i}{\mathcal{B}_N^j}. \quad (47)$$

This means that we need to evaluate an integral like Eq. (40) seven times per set of data, to compute \mathcal{B}_N^m for each m in M . Those seven numbers, together with the evidence for \mathcal{H}_N , are enough to compute all the quantities of interest.

Instead of asking about a generic deviation from GR, we may also compare GR to a particular alternative theory. For such purpose, we will usually assign equal prior weight to GR and its alternative to compute:

$$\mathcal{O}_{GR}^j = \frac{P(\mathcal{H}_j)}{P(\mathcal{H}_{GR})} \mathcal{B}_{GR}^j = \mathcal{B}_{GR}^j, \quad (48)$$

where \mathcal{H}_j may be any of the hypotheses in \tilde{M} or an even more specific hypothesis. (The latter case demands an extra execution of the inference code.)

3. Multiple data sets

So far we have assumed that the data \mathbf{B} , corresponding to one or more GW detectors, can be analyzed coherently; however, there are cases in which we would like to combine results from sets of data analyzed incoherently. Examples are data sets corresponding to different sources or observation periods. Our Bayesian framework makes it possible to combine the respective odds in order to make an overall model selection statement (in our case, about the presence of signal or the validity of GR).

For instance, we may analyze data for N_P pulsars and ask about the probability that any of them contain a signal; treating each as an independent observation, the combined probability can be constructed from the odds above. Letting $\mathcal{H}_{S_i}, \mathcal{H}_{N_i}$ respectively denote signal and noise hypotheses for the i^{th} source, while \mathcal{S}_{any} corresponds to a signal being present in *any* of the sources and \mathcal{N}_{all} corresponds to Gaussian noise in data for *all* sources:

$$\begin{aligned}
 {}^{(N_P)}\mathcal{O}_{N_{\text{all}}}^{\mathcal{S}_{\text{any}}} &= \frac{P(\mathcal{S}_{\text{any}}|\mathbf{B})}{P(\mathcal{N}_{\text{all}}|\mathbf{B})} = \frac{1 - P(\mathcal{N}_{\text{all}}|\mathbf{B})}{P(\mathcal{N}_{\text{all}}|\mathbf{B})} \\
 &= \frac{1}{P(\bigwedge_i \mathcal{H}_{N_i}|\mathbf{B})} - 1 = \left[\prod_{i=1}^{N_P} \frac{1}{P(\mathcal{H}_{N_i}|\mathbf{B}_i)} \right] - 1 \\
 &= \left[\prod_{i=1}^{N_P} \frac{P(\mathcal{H}_{S_i}|\mathbf{B}_i) + P(\mathcal{H}_{N_i}|\mathbf{B}_i)}{P(\mathcal{H}_{N_i}|\mathbf{B}_i)} \right] - 1 \\
 &= \left[\prod_{i=1}^{N_P} (\mathcal{O}_{N_i}^{\mathcal{S}_i} + 1) \right] - 1, \tag{49}
 \end{aligned}$$

where we have used the exclusivity and exhaustiveness of the signal and noise hypotheses, i.e.

$$P(\mathcal{S}_{\text{any}}|\mathbf{B}) + P(\mathcal{N}_{\text{all}}|\mathbf{B}) = 1, \tag{50}$$

$$P(\mathcal{H}_{S_i}|\mathbf{B}_i) + P(\mathcal{H}_{N_i}|\mathbf{B}_i) = 1, \tag{51}$$

with i indexing data sets. Note that the data sets for different sources (\mathbf{B}_i 's) are *not* conditionally independent under \mathcal{S}_{any} or \mathcal{N}_{all} . Also, Eq. (49) does not enforce the requirement that, if signals are present in multiple sources, they all correspond to the same model from Eq. (36); such a constraint could be implemented at this stage, but is more easily enforced by examining individual values of $\mathcal{O}_{N_i}^m$ when necessary.

The construction of Eq. (49) implicitly assigns model priors to each of the meta-hypotheses \mathcal{S}_{any} and \mathcal{N}_{all} such that:

$$\frac{P(\mathcal{S}_{\text{any}})}{P(\mathcal{N}_{\text{all}})} = \left[\frac{P(\mathcal{H}_S)}{P(\mathcal{H}_N)} + 1 \right]^{N_P} - 1, \tag{52}$$

where we have assumed the priors for signal vs noise are equal for all sources, i.e. $P(\mathcal{H}_{S_i}) = P(\mathcal{H}_S)$ and $P(\mathcal{H}_{N_i}) = P(\mathcal{H}_N)$ for all i . When making combined statements for multiple sources, we may wish to choose $P(\mathcal{H}_S)/P(\mathcal{H}_N)$ such as to produce any desired value of $P(\mathcal{S}_{\text{any}})/P(\mathcal{N}_{\text{all}})$, say $P(\mathcal{S}_{\text{any}}) = P(\mathcal{N}_{\text{all}})$. Furthermore, one may wish to weight each pulsar differently within \mathcal{S}_{any} by incorporating information about the source distance (or other parameters) into the priors via a parametrization like Eq. (17); this may improve the sensitivity of the ensemble odds to weak signals in the set, as suggested in [39]. However, using such

a parametrization generally implies committing to a specific gravitational theory (or family of theories). We choose not to take such approach in this study.

Besides combining data for multiple pulsars, for a given source, we could also (incoherently) combine the results of analyses using data from different observation periods. Since the astrophysical CWs we are considering should either be present in all N_R observation runs or in none of them, the relevant odds, generalizing Eq. (43), are

$$\begin{aligned}
 {}^{[N_R]}\mathcal{O}_N^S &= \frac{P(\mathcal{H}_S|\mathbf{B})}{P(\mathcal{H}_N|\mathbf{B})} = \sum_{m \in M} \frac{P(\mathcal{H}_m|\mathbf{B})}{P(\mathcal{H}_N|\mathbf{B})} \\
 &= \sum_{m \in M} \frac{P(\mathbf{B}|\mathcal{H}_m)P(\mathcal{H}_m)}{P(\mathbf{B}|\mathcal{H}_N)P(\mathcal{H}_N)} \\
 &= \sum_{m \in M} \frac{P(\mathcal{H}_m)}{P(\mathcal{H}_N)} \prod_{j=1}^{N_R} (\mathcal{B}_N^m)_j, \tag{53}
 \end{aligned}$$

where we have again used $\mathbf{B} = \{\mathbf{B}_j\}_{j=1}^{N_R}$ to refer to the totality of data, with j indexing observation runs. The independence of the \mathbf{B}_j 's, conditional on \mathcal{H}_m and \mathcal{H}_N , is applied on the last line to write the result in terms of the individual Bayes factors for each run, $(\mathcal{B}_N^m)_j$.

Similarly, we can use multiple data sets to make a single statement about deviations from GR. Once we have made N_P detections from different sources, the odds for a GR violation is:

$${}^{(N_P)}\mathcal{O}_{\text{GR}}^{\text{nGR}} = \sum_{m \in M} \frac{P(\mathcal{H}_m)}{P(\mathcal{H}_{\text{GR}})} \prod_{i=1}^{N_P} (\mathcal{B}_{\text{GR}}^m)_i, \tag{54}$$

where, again, i indexes sources; this is a generalization of Eq. (45). (See Sec. III D of [35] for an analogous derivation.)

4. Non-Gaussian noise

Up to this point, like most other CW studies, we have assumed that the detector noise is Gaussian. However, although previous work has indicated that this is generally a very good approximation [17,19], it is not exactly true for actual detector noise (for some frequencies more so than others). Happily, most of the model selection statements expounded so far are valid also in the presence of non-Gaussian instrumental noise, after some light reinterpretation.

If the assumption of Gaussianity does not hold, the hypotheses constructed in Sec. III A 1 are no longer exhaustive: the data may not only be explained by Gaussian noise or a signal (GR or otherwise), but also by non-Gaussian artifacts that are impossible to satisfactorily model. Nevertheless, the computation and interpretation of evidences and odds remain unchanged for all the hypotheses under consideration.

Because “noise” no longer just means “Gaussian noise”, \mathcal{O}_N^S (which compares the signal model vs *Gaussian* noise) has to be treated more carefully for detection purposes. Indeed, instrumental features that are clearly non-Gaussian (e.g. a loud, narrow-band artifact wandering across the frequency of interest) will generally result in a relatively large value of \mathcal{O}_N^S , even if there is no detectable astrophysical signal in the data. This issue affects the standard GR searches as well [17], although perhaps to a lesser degree due to the reduced signal parameter space.

It is possible to mitigate this problem by constructing a hypothesis that captures some key characteristic of instrumental features and helps discriminate those from real astrophysical signals. Perhaps the best way to do this is to take advantage of the fact that an astrophysical CW must manifest itself coherently across detectors, while the same is not true for detector artifacts [40]. We can thus define an *instrumental feature hypothesis* (\mathcal{H}_I) to encompass the cases in which the data are composed of Gaussian noise, or features that look like astrophysical signals but are not coherent across detectors (viz. they do not have a consistent phase evolution and they are best described by different waveform parameters).

Formally, we define \mathcal{H}_I by:

$$\mathcal{H}_I \equiv \bigwedge_{d=1}^{N_D} (\mathcal{H}_{S_d} \vee \mathcal{H}_{N_d}), \quad (55)$$

where the subscript d identifies detectors, and \wedge is the logical “and” junction. This definition does not explicitly encompass instrumental features that are coherent across some subset of the detectors. Also, note that Eq. (55) implicitly contains a term equivalent to the usual noise hypothesis $\mathcal{H}_N = \bigwedge_d \mathcal{H}_{N_d}$. Similarly, it also contains a term corresponding to the presence of signals in all detectors ($\bigwedge_d \mathcal{H}_{S_d}$). Importantly, such an incoherent term is not equivalent to the coherent signal hypothesis \mathcal{H}_S , as given by the multidetector likelihood of Eq. (37):

$$\mathcal{H}_S \neq \bigwedge_{d=1}^{N_D} \mathcal{H}_{S_d}. \quad (56)$$

While the evidence integral of Eq. (40) factorizes into single-detector terms for \mathcal{H}_N (due to the null parameter space), the same is not true for \mathcal{H}_S . Furthermore, because it does not demand detector coherence, the RHS of Eq. (56) is associated with a considerably larger parameter space than the LHS. Thus, in the presence of an astrophysical signal, model selection will favor \mathcal{H}_S due to its smaller Occam’s penalty. The same is true, of course, when comparing \mathcal{H}_S to \mathcal{H}_I as a whole.

From Eq. (55), it is straightforward to write the evidence for \mathcal{H}_I as

$$P(\mathbf{B}|\mathcal{H}_I) = \prod_{d=1}^{N_D} [P(\mathbf{B}_d|\mathcal{H}_{S_d})P(\mathcal{H}_{S_d}|\mathcal{H}_I) + P(\mathbf{B}_d|\mathcal{H}_{N_d})P(\mathcal{H}_{N_d}|\mathcal{H}_I)] \quad (57)$$

and use this to construct the odds comparing against \mathcal{H}_S :

$$\mathcal{O}_I^S = \frac{P(\mathcal{H}_S)}{P(\mathcal{H}_I)} \frac{\mathcal{B}_N^S}{\prod_{d=1}^{N_D} [P(\mathcal{H}_{S_d}|\mathcal{H}_I)(\mathcal{B}_{N_d}^{S_d} - 1) + 1]}. \quad (58)$$

Here we have used Eq. (57), together with the fact that $P(\mathcal{H}_{S_d}|\mathcal{H}_I) + P(\mathcal{H}_{N_d}|\mathcal{H}_I) = 1$ and $P(\mathbf{B}|\mathcal{H}_N) = \prod_d P(\mathbf{B}|\mathcal{H}_{N_d})$, to write \mathcal{O}_I^S as a function of the detector-coherent signal vs noise Bayes factor \mathcal{B}_N^S , the single-detector signal vs noise Bayes factors $\mathcal{B}_{N_d}^{S_d}$, and model priors $P(\mathcal{H}_S)$, $P(\mathcal{H}_I)$ and $P(\mathcal{H}_{S_d}|\mathcal{H}_I)$.

As usual, we are free to choose the model priors to give more or less weight to different hypotheses. For example, we recover the choice of [17] (Appendix A3) by setting $P(\mathcal{H}_{S_d}|\mathcal{H}_I) = 0.5$ for all d and $P(\mathcal{H}_S) = P(\mathcal{H}_I) \times 0.5^{N_D}$ such that:

$$\ln \mathcal{O}_I^S = \ln \mathcal{B}_N^S - \sum_{d=1}^{N_D} \ln (\mathcal{B}_{N_d}^{S_d} + 1). \quad (59)$$

(When comparing to Appendix A3 of [17], however, note that in that work “I” is used to denote both the background information and the “incoherent-signal-or-noise” hypothesis, which can be identified with our \mathcal{H}_I .)

There is reason to believe that $\ln \mathcal{O}_I^S$, with model priors as in Eq. (59), is quite good at picking out instrumental features, even for data from just two instruments [17]. (Note that we would expect the discriminatory power of $\ln \mathcal{O}_I^S$ to grow with the number of detectors available.) However, at the end of the day, we can never be fully confident that \mathcal{H}_I will indeed capture all nonastrophysical disturbances. To address this, we may always treat $\ln \mathcal{O}_N^S$ and $\ln \mathcal{O}_I^S$ as any generic detection statistic and use estimates of the background distribution to establish significance.

B. Parameter estimation

Besides choosing between different models, we can use Bayesian statistics to obtain posterior probability density functions (PDFs) on the parameters of a given template (first-level inference). In the absence of a loud signal, this can be used to obtain credible intervals that yield upper limits on the amplitudes of GR deviations.

For a model \mathcal{H} with N parameters, an N -dimensional posterior PDF covering the parameter space Θ can be obtained from Bayes’ theorem:

$$p(\vec{\theta}|\mathbf{B}, \mathcal{H}) = \frac{p(\mathbf{B}|\vec{\theta}, \mathcal{H})p(\vec{\theta}|\mathcal{H})}{P(\mathbf{B}|\mathcal{H})}, \quad (60)$$

for $\vec{\theta}$ in Θ , and with $p(\vec{\theta}|\mathcal{H})$ the prior over Θ . To obtain a one-dimensional PDF for a single parameter (call it θ_i), the N -dimensional distribution must be marginalized over all nuisance parameters (viz. all parameters except θ_i):

$$\begin{aligned} p(\theta_i|\mathbf{B}, \mathcal{H}) &= \int_{\Theta'} p(\vec{\theta}|\mathbf{B}, \mathcal{H}) d^{N-1}\theta_j \\ &\propto \int_{\Theta'} p(\mathbf{B}|\vec{\theta}, \mathcal{H}) p(\vec{\theta}|\mathcal{H}) d^{N-1}\theta_j, \end{aligned} \quad (61)$$

where $0 < j \leq N$, such that $j \neq i$, and Θ' denotes the parameter space Θ with the i^{th} dimension removed. Note that the equality has been replaced by a relation of proportionality because we have excluded the evidence $P(\mathbf{B}|\mathcal{H})$ from the expression. (Although of great importance for model selection, this quantity is uninteresting for the purposes of parameter estimation and can be treated as a simple normalization constant.) As discussed in Sec. IV, we evaluate Eq. (61) with the same algorithm used to compute the evidence.

Equation (61) can be used to place upper limits on model parameters; in particular, we will use it to place limits on the amplitude of GR deviations. Consider, for instance, the case of a scalar-tensor theory that can be encapsulated by our GR + s model as described in the previous section; the 95%-credible upper limit on the strength of the breathing mode is $h_s^{95\%}$, defined by:

$$0.95 = \int_{\min(h_s)}^{h_s^{95\%}} p(h_s|\mathbf{B}, \mathcal{H}_{\text{GR+s}}) dh_s, \quad (62)$$

where $\min(h_s)$ is the minimum value of h_s allowed by the prior.

Note that there may be reasons to compute posteriors under different priors than when computing Bayes factors. In particular, it is conventional to present upper limits obtained using a uniform prior over some broad range of the amplitude parameters. With a uniform prior, the posterior is trivially related to the likelihood. This approach produces a more conservative upper limit than other choices, e.g. a Jeffreys prior (see Appendix B).

IV. ANALYSIS

We quantify our ability to use Bayesian model selection to detect CW signals and determine their polarization content as described above. To do this, we use one year of simulated data from three advanced interferometric detectors at design sensitivity: LIGO Hanford (H1), LIGO Livingston (L1) and Virgo (V1). Detector noise is simulated by drawing from a Gaussian distribution with zero mean and variance corresponding to the power

spectral density (PSD) of each detector at the GW frequency of the pulsar. (Previous work has shown that these are good assumptions for actual reduced detector data [17,19].)

As mentioned in the previous section, the key step in our analysis is the computation of the evidence integral of Eq. (40) for the hypotheses under consideration (one noise model, plus seven signal submodels; see Sec. III A 1). We carry this out using a version of the Bayesian inference code used for the targeted pulsar search in [33,41], which we modified to handle signals from theories other than GR. This inference code is itself built on the implementation of Skilling's nested-sampling algorithm [42] in the LALInference package [43], part of the LIGO Algorithm Library Suite [44]. This is the same inference software used for CBC analyses, including GW150914 [45].

In computing likelihoods, we take source location, frequency and frequency derivatives as known quantities (relevant uncertainties are negligible for this analysis). Unless otherwise stated, priors uniform in the logarithm are used for amplitude parameters (h_0 or h_p 's), since these are the least informative priors for scaling coefficients (also known as ‘‘Jeffreys priors’’) [46]; we make the somewhat arbitrary choice of restricting the strain amplitudes to the $10^{-28} - 10^{-24}$ range (this is of little consequence for model selection, as explained in Appendix B). Flat priors are placed over all phase offsets (ϕ_0 and all the ϕ_p 's).

All plots for the Crab pulsar (PSR J0534 + 2200) in Sec. V are produced using known values of its orientation parameters, $\cos i$ and ψ , and with the triaxial parametrization of tensor modes; for other pulsars, however, the free-tensor parametrization is used instead. (See Sec. III A 1 and Appendix A.)

We follow common practice by adopting the principle of indifference (see e.g. Chap. 5 of [36]) in assigning equal prior probability to the signal and noise models, i.e. we let

$$P(\mathcal{H}_S) = P(\mathcal{H}_N) = 1/2. \quad (63)$$

We must also decide how to split the prior among the different \mathcal{H}_m 's when computing \mathcal{O}_N^S and $\mathcal{O}_{\text{GR}}^{\text{nGR}}$. In the former case, we choose to distribute the prior weight uniformly among all signal models, so that:

$$P(\mathcal{H}_m) = |M|^{-1}/2 = 1/14, \quad (64)$$

with $|M| = 7$ the cardinality of M [i.e. the number of signal models that go into the construction of \mathcal{H}_S , see Eq. (36)]. In the latter, however, we prioritize GR by setting:

$$P(\mathcal{H}_{\text{GR}}|\mathcal{H}_S) = 1/2, \quad (65)$$

$$P(\mathcal{H}_m|\mathcal{H}_S) = |\tilde{M}|^{-1}/2 = 1/12. \quad (66)$$

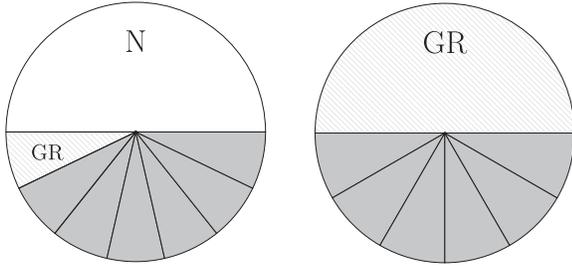


FIG. 4. Model priors. Distribution of prior probability over subhypotheses for the construction of \mathcal{O}_N^S (left) and \mathcal{O}_{GR}^{nGR} (right), according to Eqs. (43) and (45), respectively. For \mathcal{O}_N^S , we assign equal weight to the \mathcal{H}_N (white) and \mathcal{H}_S (gray); as in Eq. (64), we make no *a priori* distinction between non-GR models (solid) and GR (hatched). For \mathcal{O}_{GR}^{nGR} , we set equal prior probability for \mathcal{H}_{GR} and \mathcal{H}_{nGR} , distributing the prior equally among non-GR models, as in Eq. (66).

This distribution is illustrated schematically in Fig. 4. Note that these are not the only justifiable options; for example, we might want to prioritize \mathcal{H}_{GR} when constructing \mathcal{H}_S in order to better handle a noise background that does not conform to our assumption of Gaussianity. (Other strategies to tackle non-Gaussian noise are discussed in Sec. III A 4.) In any case, the code is sufficiently flexible to make different choices for the model priors if desired.

To study our method in the presence of signal, we perform several injections of scalar, vector and tensor polarizations (and combinations thereof) for all the 200 pulsars analyzed in [17]. The simulated signals have a range of signal-to-noise ratios (SNRs), which we proxy below by their *effective strain amplitudes*. We define these in terms of the a_p 's from Eq. (13) by

$$h_t \equiv \sqrt{a_+^2 + a_x^2}, \quad (67)$$

$$h_v \equiv \sqrt{a_x^2 + a_y^2}, \quad (68)$$

$$h_s \equiv a_s, \quad (69)$$

for tensor, vector and scalar signals respectively. Each simulated signal is generated with a random value of the nuisance phase parameters (ϕ_0 or ϕ_p 's). GR injections are always carried out using the triaxial template of Eq. (29), with random orientation parameters (ψ and ι) when those are not known. Location is always taken to be fixed at the known value for each pulsar.

V. RESULTS

A. Model selection

1. Signal vs noise

We first show that \mathcal{O}_N^S , as defined in Eq. (43), can be used to discriminate signals of any polarization from Gaussian

noise, without significant loss of sensitivity to GR signals. The black histogram in Fig. 5 shows the distribution of the natural logarithm of this quantity (henceforth, “log-odds”), obtained from the analysis of an ensemble of noise instantiations corresponding to a single source—in this case, the Crab pulsar. For comparison, the gray (unhatched) histogram in Fig. 5 is the analogous distribution for $\ln \mathcal{B}_N^{GR}$ [note that $\mathcal{B}_N^{GR} = \mathcal{O}_N^{GR}$ if we assign equal priors to the GR and Gaussian noise models, cf. Eq. (43) with $m = GR$]; this is the value computed in regular, GR-only targeted pulsar searches, although with different signal amplitude priors [17]. Note that odds carry an intrinsic probabilistic meaning in terms of gambling probabilities, and a background histogram like this is not required to interpret their value (see e.g. [36]).

For both quantities shown in Fig. 5, a negative value marks a preference for the noise model (\mathcal{H}_N , as defined at the beginning of Sec. III A 1). However, note that a conservative (as determined by the priors) analysis should not be expected to strongly favor \mathcal{H}_N , since the presence of a weak signal below the noise threshold cannot be discarded; this explains why the ranges in the plots of Fig. 5 do not extend to more negative values. Generally speaking, the magnitude of the signal prior volume (viz. the volume of parameter space allowed by the signal model, weighted by the prior function) will determine the mean of background distributions like Fig. 5, which will be more negative the greater the signal volume. This is a manifestation of an implicit Occam’s penalty.

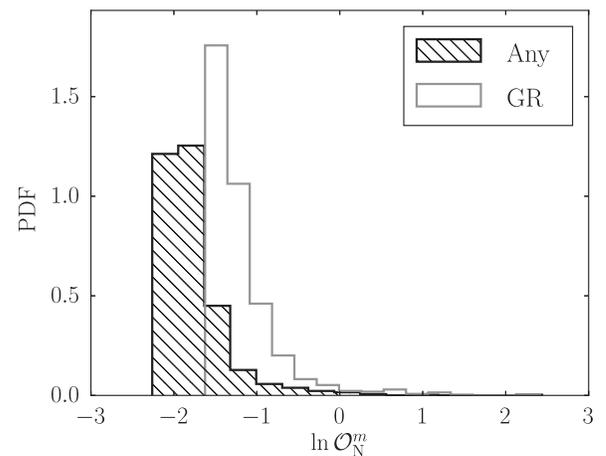


FIG. 5. Signal vs noise log-odds background distributions for any-signal and GR hypotheses. Histograms of $\ln \mathcal{O}_N^S$ (black line, hatched) and $\ln \mathcal{O}_N^{GR}$ (gray line) over an ensemble of 1000 simulated noise instantiations corresponding to the Crab pulsar. For each instantiation, three time series of Gaussian noise were produced using the design noise spectra of H1, L1 and V1, as outlined in Sec. IV; the data are analyzed coherently across detectors. (Note that here $\ln \mathcal{O}_N^{GR} = \ln \mathcal{B}_N^{GR}$, since we assign equal weight to both models.)

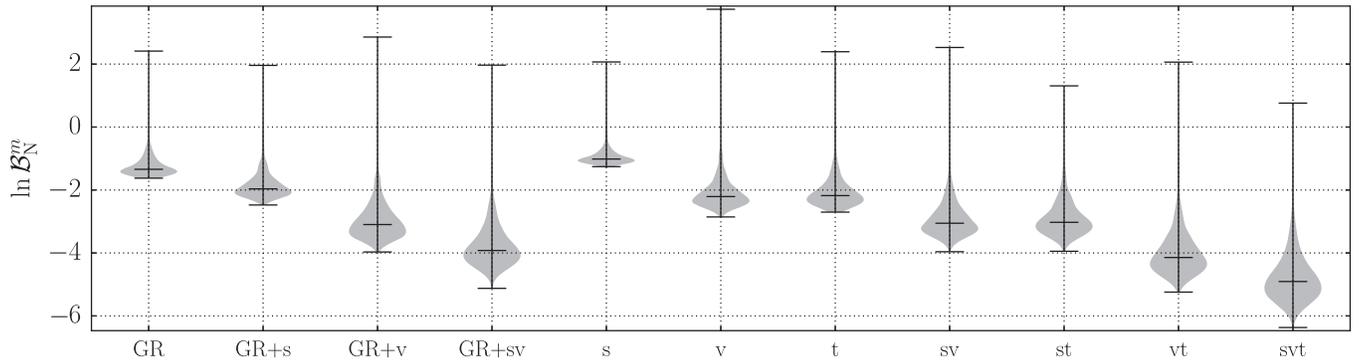


FIG. 6. Signal vs noise log-Bayes background distributions for all subhypotheses. Violin plots representing histograms of the log-Bayes of several models vs noise, computed over an ensemble of 1000 simulated noise instantiations each corresponding to H1, L1 and V1 design data prepared for the Crab pulsar; the data are analyzed coherently across detectors. The labels on the x axis indicate which hypothesis is being compared against noise; the “GR” label indicates tensor modes parametrized by Eq. (29) with fixed ψ and ι . Black lines mark the range and median of each distribution. (The gray histogram in Fig. 5 corresponds to the leftmost distribution here.)

The relationship between the Bayes factors for different signal hypotheses vs noise is illustrated in Fig. 6, which shows violin plots representing the noise-ensemble distributions of $\ln \mathcal{B}_N^m$ for all models discussed in III A 1. The values for $m \in \{s, v, sv, GR, GR + s, GR + t, GR + sv\}$ are combined to produce $\ln \mathcal{O}_N^S$ in Fig. 5. As explained above, the “GR” label indicates that the tensor modes have been parametrized using the triaxial model of Eq. (29), with orientation parameters fixed at the known values for the Crab pulsar; on the other hand, the “t” label corresponds to the free-tensor template of Eq. (28). We include both parametrizations to demonstrate the effect of assuming a triaxial emission mechanism and restricting the orientation of the source (see also Appendix A).

Interestingly, Fig. 6 reveals the relationship between $\ln \mathcal{B}_N^m$ and the number of degrees of freedom (a proxy for the prior volume) of model m : models with more degrees of freedom have a greater prior volume and are correspondingly downweighted, resulting in more negative values of $\ln \mathcal{B}_N^m$; this is a manifestation of the Occam’s penalty automatically applied by the Bayesian analysis (see e.g. Chap. 28 in [37]). We underscore that this feature arises naturally from the computation of the evidence integral, and not from manually downweighting either model *a priori*.

If the data contain a sufficiently loud signal of any polarization, the evidence for \mathcal{H}_S will surpass that for \mathcal{H}_N , and this can be used to establish a detection. Figure 7 shows the response of $\ln \mathcal{O}_N^S$ and $\ln \mathcal{B}_N^{GR}$ to the presence of GR and non-GR signals. In particular, the second panel in Fig. 7 shows results for injected signals of the vector-only model of [21], but the behavior would be the same for scalar-only signals. The general features of these plots confirm our expectations that for weak, subthreshold signals, the analysis should not be able to distinguish between the signal and noise models, yielding a Bayes factor close to unity (more precisely, a value of $\ln \mathcal{O}_N^S$ consistent with the

background distributions of Fig. 5). Note that, in agreement with Fig. 5, the noise baseline for $\ln \mathcal{O}_N^S$ lies below that of $\ln \mathcal{B}_N^{GR}$, due to its greater prior volume.

For stronger (detectable) signals, the basic form of our likelihood functions, Eq. (25), leads us to expect $\ln \mathcal{O}_N^S$ to scale linearly with the square of the signal-to-noise ratio (SNR):

$$\ln \mathcal{B}_N^m \sim (\mathbf{B} \cdot \mathbf{\Lambda}_{MP} - |\mathbf{\Lambda}_{MP}|^2/2)/\sigma^2 \propto (h_{inj}/\sigma)^2, \quad (70)$$

where the variance σ^2 proxies the PSD and we let $\mathbf{\Lambda}_{MP}$ be the time-series vector corresponding to the maximum probability template; for a stationary PSD, this implies $\ln \mathcal{B}_N^m \propto h_{inj}^2$, as observed in Fig. 7. The spread around the trendline is due to the individual features of each noise instantiation and (much less so) to numerical errors in the computation of the evidence, Eq. (40). For details on numerical uncertainty, see Appendix C.

From the left panel of Fig. 7, we conclude that $\ln \mathcal{O}_N^S$ can be as good an indicator of the presence of GR signals as $\ln \mathcal{B}_N^{GR}$ itself; this implies that we may include non-GR polarizations in our search without significantly sacrificing sensitivity to GR signals. However, the power of $\ln \mathcal{O}_N^S$ lies in responding also to non-GR signals. As an example of this, the right panel in Fig. 7 shows $\ln \mathcal{O}_N^S$ and $\ln \mathcal{B}_N^{GR}$ as a function of the amplitude of a fully non-GR injection. Here, we have chosen to inject a particular model of vector signal developed in [21], but the results are generic.

Note that, for sufficiently loud signals, \mathcal{H}_{GR} becomes preferable over \mathcal{H}_N (hence $\ln \mathcal{B}_N^{GR} > 0$), even when the injection model does not match the search; this is because the noise evidence drops faster than GR’s and becomes very small (i.e. the data do not look at all like Gaussian noise, although they do not match the expected GR signal well either). The particular SNR at which this occurs will depend on the overlap between the antenna patterns of the

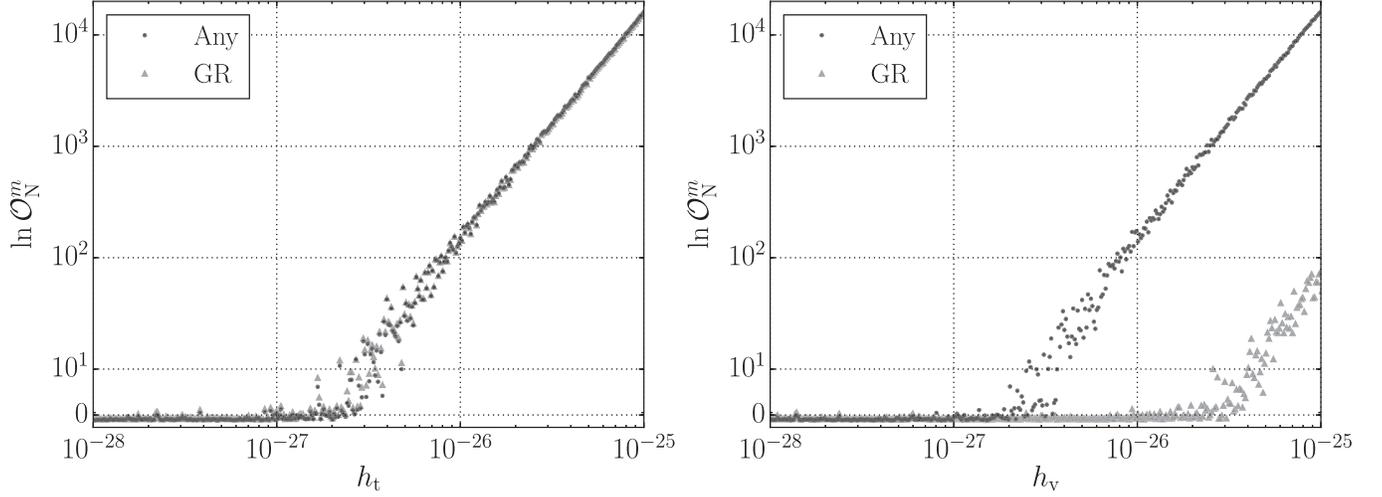


FIG. 7. Expected sensitivity to GR and vector injections. Log-odds of any-signal (\mathcal{H}_S , black circles) and GR (\mathcal{H}_{GR} , gray triangles) vs noise (\mathcal{H}_N) hypotheses, as a function of injection amplitude, for signals corresponding to both GR (left) and the vector-only model from [21] (right). The any-signal odds is defined in Eq. (43). Each of the 500 points corresponds to a data instantiation (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar signal of the indicated strength. The injections were performed with random values of the nuisance phase parameters, and the data were analyzed coherently across detectors. A logarithmic scale is used for the y axis, except for a linear stretch corresponding to the first decade.

injection and those of GR, and will consequently vary among sources.

For the interesting case of scalar-tensor theories (here, templates composed of GR plus an extra breathing component, and denoted “GR + s”), the behavior is slightly different. This is both because GR + s has an extra amplitude degree of freedom (a_s) and, as discussed in Sec. III A 1, because \mathcal{H}_{GR} can be recovered as a special case of \mathcal{H}_{GR+s} (namely, when $a_s \rightarrow 0$). In Fig. 8, we present the log-odds of signal vs noise hypotheses as a function of injected GR (x axis) and scalar (y axis) strengths. These plots divide the h_s – h_t plane in roughly two regions where the associated signal model (\mathcal{H}_S , \mathcal{H}_{GR} or \mathcal{H}_{GR+s}) is preferred (black) and where it is not (red). The latter

corresponds to the area of parameter space associated with subthreshold signals that cannot be detected.

As expected, the best coverage is obtained when analyzing the data using the model matching the injection, GR + s, (rightmost plot) or the all-signal model (leftmost plot). In both these cases, the results improve with either scalar or tensor SNR. In contrast, the GR analysis (center plot) is sensitive to tensor strain, but, as evidenced by the extended red region in the central plot, it misidentifies strong scalar signals as noise. Nevertheless, if the scalar component is larger than $\sim 5 \times 10^{-26}$, the GR analysis will disfavor the noise hypothesis, even for a small tensor component, as in the right panel of Fig. 7; this is the same behavior observed in Fig. 7. In contrast, the any-signal

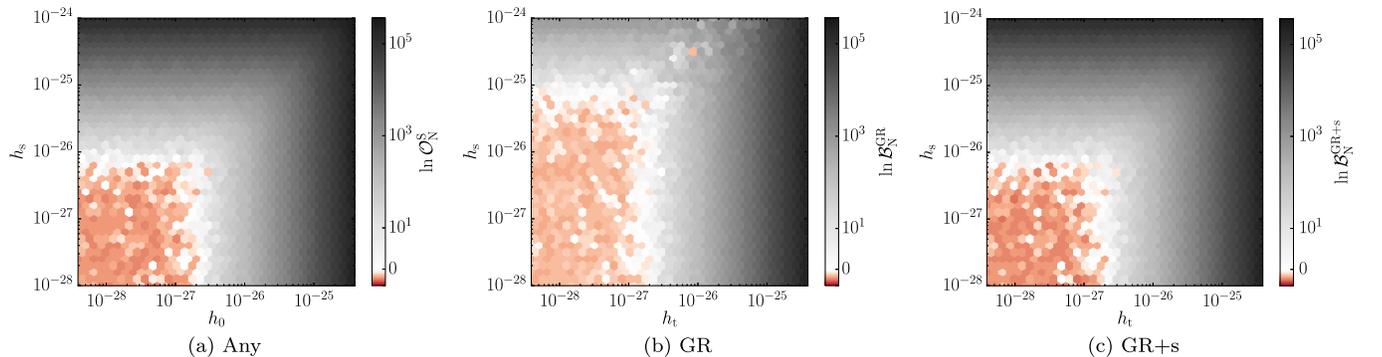


FIG. 8. Expected sensitivity to scalar-tensor injections. Log-odds of any-signal (\mathcal{H}_S , left), GR (\mathcal{H}_{GR} , center) and GR + s (\mathcal{H}_{GR+s} , right) hypotheses vs noise. The any-signal odds is defined in Eq. (43). Each plot was produced by analyzing 2500 instantiations of data (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar GR + s signal of the indicated tensor (x axis) and scalar (y axis) amplitudes. The color of each hexagon represents the average value of the log-odds in that region of parameter space; color is normalized logarithmically, except for a linear stretch in the $(-1, 1)$ range.

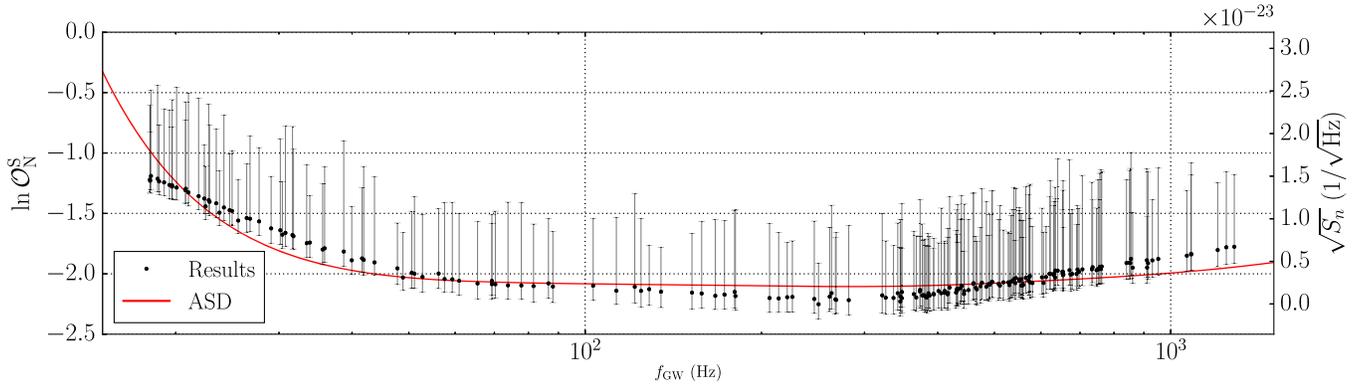


FIG. 9. Signal log-odds vs GW frequency for noise-only data. Circles mark the mean of the distribution of $\ln \mathcal{O}_N^S$, as a function of the expected GW frequency for each pulsar in our set; vertical lines indicate one-sided standard deviations for each source. Each data point and corresponding bars summarize the shape of a distribution like Fig. 5 for each of the pulsars, but produced from only 100 runs per source. The effective noise amplitude spectral density $\sqrt{S_n}$ (ASD, red curve), corresponding to the harmonic mean of each detector PSD, is overlaid for comparison (scaling obtained from a linear regression).

analysis is sensitive to the total power of the injected signal, regardless of polarization.

We have produced distributions of background $\ln \mathcal{O}_N^S$, like those of Fig. 5, for all 200 known pulsars in the sensitive band of the three detectors under consideration (same set analyzed in [17]). In Fig. 9, these are represented by their respective means and one-sided standard deviations as a function of the pulsar's GW frequency. The frequency dependence is explained by variations in the instrumental noise spectra. This is explained by the fact that, for a particular prior choice, more information is gained from the data if the noise floor is lower: with less noise it is possible to discard the presence of weaker signals, so the value of $\ln \mathcal{O}_N^S$ decreases.

2. GR vs non-GR

In the presence of a signal, \mathcal{O}_{GR}^{nGR} , as defined by Eq. (45), indicates whether there is reason to believe there is a GR violation or not. Because there could always be an unresolvably small departure from GR, we do not expect our analysis (with priors as chosen) to ever strongly favor the GR hypothesis; rather, in the presence of a GR signal we will find that $\ln \mathcal{O}_{GR}^{nGR}$ remains relatively close to zero, simply meaning that there is no strong evidence for or against non-GR features. This is indeed the behavior observed in the left panel of Fig. 10, where $\ln \mathcal{O}_{GR}^{nGR}$ is shown to be roughly insensitive to tensor injection amplitude. For values of h_t below certain threshold (which, in this case, is around 3×10^{-27}), the search does not detect a

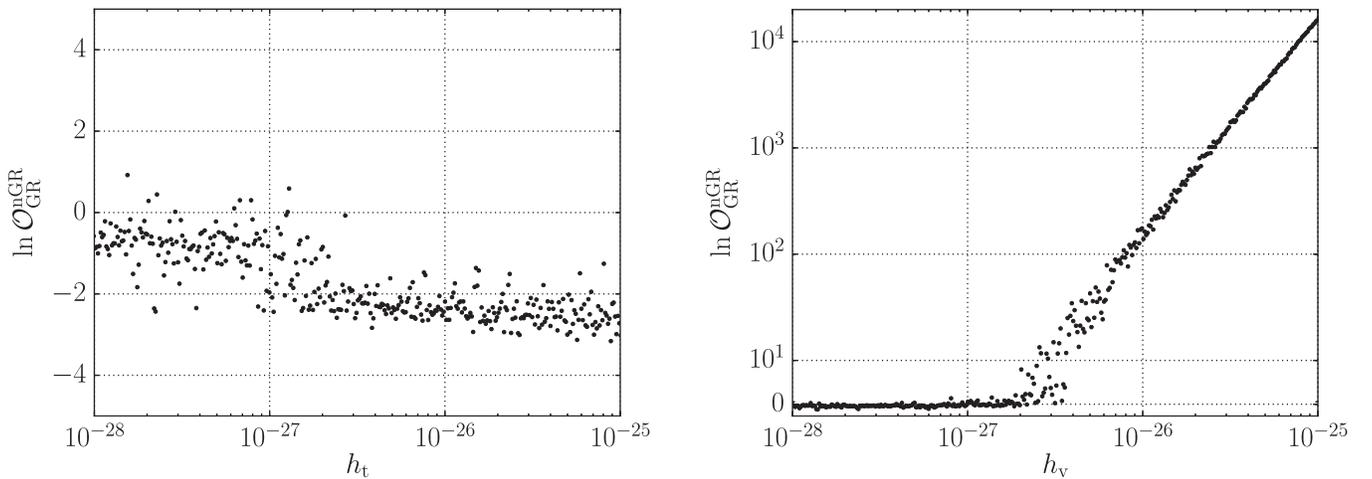


FIG. 10. Categorizing tensor and vector injections (\mathcal{H}_{nGR} vs \mathcal{H}_{GR}). Non-GR vs GR log-odds, as a function of effective injection amplitude, for both GR (left) and the vector-only model from [21] (right). Each of the 500 points corresponds to a data instantiation (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar signal of the indicated strength. The injections were performed with random values of the nuisance phase parameters, and the data were analyzed coherently across detectors. Note that, on the right, a logarithmic scale is used for the y axis, except for a linear stretch corresponding to the first decade.

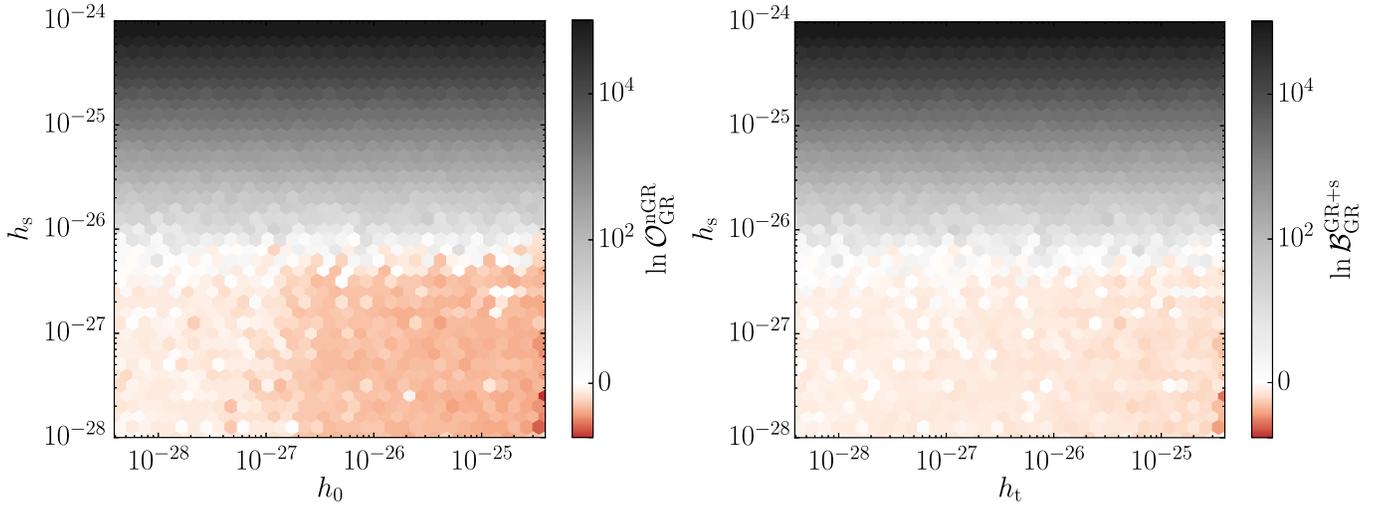


FIG. 11. Categorizing scalar-tensor injections (\mathcal{H}_{nGR} & $\mathcal{H}_{\text{GR}+s}$ vs \mathcal{H}_{GR}). Log-odds comparing the non-GR and GR + s hypotheses to GR. The non-GR odds is defined in Eq. (45). Each plot was produced by analyzing 2500 instantiations of data (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar GR + s signal of the indicated tensor (x axis) and scalar (y axis) amplitudes. The color of each hexagon represents the average value of the log-odds in that region of parameter space; color is normalized logarithmically, except for a linear stretch in the $(-1, 1)$ range.

signal and, consequently, no information is gained for or against \mathcal{H}_{GR} , i.e. $\ln \mathcal{O}_{\text{GR}}^{\text{nGR}} \sim 0$. The difference between the two populations (below and above threshold) is determined mainly by the choice of amplitude priors.

The behavior of $\mathcal{O}_{\text{GR}}^{\text{nGR}}$ is less ambiguous in the presence of a non-GR signal. For instance, if the data contain a detectable signal that completely lacks tensor components, then $\mathcal{O}_{\text{GR}}^{\text{nGR}}$ will unequivocally reflect this. This is evidenced by the growth of $\ln \mathcal{O}_{\text{GR}}^{\text{nGR}}$ with injected nontensorial SNR in the right panel of Fig. 10. In other words, while the analysis is inconclusive for GR injections because it cannot discard

the presence of subthreshold non-GR components hidden by the noise, vector signals are clearly identified as not conforming to GR. This is a reflection of the fact that, as mentioned in the introduction, any evidence of a nontensorial component is fatal for GR, but absence of non-GR components does not mean Einstein's theory is necessarily correct.

As might be expected, $\mathcal{O}_{\text{GR}}^{\text{nGR}}$ responds to non-GR signals that include a tensor component with a combination of features from both panels of Fig. 10. As an example, the left plot of Fig. 11 shows $\mathcal{O}_{\text{GR}}^{\text{nGR}}$ in the presence of GR + s

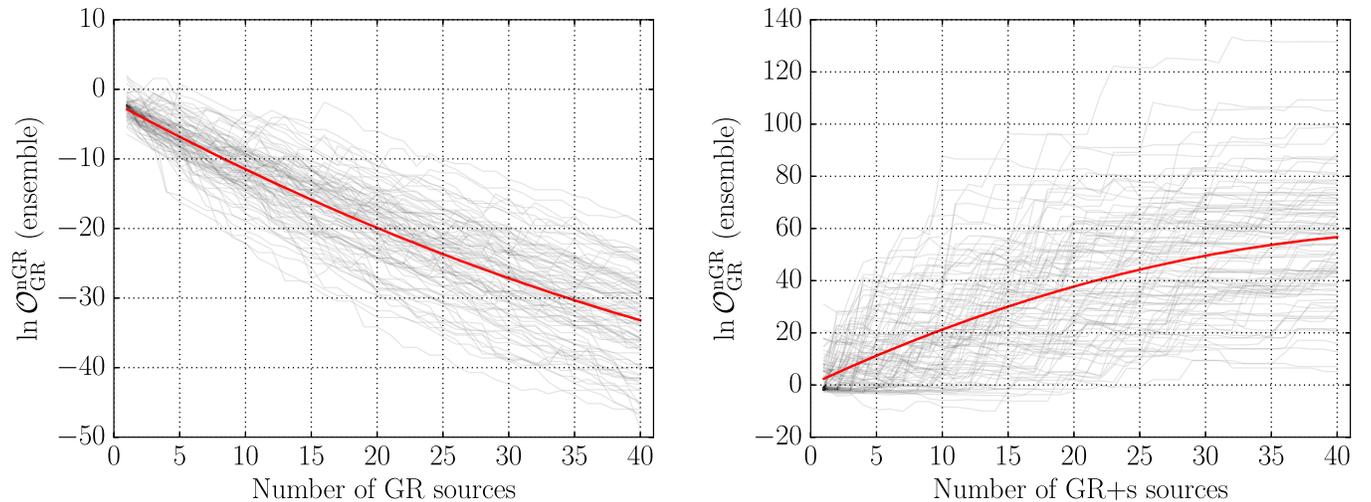


FIG. 12. Ensemble non-GR vs GR log-odds. Non-GR vs GR log-odds computed from data for multiple sources, vs the number of sources in the set. Each light-gray trace marks a possible progression of the ensemble log-odds as new sources are added; the red line corresponds to the best quadratic fit. For each pulsar, we chose an arbitrary data instantiation containing a GR (left) or GR + s (right); GR signals are restricted to $10^{-27} < h_t < 10^{-26}$, while GR + s signals also satisfy $0.3 < h_s/h_t < 1$. We compute the value of $\ln \mathcal{O}_{\text{GR}}^{\text{nGR}}$ for each signal in the set and combine them according to Eq. (54) to obtain the ensemble value plotted in the y axis.

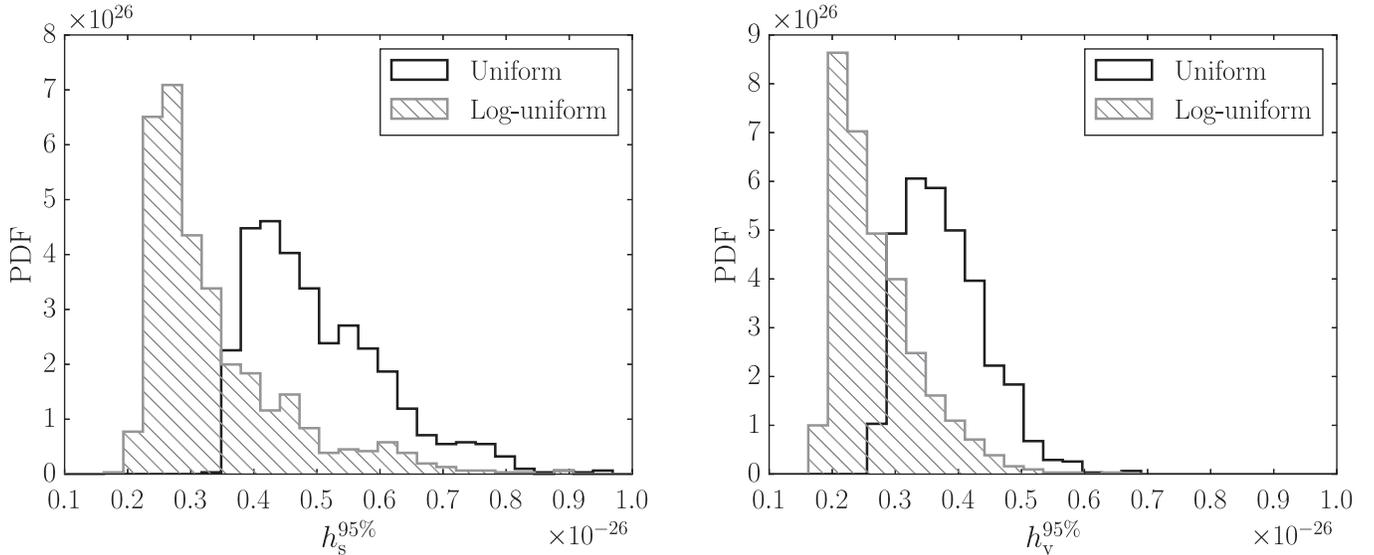


FIG. 13. Expected Crab non-GR upper limits in absence of signal. Histogram of 95%-credible upper limits for the scalar (left) and vector (right) amplitudes, for a set of 1000 noise-only data sets, computed using priors uniform in the amplitude (black) or uniform in the logarithm of the amplitude (hatched gray); the differences between these two priors are discussed in detail in Appendix B. Each instantiation (one time series for each detector: H1, L1 and V1) is made up of simulated Gaussian noise with standard deviation given by the advanced design PSDs. Scalar and vector upper limits are produced using the GR + s and GR + v models respectively.

injections, as a function of injected tensor and scalar amplitudes. This plot can be split into three clearly demarcated regions: one in which the signal is not detected (light red, bottom left), one in which the signal is detected and the non-GR model is preferred (black, top), and one in which the signal is detected but where the evidence for a deviation from GR is not clear due to the predominance of the tensorial component (darker red, bottom right). The first corresponds to the subthreshold population on either side of Fig. 10, while the second and third correspond to the above-

threshold populations on the right and left sides of Fig. 10, respectively; indeed, note that a horizontal slice taken over the red region of the left plot produces a series of points like those in the left panel of Fig. 10. For reference, Fig. 11 also includes the direct comparison of GR + s and GR on the right.

We can make a stronger statement about the agreement of the data with GR by making use of signals from multiple sources, as discussed in Sec. III A 3. The power of combining multiple signals is illustrated in Fig. 12,

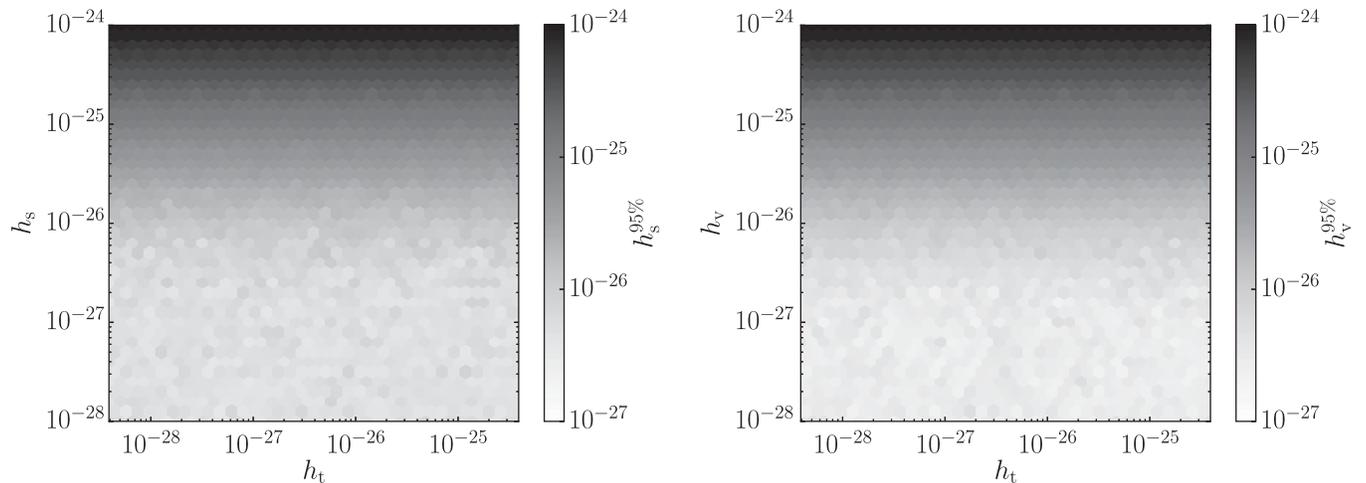


FIG. 14. Expected Crab scalar and vector upper limits in presence of GR + s and GR + v signals. Shading represents the 95%-credible upper limit for the scalar ($h_s^{95\%}$, left) and vector ($h_v^{95\%}$, right) amplitudes, vs the amplitude of injected GR (x axis) and corresponding non-GR (y axis) components. Each plot was produced by analyzing 2500 instantiations of data (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar GR + s (left) or GR + v (right) signal with indicated strains. The color of each hexagon represents the average value of the upper limit in that region of parameter space.

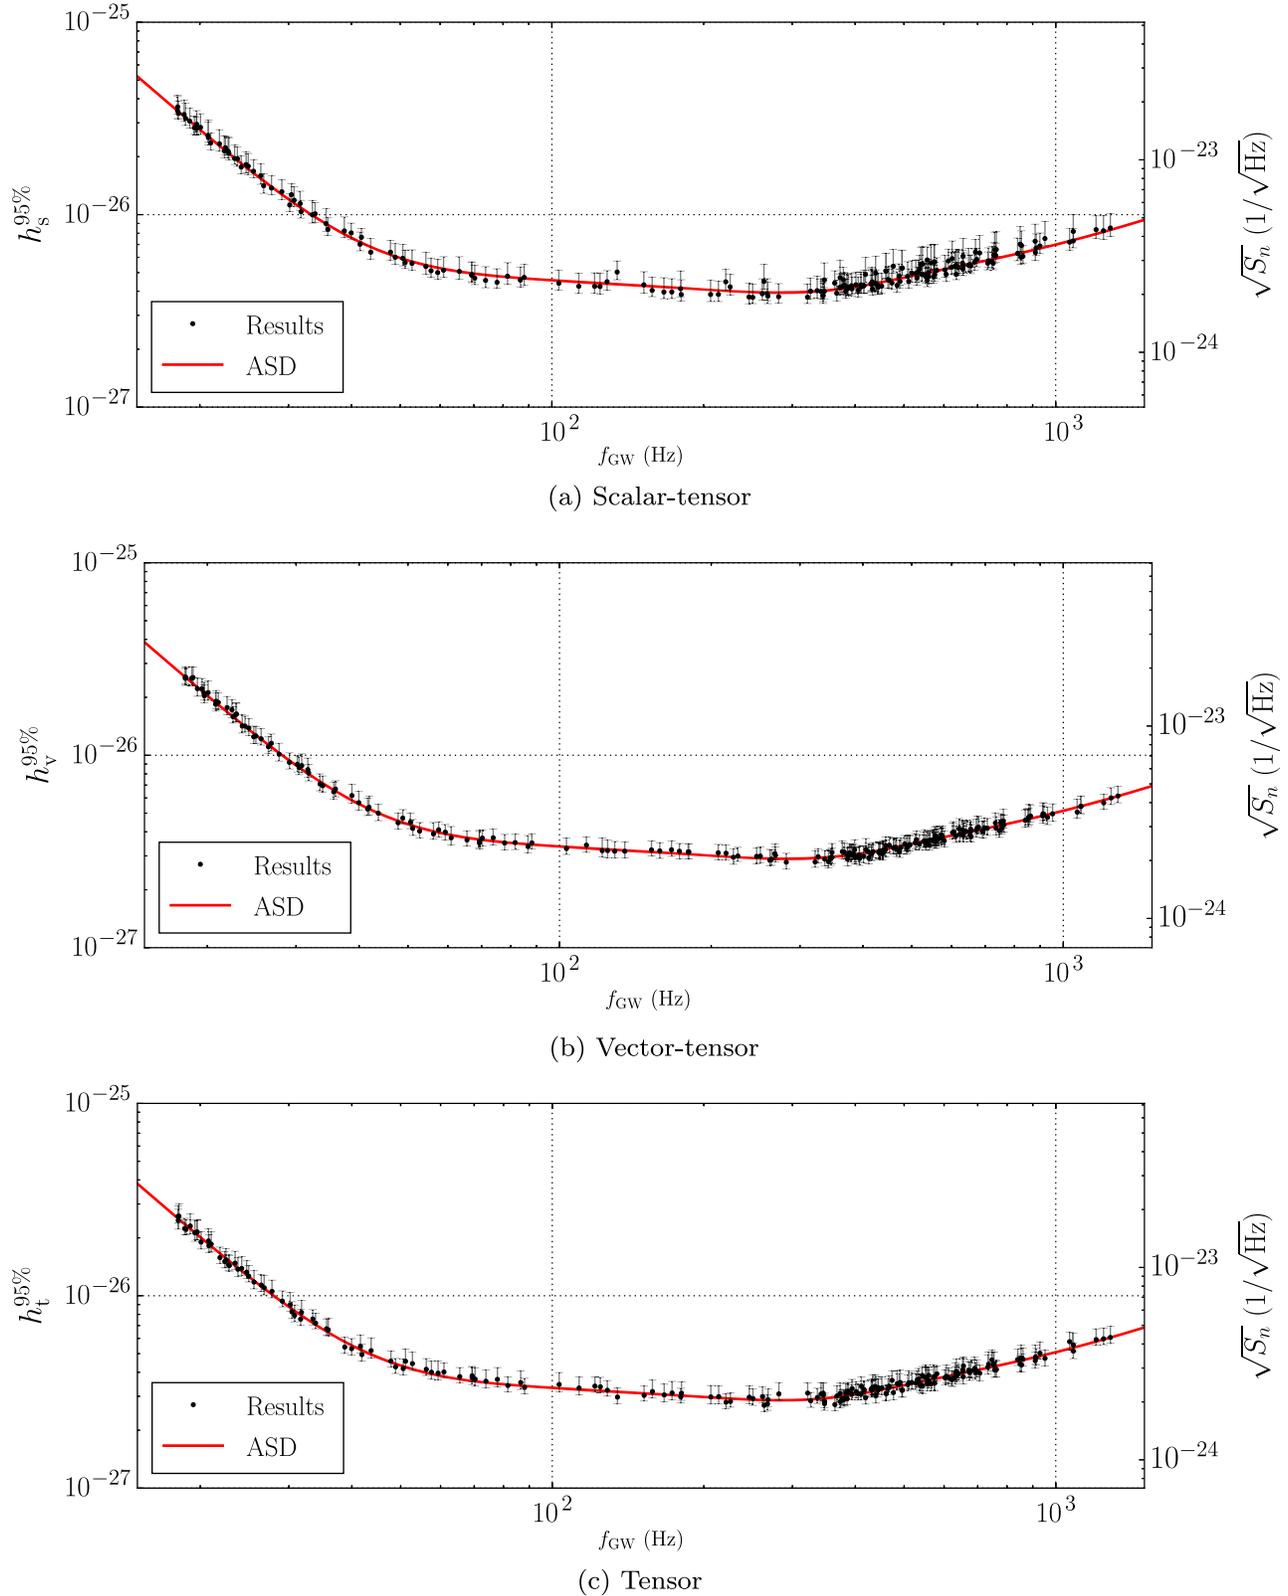


FIG. 15. Expected upper limits in absence of signal vs GW frequency. Circles mark the mean of the distribution of $h_s^{95\%}$ (top), $h_v^{95\%}$ (middle) and $h_t^{95\%}$ (bottom), as a function of expected GW frequency for each pulsar in our set; vertical lines mark one-sided standard deviations for each source. Each datapoint and corresponding bars summarize the shape of a distribution like those of Fig. 13, but produced from 100 noise instantiations each. The scalar, vector and tensor upper limits were produced assuming st, vt and t models respectively. We use uniform priors in all amplitude parameters (see Fig. 13 and Appendix B).

where $\ln \mathcal{O}_{\text{GR}}^{\text{nGR}}$, as defined in Eq. (54), is plotted vs number of GR (left) and GR + s (right) signals detected. Note that this presumes that, for each source, the presence of a signal has already been established from the value of $\ln \mathcal{O}_{\text{N}}^{\text{S}}$. Computing the ensemble $\ln \mathcal{O}_{\text{GR}}^{\text{nGR}}$, as done here, is a good way of summarizing the information contained in the data about the relative likelihoods between the two models, but it provides no information not already present in the set of individual single-source odds.

B. Parameter estimation

When no conclusive evidence for a CW is found in the data, we are still interested in placing upper limits on the strength of possible signals (up to some credibility), and this is done as explained in Sec. III B. By the same token, if a signal consistent with GR is detected, we can always place an upper limit on the amplitude of non-GR modes, even if the odds indicate there is no clear sign of a GR violation.

For instance, we can get a quantitative estimate of our sensitivity to scalar modes from a given source by looking at the distribution of $h_s^{95\%}$, defined in Eq. (62), computed for a set of noise-only data instantiations. Such distribution for the Crab pulsar is presented in the left panel of Fig. 13. Similarly, the right panel presents estimates for the sensitivity to vector modes coming from the Crab pulsar, assuming a vector-tensor model. In this case, however, the quantity plotted is the upper limit on total, effective vector strain amplitude h_v , Eq. (68). These plots include distributions produced using the same log-uniform prior used to obtain Bayes factors, as well as more conservative ones obtained using uniform amplitude priors (see Appendix B). In either case, the magnitude of $h_v^{95\%}$ is comparable to that of $h_s^{95\%}$.

Interestingly, our ability to measure scalar and vector amplitudes is unaffected by the presence of other modes. We illustrate this for the Crab pulsar in Fig. 14, which results from analyzing data with GR + s (left) and GR + v (right) injections. There we plot $h_s^{95\%}$ as a function of scalar and tensor injection amplitudes on the left, and $h_v^{95\%}$ as a function of vector and tensor injection amplitudes on the right. From these plots, one can conclude that $h_s^{95\%}$ and $h_v^{95\%}$ are sensitive only to the corresponding scalar and vector components, and not by h_t . (It is worth emphasized that the upper limits, $h_s^{95\%}$ and $h_v^{95\%}$, are well-defined even when the non-GR component is strong enough to be detected, as is the case for the darker-colored regions.)

As shown previously in the literature, the mean of distributions like those of Fig. 13 will scale with $\sqrt{S_n(f)}/T$, where $S_n(f)$ is the effective PSD of the detector noise at the expected GW frequency f , and T is the integration time (cf. Eq. (26) and Fig. 1 of [32]). Because of this, the mean of this distribution will vary with the source's expected GW frequency, as shown in Fig. 15. Following convention, these upper limits are computed

using uniform amplitude priors, which means they are a factor of a few less stringent than those obtained with a log-uniform prior (see Fig. 13 and Appendix B). Also, for completeness, Fig. 15 also includes the expected tensor upper limits, $h_t^{95\%}$. Note that those values are not the same as would be obtained by the standard GR-only search, because that analysis looks at the triaxial h_0 of Eq. (29), rather than h_t .

In order to compare our sensitivity to the different polarizations, in Fig. 16 we histogram the tensor and vector upper limits as a ratio of the scalar upper limits—this includes the $h_t^{95\%}$ and $h_v^{95\%}$ values shown in Fig. 15, as well as the limits on the individual amplitudes from which they are constructed ($h_+^{95\%}$, $h_\times^{95\%}$, $h_x^{95\%}$ and $h_y^{95\%}$). The mean

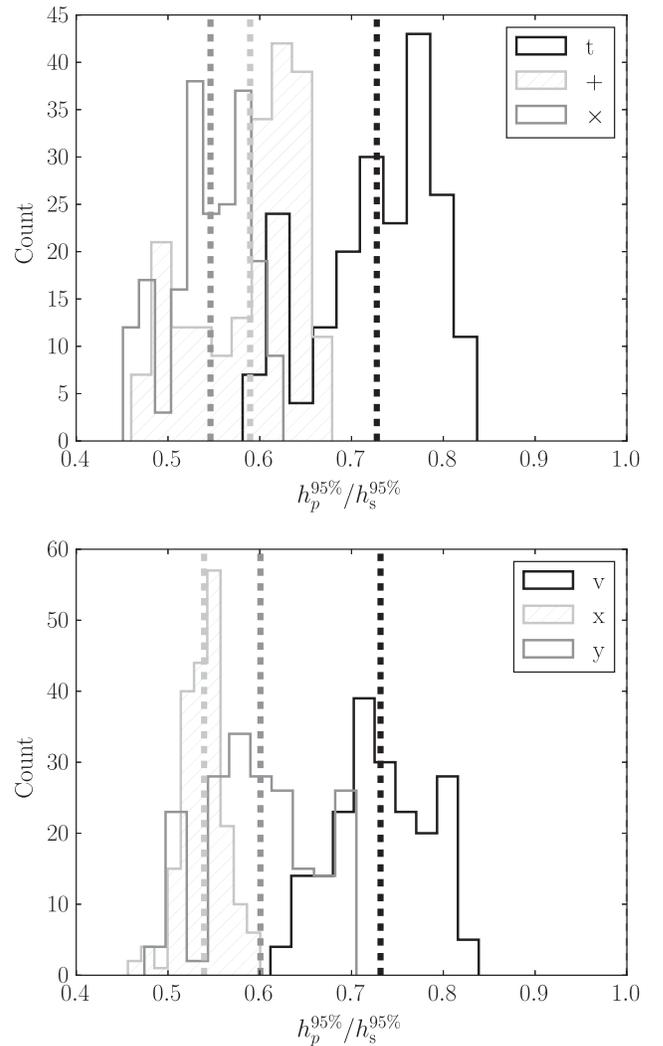


FIG. 16. Tensor and vector upper limits as a ratio of scalar upper limits. Histogram of tensor (top) and vector (bottom) upper limits divided by the scalar upper limit for each pulsar. The top plot shows ratios for $h_t^{95\%}$ (black), $h_+^{95\%}$ (light gray, hatched), and $h_\times^{95\%}$ (dark gray); the bottom plot shows ratios for $h_v^{95\%}$ (black), $h_x^{95\%}$ (light gray, hatched), and $h_y^{95\%}$ (dark gray). Vertical dashed lines mark the mean of each distribution.

of these distributions (vertical dashed lines in Fig. 16) indicate that, for most pulsars, the scalar upper limit is slightly larger in magnitude than those for the $+$, \times , x or y modes; this systematic effect is a manifestation of the decreased sensitivity of quadrupolar detectors to scalar waves, which was discussed in Sec. II A (see, in particular, Fig. 2). The fact that the difference between $h_s^{95\%}$ and $h_t^{95\%}$, or $h_v^{95\%}$, is less pronounced can be easily be explained as a statistical factor arising from the definitions of h_t and h_v as square-roots of sums of squares, Eqs. (67) and (68). Both these scalings are discussed in more detail in Appendix D.

VI. CONCLUSION

We have developed a Bayesian framework to detect CW signals from known sources regardless of polarization content, to disentangle the modes present in a given signal, and to constrain the amplitudes of extra polarizations that may be hiding under the noise. We have implemented this as an extension of LIGO’s Bayesian targeted CW search pipeline [33], and thus benefit from the power of the nested sampling algorithm on which it is based.

We have tested our methods on one year of simulated noise for three advanced-era detectors at design sensitivity (H1, L1, V1), and prepared for a set of multiple known sources in their frequency band. This allows us to estimate our future sensitivity to CW polarizations, in this most optimistic case. Under these conditions and for the Crab pulsar in particular, we expect signals of any polarization to become detectable for characteristic strain amplitudes $h \gtrsim 3 \times 10^{-27}$ (Figs. 7 and 8); this threshold will vary among sources, due to differences in position (sky location and orientation) and detector PSD at the expected GW frequency (cf. e.g. Fig. 9). Furthermore, the value of this threshold will decrease linearly with the square-root of the observation time [32].

A signal louder than the detection threshold will allow us to determine whether its polarization content is consistent with GR or not, and the strength of this statement will depend almost exclusively on the power of the non-GR component (Figs. 10 and 11). In other words, from a model-selection standpoint, the non-GR hypothesis will only be unequivocally favored if the total power in non-GR modes is greater than the threshold value, regardless of the strength of the GR modes. However, for signals that do not satisfy this, we may always place upper limits on non-tensorial amplitudes and thus constrain deviations from GR; for instance, Fig. 15 presents the most optimistic projections for 95%-credible upper limits for scalar and vector amplitudes of CW signals from all pulsars in our set ($h_s^{95\%} \sim 4 \times 10^{-27}$ and $h_v^{95\%} \sim 3 \times 10^{-27}$, in the best case). As far as we are aware, these are the first generic estimates of sensitivity to scalar and vector polarizations ever published [47].

From our projected upper limits, we have found that, at design sensitivity, the LIGO-Virgo network will be generally less sensitive to continuous scalar signals than to the individual vector or tensor modes by factors of 0.45–0.7, depending on the location of the source (Fig. 16); this diminished sensitivity to scalar modes stems from the quadrupolar nature of the detector antenna patterns (Fig. 2 and Appendix D). Also, our injection studies indicate that the upperlimits on the amplitudes of non-tensorial modes will be roughly unaffected by the presence or absence of a tensor signal in the data (Fig. 14).

Although the results presented here made use of simulated Gaussian noise, the procedure is identical for actual detector data. Furthermore, the assumption of Gaussianity has been shown to hold relatively well for real CW data [19], so the actual sensitivity limits should not be far from those presented here. If the data are strongly non-Gaussian, however, one must be careful in using $\ln \mathcal{O}_N^S$ for detection purposes and may instead wish to adopt one of the strategies suggested in Sec. III A 4.

Another important limitation of our results is that here we only consider CW signals emitted at $f = 2f_{\text{rot}}$, while it is to be expected that other mechanisms (within GR or not) allow emission at other harmonics, $f = f_{\text{rot}}$ in particular. Yet, the only change required to account for this is to modify the template in Eq. (12) to include terms at different harmonics; the ability to do this already exists within our current infrastructure. We also assume that other aspects of the waves, like their speed, remain in agreement with the GR prediction, an assumption that will be relaxed in a future study.

ACKNOWLEDGMENTS

The authors would like to thank Ian Jones and Walter Del Pozzo for carefully reading this manuscript and providing insightful suggestions; we also thank Tjonnje Li and Carver Mead, as well as many colleagues in the LIGO Scientific Collaboration Continuous Waves group, for many useful comments. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative Agreement No. PHY-0757058. We are grateful for computational resources provided by Cardiff University and funded by an STFC grant supporting UK Involvement in the Operation of Advanced LIGO. M. P. is funded by the STFC under Grant No. ST/N005422/1. Plots produced using Matplotlib [48]. This paper carries LIGO Document No. LIGO-P1600305.

APPENDIX A: TENSOR MODELS

A conceptual distinction can be drawn between the triaxial GR model and a free-tensor model that includes $+$ and \times but does not restrict their relative amplitudes (denoted “t”). The former has four free parameters (overall

amplitude, h_0 ; overall phase, ϕ_0 ; inclination, ι ; polarization, ψ) and corresponds to a signal template of the form [same as Eq. (29)]:

$$\Lambda_{\text{GR}}(t) = \frac{1}{2} h_0 e^{i\phi_0} \left[\frac{1}{2} (1 + \cos^2 \iota) F_+(t; \psi) - i \cos \iota F_\times(t; \psi) \right]. \quad (\text{A1})$$

This is a reparametrization of the free-tensor model, which also has four parameters (plus amplitude, a_+ ; cross amplitude, a_\times ; plus phase, ϕ_+ ; cross phase, ϕ_\times) and whose template is [same as Eq. (28)]:

$$\Lambda_t(t) = \frac{1}{2} [a_+ e^{i\phi_+} F_+(t; \psi = 0) + a_\times e^{i\phi_\times} F_\times(t; \psi = 0)]. \quad (\text{A2})$$

If ψ and ι are known, it is clear that the two models are different, since \mathcal{H}_{GR} has two free parameters (h_0 , ϕ_0) and \mathcal{H}_t has four (a_+ , a_\times , ϕ_+ , ϕ_\times). If the orientation is not fixed, however, the two models span the same signal space. This is because there is a degeneracy between ψ and a_+ , a_\times due to the way the antenna patterns transform under changes in ψ :

$$F_+(t; \psi') = F_+(t; \psi) \cos 2\Delta\psi + F_\times(t; \psi) \sin 2\Delta\psi, \quad (\text{A3})$$

$$F_\times(t; \psi') = F_\times(t; \psi) \cos 2\Delta\psi - F_+(t; \psi) \sin 2\Delta\psi, \quad (\text{A4})$$

with $\psi' = \psi + \Delta\psi$. Equations (A3) and (A4) can be derived from Eqs. (1) and (2), respectively, as in [23] (or see, e.g., Sec. 9.2.2 in [18]). Consequently, changing $\psi \rightarrow \psi'$ in Eq. (A2) is equivalent to leaving ψ fixed [at, say, $\psi = 0$ as in Eq. (A2)] while replacing the plus and cross complex amplitudes by

$$a_+' e^{i\phi_+'} = a_+ e^{i\phi_+} \cos 2\Delta\psi - a_\times e^{i\phi_\times} \sin 2\Delta\psi, \quad (\text{A5})$$

$$a_\times' e^{i\phi_\times'} = a_\times e^{i\phi_\times} \cos 2\Delta\psi + a_+ e^{i\phi_+} \sin 2\Delta\psi. \quad (\text{A6})$$

This is illustrated in Fig. 17.

These rotational properties are easily understood by recalling that GW polarizations can be defined in *any* frame, although a given signal might look more or less simple given the choice of frame. Equations (A3) and (A4) provide the transformation between frames that are coaligned except for a rotation of $\Delta\psi$ around their z axes. Because waveform predictions, e.g. Eq. (A1), are made in specific frames, it is important to orient the wave frame appropriately when working with a given theory and emission mechanism. However, if the signal parametrization is independent of any theory, e.g. Eq. (A2), one is free to pick any frame (i.e. any ψ).

The relationship between the different tensor model parameters is reflected in the posterior probability plots

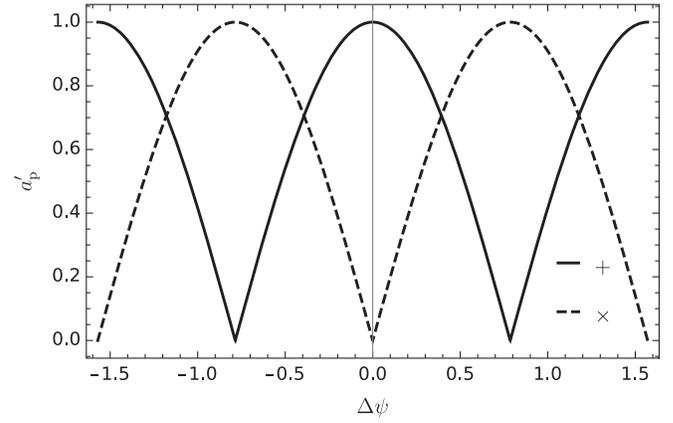


FIG. 17. Effect of changing polarization angle. Norm of the complex plus (a_+ , solid line) and cross (a_\times , dashed line) weights after rotating the source by $\Delta\psi$ in the plane of the sky, i.e. letting $\psi \rightarrow \psi' = \psi + \Delta\psi$; this transformation is expressed in Eqs. (A5) and (A6). In this case, we start from $a_+ = 1$, $a_\times = 0$ and $\psi = 0$.

of Fig. 18. For fixed orientation, both the triaxial (a) and free-tensor (b) analyses accurately determine the amplitude and phase of the injected signal. In panel (b), a_+ and a_\times are constrained to lie within a region consistent with $h_{\text{inj}}^2 = a_+^2 + a_\times^2$ and $a_+/a_\times = (1 + \cos^2 \iota)/(2 \cos \iota)$, for an effective injection amplitude given by $h_{\text{inj}}^2 = h_0^2 (1 + \cos^2 \iota)^2 / 4 + h_0^2 \cos^2 \iota$, as in Eq. (67). When the orientation is allowed to vary, we observe the expected correlations between the recovered triaxial amplitude and the orientation parameters in panel (c); in this case, ψ and $\cos \iota$ will also become correlated, as better shown in Fig. 19.

The degeneracy between ψ and a_+ , a_\times is particularly evident in Fig. 18(d), where the one-dimensional PDF for ψ shows that this parameter cannot be constrained, even for a loud signal. Furthermore, joint posteriors between ψ and a_+ & a_\times confirm that this is due to the degeneracy from Eqs. (A5) and (A6), as seen by comparing these two-dimensional PDFs to Fig. 17. Physically, this is a consequence of the fact that we are free to orient the polarization frame as we wish.

Because their signal templates are degenerate when ψ and ι are allowed to vary, the distinction between \mathcal{H}_{GR} and \mathcal{H}_t is not really meaningful for unfixed orientation. This can be seen from the values of $\ln \mathcal{B}_{\text{GR}}^l$ in the cases of known and unknown orientations, as in Fig. 20. On the left panel, \mathcal{H}_{GR} is defined with specific values of ψ and $\cos \iota$ that match the injections; on the right, the \mathcal{H}_{GR} priors allow ψ and $\cos \iota$ to range over their full ranges, and the injections are performed with random values of both. When the orientation is fixed, \mathcal{H}_{GR} will always be preferred to \mathcal{H}_t for resolvable signals because of its lower Occam's penalty; however, that is not true for unfixed orientation. Note that, in the strictest sense, the two hypotheses are not logically equivalent unless their parameter priors are related by the Jacobian of the coordinate transformation between the two

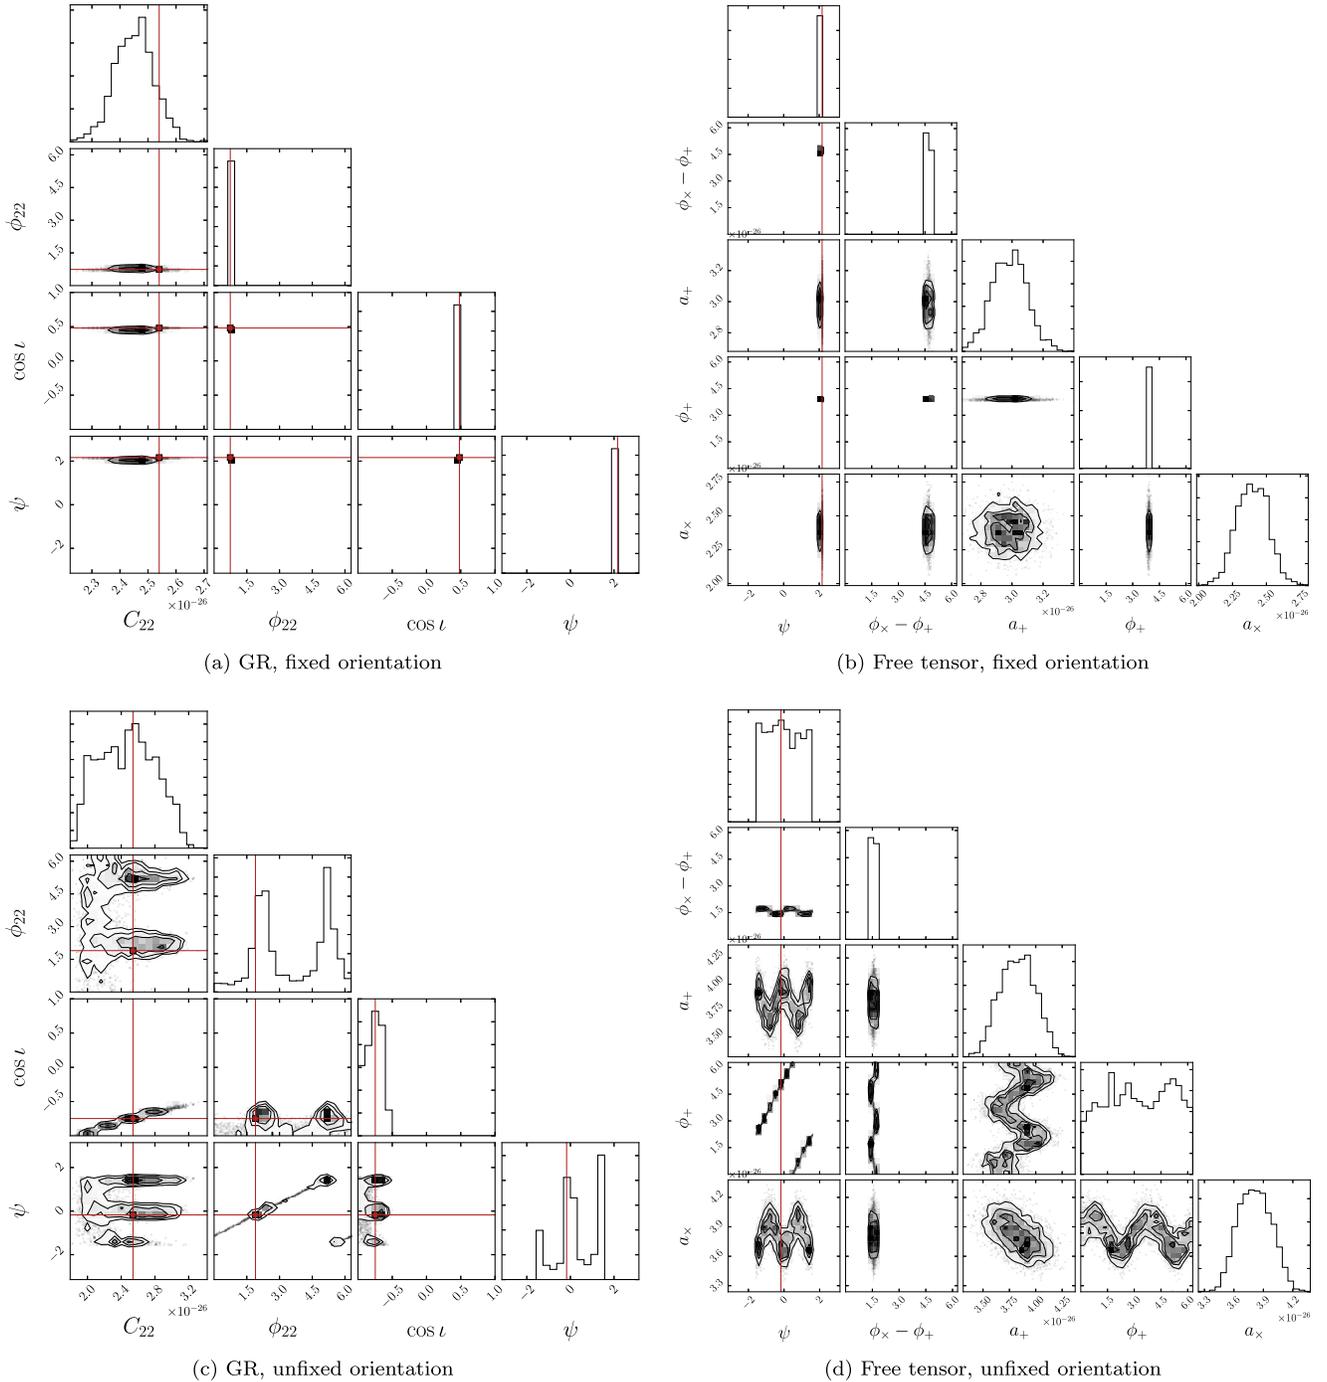


FIG. 18. Tensor posteriors in presence of signal. Posterior PDFs for parameters of \mathcal{H}_{GR} (left) and \mathcal{H}_t (right) with fixed (top) and unfixed (bottom) source orientation (ψ, ι). Each panel consists of a corner plot displaying the two-dimensional posteriors for each pair of parameters as indicated by the x and y labels, with the diagonals showing a histogram of the one-dimensional PDF for each parameter [i.e. the one-dimensional PDF obtained after marginalization of the multidimensional posterior PDF all other quantities, as in Eq. (61)]. The data analyzed contain signals with parameters indicated by the red lines; note that $C_{22} = h_0/2$ is the quantity that was actually used to parametrize GR triaxial amplitudes in the code [33]. In the top row, $\cos \iota$ and ψ are fully known, and their resolution in these plots is limited by binning only. These plots were produced using the `corner.py` package [49].

parametrizations, Eqs. (A1), (A2); only in that case all regions of signal space are treated equally by \mathcal{H}_{GR} and \mathcal{H}_t . This explains the variation in values of $\ln \mathcal{B}_{\text{GR}}^t$ on the right of Fig. 20.

If one knew the source orientation and one believed that the only viable mechanism for producing CWs at the assumed frequency in GR is the triaxial model embodied by Eq. (A1), then one could include the free-tensor hypothesis

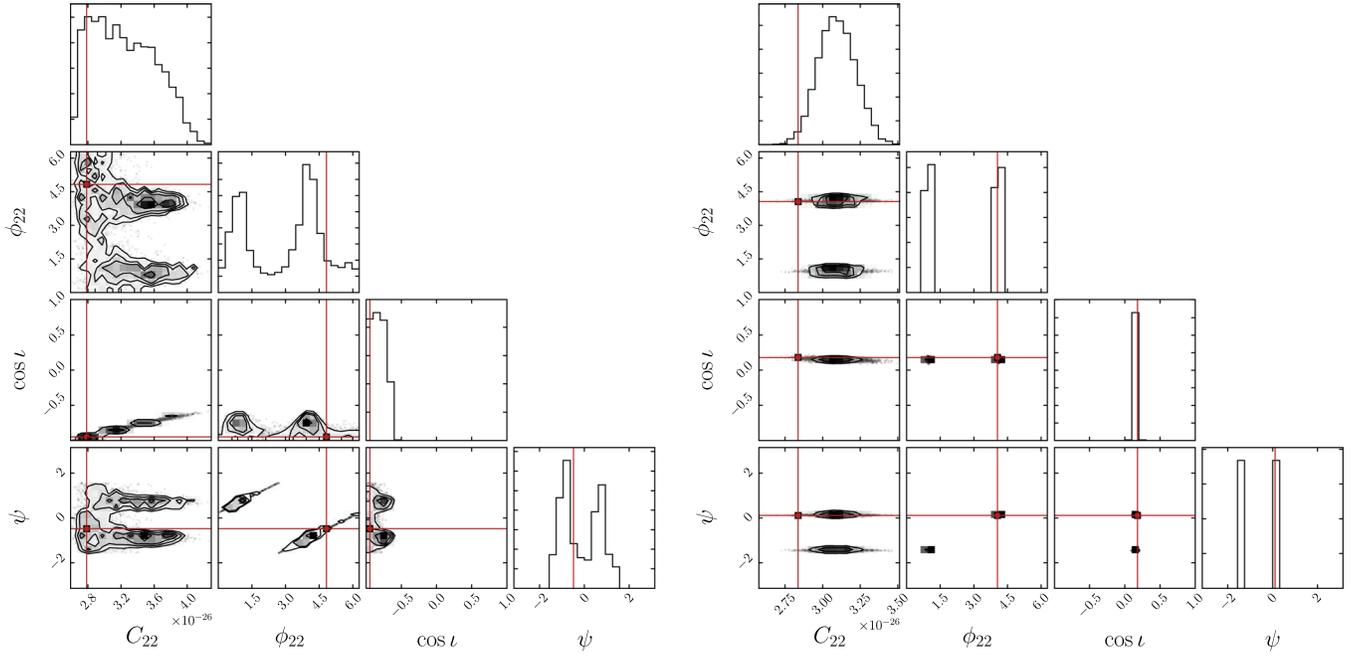


FIG. 19. Effect of inclination. Posterior PDFs for parameters of \mathcal{H}_{GR} with unfixed source orientation (ψ, ι). Each panel consists of a corner plot displaying the two-dimensional posteriors for each pair of parameters as indicated by the x and y labels, with the diagonals showing a histogram of the one-dimensional PDF for each parameter. The data sets analyzed contain signals with parameters indicated by the red lines; note that $C_{22} = h_0/2$ is the quantity that was actually used to parametrize GR triaxial amplitudes in the code [33]. On the left, the injected signal corresponds to a face-off source ($\cos \iota \approx -1$), making it difficult to constrain the polarization angle ψ ; on the right, the injection has similar amplitude but corresponds to an edge-on source ($\cos \iota \approx 0$), making it easy to constrain ψ [modulo $\pi/2$ due to the $2\Delta\psi$ dependence of Eqs. (A3) and (A4)]. These plots were produced using the `corner.py` package [49].

and all of its derivatives (i.e. t, st, vt, stv) in the non-GR set \tilde{M} , on top of $\{s, v, sv, \text{GR} + s, \text{GR} + v, \text{GR} + sv\}$. Doing so would mean treating a tensor-only signal that does not conform to Eq. (A1) as evidence of a GR violation, rather

than of a different emission mechanism within GR. Given the many simplifications intrinsic to the triaxial model, however, having that much confidence in its validity seems unwarranted; hence we choose to not take that approach.

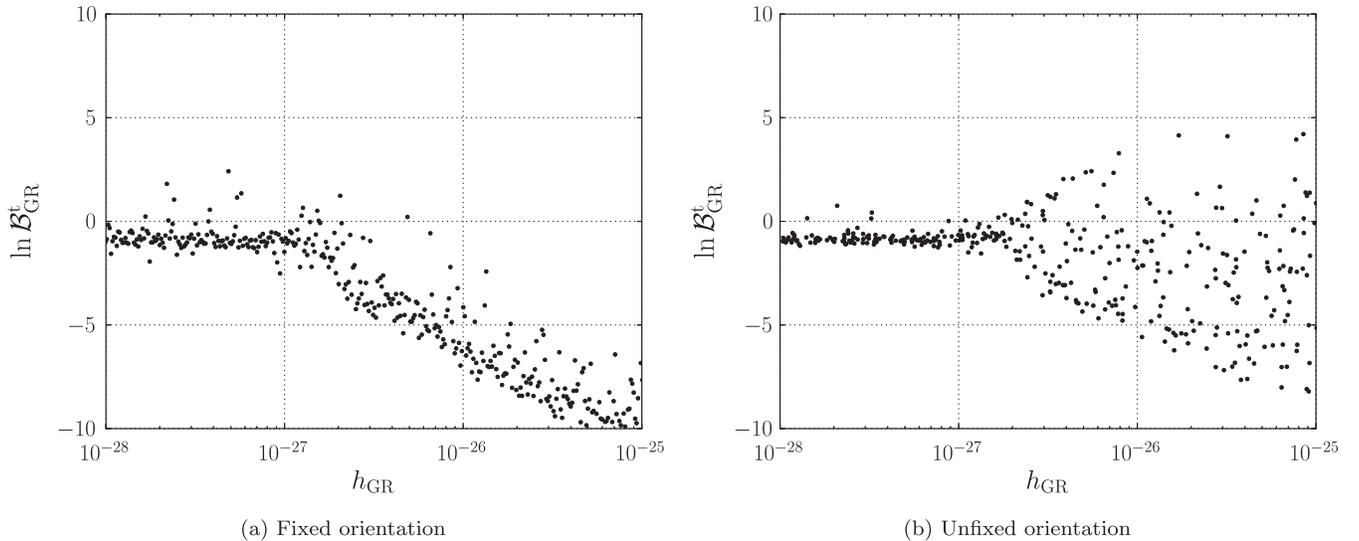


FIG. 20. Free-tensor vs GR. Natural logarithm of the Bayes factor comparing \mathcal{H}_t to \mathcal{H}_{GR} , as a function of GR injection amplitude for fixed (left) and unfixed (right) source orientation. On the left, the analysis correctly gives preference to \mathcal{H}_{GR} for signals above the detection threshold; on the right, however, the analysis is unable to satisfactorily distinguish between \mathcal{H}_t and \mathcal{H}_{GR} , due to the orientation degeneracies discussed in Appendix A.

APPENDIX B: AMPLITUDE PRIORS

Previous CW Bayesian searches targeted to known pulsars have always applied a flat prior on the signal amplitude parameter [17]. This is because flat priors, if wide enough, cause the posterior to be only determined by the likelihood (up to normalization), yielding more conservative upper limits on the signal strength. Furthermore, unlike with priors uniform in the logarithm of the quantity, upper limits derived with flat priors will generally not depend on the limits set by the prior (again, assuming the range allowed extends from zero amplitude to some large value that does not truncate the likelihood).

Upper limits obtained using log-uniform priors (uniform in the logarithm of the quantity) will, generally, be dependent on the range of the prior, although not strongly. For example, consider a one-dimensional problem on some positive parameter x . For simplicity, further assume we have a flat likelihood between $x = 0$ and an upper cutoff at $x = x_{\max}$; then, x_{\max} will necessarily also be an upper bound for the posterior. Because the likelihood is uniform, below the cutoff the posterior will be determined, up to normalization, by the prior only, i.e. for $x < x_{\max}$,

$$p(x|\mathbf{B}, \mathcal{H}) \propto p(x|\mathcal{H}). \quad (\text{B1})$$

Now consider a log-uniform prior $p(x|\mathcal{H}) \propto d(\log x) \propto 1/x$, with a lower bound x_{\min} , such that $0 < x_{\min} < x_{\max}$. Because such prior is uniform in the $\log x$, this implies that the 95%-credible upper limit on x will be given by:

$$\begin{aligned} \log x^{95\%} &= \log x_{\min} + 0.95(\log x_{\max} - \log x_{\min}) \\ &= \log(x_{\max}^{0.95}/x_{\min}^{0.95-1}). \end{aligned} \quad (\text{B2})$$

Since x_{\max} is set by the likelihood (by construction), if the prior is changed by rescaling x_{\min} by a factor α ,

$$x_{\min} \rightarrow x'_{\min} = \alpha x_{\min}, \quad (\text{B3})$$

then, for a given set of data, the upper limit becomes $x_{\alpha}^{95\%}$, satisfying:

$$x_{\alpha}^{95\%}/x^{95\%} = \alpha^{0.05}. \quad (\text{B4})$$

Thus, the dependence of the upper limit on the range defined by the log-uniform prior is quite weak, as illustrated in Fig. 21. This explains why upper limits obtained with a log-uniform prior differ only by a factor of a few from those obtained with a flat one, as seen in Fig. 13.

However, the flat priors do not properly represent our ignorance of the scale of the signal amplitude. This problem manifests itself in negative Bayes factors that too quickly favor the noise hypothesis if no loud signal is clearly present, rather than reflecting our expectation that a signal might be hiding under the noise. This can be seen in Fig. 22, where we show the distributions of $\ln \mathcal{B}_N^{\text{GR}}$, obtained for several noise-only data instantiations for the

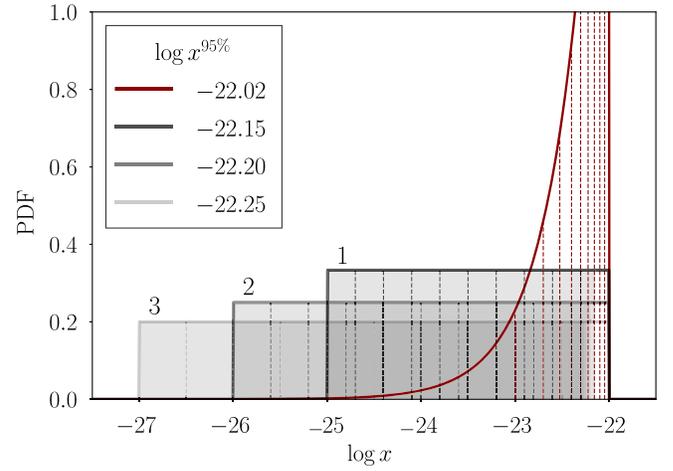


FIG. 21. Log-uniform prior and upper limits. For a one-dimensional random variable x , we show the probability densities corresponding to a uniform likelihood with upper cutoff $\log x_{\max} = -22$ (red) and log-uniform priors with different lower cutoffs ($\log x_{\min} = -25$ for box 1, $\log x_{\min} = -26$ for box 2 and $\log x_{\min} = -27$ for box 3). Vertical dashed lines mark areas of equal probability mass for each distribution. The combined effect of the likelihood and each of the prior distributions is to produce 95%-credible upper limits on x with values shown in the legend. The value obtained using only the likelihood corresponds to that obtained with a uniform prior with a broad enough range. As expected from Eq. (B4), the upper limit is not very sensitive to the lower bound set by the prior.

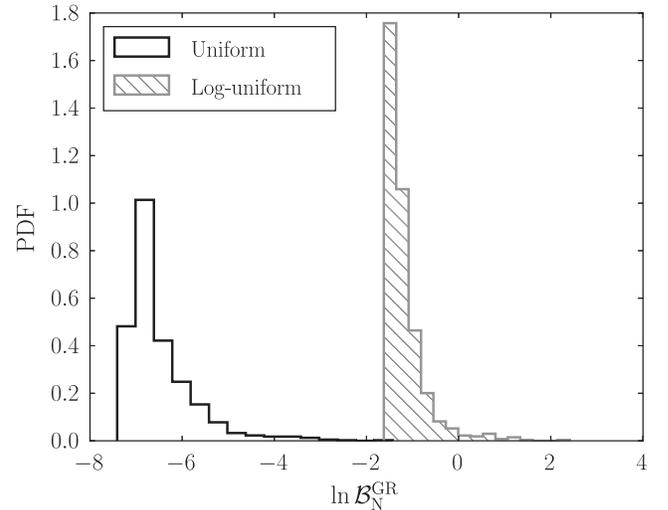


FIG. 22. Log-uniform vs flat amplitude priors. The logarithm of the GR vs noise Bayes factor is computed for 1000 instantiations of Crab pulsar noise. For the GR amplitude h_0 , we apply priors uniform in the quantity (black) and uniform in the logarithm of the quantity (hatched gray). The flat prior causes one to more strongly favor the noise model, due to a larger implicit Occam's penalty.

Crab pulsar, corresponding to flat and log-uniform priors in the GR amplitude parameter, h_0 ; a uniform prior results in lower values of $\ln \mathcal{B}_N^{\text{GR}}$ that strongly favor \mathcal{H}_N . This behavior is not specific to the GR model.

For most of our analysis, we choose to apply priors uniform in the logarithm of all amplitude quantities. However, for the sake of consistency with previous searches and in order to make our limits more conservative, we also present upper limits produced using flat amplitude priors, as shown in Fig. 13.

APPENDIX C: NUMERICAL ERROR

The fractional numerical error in the computation of the natural logarithm of the evidence by nested sampling is usually estimated by

$$\delta[\ln P(\mathbf{B}|\mathcal{H})] \sim \sqrt{H/N_{\text{live}}}, \quad (\text{C1})$$

where N_{live} is the number of live points and H is the information gained in the analysis,

$$H \equiv \int_{\Theta} p(\vec{\theta}|\mathbf{B}, \mathcal{H}) \ln \frac{p(\vec{\theta}|\mathbf{B}, \mathcal{H})}{p(\vec{\theta}|\mathcal{H})} d\vec{\theta}, \quad (\text{C2})$$

a quantity that is easy to estimate from the output of the nested sampling code [42,50].

An example of the actual statistical error as function of SNR is presented in Figs. 23 and 24, where the injected GR signal amplitude serves as proxy for ρ (for fixed PSD). From these plots it becomes apparent that, although the actual error might exceed the estimator of Eq. (C1), its absolute magnitude is quite small and should not affect our results. In any case, Eq. (C1) indicates that any level of accuracy may be achieved by increasing the number of live points (at the cost of increased computational burden). For

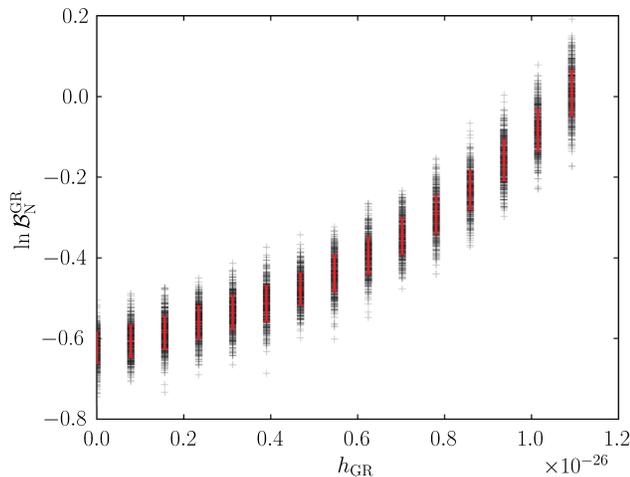


FIG. 23. Numerical error in Bayes factor computation. The logarithm of the GR vs noise Bayes factor is computed 500 times for different values of injected GR signal amplitude. The noise realization is not varied between computations with the same injection strength, only the seed for the random number generator used by the nested sampling algorithm. The red bars mark one standard deviation around the mean.

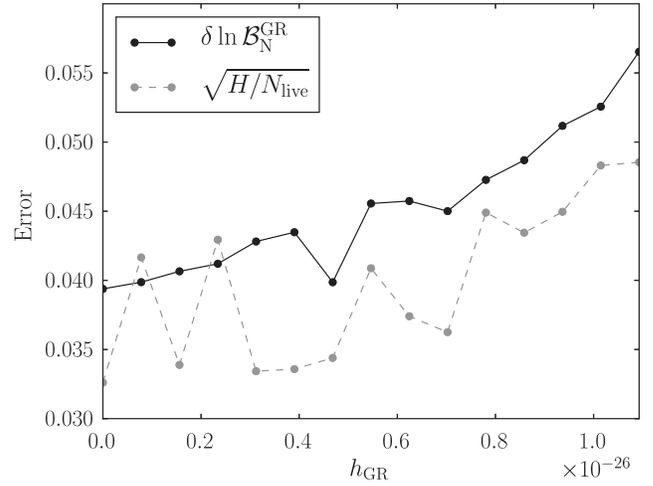


FIG. 24. Observed error and prediction. Error in the computation of the logarithm of the GR vs noise Bayes factor as a function of injected GR signal amplitude. The solid, black curve corresponds to measured standard deviations from the computation of $\ln \mathcal{B}_N^{\text{GR}}$ 500 times per injection strength (red bars in Fig. 23). The dashed, gray curve shows the theoretical prediction for the error in the logarithm of the evidence, Eq. (C1).

more details on the numerical error of the nested sampling algorithm in LALInference, we refer the reader to Sec. IV B of [51].

APPENDIX D: UPPER-LIMIT RATIOS

When comparing upper limits for the different modes, as in Fig. 16, two scalings become apparent: first, the +, \times , x, and y upper limits are, on average, more stringent than those for the scalar polarization by a factor of ~ 1.8 ; second, the upper limits on h_t (h_i) are a factor of ~ 1.3 larger than those on the individual + and \times (x and y) amplitudes.

The scaling between the scalar upper limit and those for the other individual strain amplitudes can be accounted for by the decreased sensitivity of quadrupolar GW detectors to scalar waves. For a single instrument (that is, not a network), this can be appreciated visually from Fig. 2, by noting that for most sky locations the magnitude of the scalar response is considerably less than for the other modes.

To properly evaluate the effect of the detector geometry on the analysis, however, it is necessary to look at the relative SNRs of unit-amplitude scalar, vector and tensor GWs from a given source, as they are received by the detector network under consideration (H1, L1, V1) after some fixed observation time. Assuming all detectors have comparably noise levels, the network SNR can be proxied by the root-mean-square (rms) amplitude of the effective network antenna patterns, defined by

$$F_{p,\text{net}}^{\text{rms}} = \sqrt{\frac{1}{T} \int_0^T \sum_d F_{p,d}^2(t)}, \quad (\text{D1})$$

for each polarization p , some long observation time T , and where the sum is over detectors d . [Here we have fixed the source and detector parameters so that the F_p 's of Eqs. (1)–(5) are now just simple functions of time.] We may then compute this for all five polarizations and for multiple sources to obtain a sky-average of the ratio of the scalar rms antenna pattern to those of the other polarizations. We find this ratio to be roughly ~ 0.55 for all polarizations, in agreement with Fig. 16, since we should expect

$$\left\langle \frac{F_{s,\text{net}}^{\text{rms}}}{F_{p,\text{net}}^{\text{rms}}} \right\rangle \sim \left\langle \frac{h_p^{95\%}}{h_s^{95\%}} \right\rangle, \quad (\text{D2})$$

where the average $\langle \cdot \rangle$ is taken over multiple sources distributed across the sky.

The relation between the h_t (h_v) upper limits and those for their component amplitudes, $+$ and \times (x and y), can be easily understood by noting that, if using flat priors and in the absence of signal, the marginalized posteriors for each of the component amplitudes (h_+ , h_\times , h_x , h_y) will roughly be described by a one-sided normal distribution. Consequently, it can be shown that posterior for the square-root of the sum of the squares of two of these quantities will be given by a chi distribution with two degrees of freedom. Considering the definitions of Eqs. (67) and (68). It is straightforward to show (numerically or analytically) that this explains the observed factor of ~ 1.3 difference between $h_t^{95\%}$ ($h_v^{95\%}$) and $h_+^{95\%}$ or $h_\times^{95\%}$ ($h_x^{95\%}$ or $h_y^{95\%}$).

-
- [1] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241103 (2016).
- [3] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 221101 (2016).
- [4] C. M. Will, *Theory and Experiment in Gravitational Physics*, revised ed. (Cambridge University Press, Cambridge, England, 1993).
- [5] C. M. Will, *Living Rev. Relativ.* **17**, 4 (2014).
- [6] D. M. Eardley, D. L. Lee, A. P. Lightman, R. V. Wagoner, and C. M. Will, *Phys. Rev. Lett.* **30**, 884 (1973).
- [7] D. Eardley, D. Lee, and A. Lightman, *Phys. Rev. D* **8**, 3308 (1973).
- [8] J. M. Weisberg, D. J. Nice, and J. H. Taylor, *Astrophys. J.* **722**, 1030 (2010).
- [9] P. C. C. Freire, N. Wex, G. Esposito-Farèse, J. P. W. Verbiest, M. Bailes, B. A. Jacoby, M. Kramer, I. H. Stairs, J. Antoniadis, and G. H. Janssen, *Mon. Not. R. Astron. Soc.* **423**, 3328 (2012).
- [10] I. H. Stairs, *Living Rev. Relativ.* **6**, 5 (2003).
- [11] N. Wex, arXiv:1402.5594.
- [12] B. P. Abbott (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **94**, 102002 (2016).
- [13] J. Aasi (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J.* **813**, 39 (2015).
- [14] B. P. Abbott (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **94**, 042002 (2016).
- [15] J. Aasi (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **90**, 062010 (2014).
- [16] J. Aasi (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **93**, 042007 (2016).
- [17] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J.* **839**, 12 (2017).
- [18] K. S. Thorne, in *Three Hundred Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, England, 1987), Chap. 9, p. 330.
- [19] M. Isi, A. J. Weinstein, C. Mead, and M. Pitkin, *Phys. Rev. D* **91**, 082002 (2015).
- [20] K. Chatziioannou, N. Yunes, and N. Cornish, *Phys. Rev. D* **86**, 022004 (2012).
- [21] C. Mead, arXiv:1503.04866.
- [22] A. Nishizawa, A. Taruya, K. Hayama, S. Kawamura, and M.-a. Sakagami, *Phys. Rev. D* **79**, 082002 (2009).
- [23] A. Blaut, *Phys. Rev. D* **85**, 043005 (2012).
- [24] E. Poisson and C. M. Will, *Gravity: Newtonian, Post-Newtonian, Relativistic* (Cambridge University Press, Cambridge, England, 2014).
- [25] D. I. Jones and N. Andersson, *Mon. Not. R. Astron. Soc.* **331**, 203 (2002).
- [26] M. Zimmermann and E. Szedenits, *Phys. Rev. D* **20**, 351 (1979).
- [27] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, *Phys. Rev. D* **58**, 084020 (1998).
- [28] R. Bondarescu, S. A. Teukolsky, and I. Wasserman, *Phys. Rev. D* **79**, 104003 (2009).
- [29] D. I. Jones, *Mon. Not. R. Astron. Soc.* **453**, 53 (2015).
- [30] R. T. Edwards, G. B. Hobbs, and R. N. Manchester, *Mon. Not. R. Astron. Soc.* **372**, 1549 (2006).
- [31] T. M. Niebauer, A. Rüdiger, R. Schilling, L. Schnupp, W. Winkler, and K. Danzmann, *Phys. Rev. D* **47**, 3106 (1993).
- [32] R. Dupuis and G. Woan, *Phys. Rev. D* **72**, 102002 (2005).
- [33] M. Pitkin, M. Isi, J. Veitch, and G. Woan, arXiv:1705.08978.
- [34] J. D. Scargle, *Astrophys. J.* **504**, 405 (1998).
- [35] T. G. F. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, M. Agathos, J. Veitch, K. Grover, T. Sidery, R. Sturani, and A. Vecchio, *Phys. Rev. D* **85**, 082003 (2012).
- [36] D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial*, 2nd ed. (Oxford University Press, New York, 2006), p. 264.
- [37] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms David J.C. MacKay* (Cambridge University Press, Cambridge, England, 2003).

- [38] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Astrophys. J.* **785**, 119 (2014).
- [39] X. Fan, Y. Chen, and C. Messenger, *Phys. Rev. D* **94**, 084029 (2016).
- [40] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, *Phys. Rev. D* **89**, 064023 (2014).
- [41] M. Pitkin, C. Gill, J. Veitch, E. Macdonald, and G. Woan, *J. Phys. Conf. Ser.* **363**, 012041 (2012).
- [42] J. Skilling, *Bayesian Anal.* **1**, 833 (2006).
- [43] J. Veitch *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
- [44] <https://wiki.ligo.org/DASWG/LALSuite>.
- [45] B.P. Abbott (LIGO Scientific Collaboration, and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241102 (2016).
- [46] E. Jaynes, *IEEE Trans. Syst. Sci. Cybern.* **4**, 227 (1968).
- [47] Note that sensitivity estimates presented in [19] were restricted to the specific vector-only model of [21].
- [48] J.D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [49] D. Foreman-Mackey, *J. Open Source Software* **1**, 1 (2016).
- [50] C.R. Keeton, *Mon. Not. R. Astron. Soc.* **414**, 1418 (2011).
- [51] J. Veitch and A. Vecchio, *Phys. Rev. D* **81**, 062003 (2010).