

Towards a generic test of the strong field dynamics of general relativity using compact binary coalescence

T. G. F. Li,¹ W. Del Pozzo,¹ S. Vitale,¹ C. Van Den Broeck,¹ M. Agathos,¹ J. Veitch,^{1,2} K. Grover,³ T. Sidery,³ R. Sturani,^{4,5} and A. Vecchio³

¹*Nikhef-National Institute for Subatomic Physics, Science Park 105, 1098 XG Amsterdam, The Netherlands*

²*School of Physics and Astronomy, Cardiff University, Queen's Buildings, The Parade, Cardiff CF24 3AA, United Kingdom*

³*School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom*

⁴*Dipartimento di Scienze di Base e Fondamenti, Università di Urbino, I-61029 Urbino, Italy*

⁵*INFN, Sezione di Firenze, I-50019 Sesto Fiorentino, Italy*

(Received 6 October 2011; published 26 April 2012)

Coalescences of binary neutron stars and/or black holes are amongst the most likely gravitational-wave signals to be observed in ground-based interferometric detectors. Apart from the astrophysical importance of their detection, they will also provide us with our very first empirical access to the genuine strong-field dynamics of general relativity (GR). We present a new framework based on Bayesian model selection aimed at detecting deviations from GR, subject to the constraints of the Advanced Virgo and LIGO detectors. The method tests the consistency of coefficients appearing in the waveform with the predictions made by GR, without relying on any specific alternative theory of gravity. The framework is suitable for low signal-to-noise ratio events through the construction of multiple subtests, most of which involve only a limited number of coefficients. It also naturally allows for the combination of information from multiple sources to increase one's confidence in GR or a violation thereof. We expect it to be capable of finding a wide range of possible deviations from GR, including ones which, in principle, cannot be accommodated by the model waveforms, on the condition that the induced change in phase at frequencies where the detectors are the most sensitive is comparable to the effect of a few percent change in one or more of the low-order post-Newtonian phase coefficients. In principle, the framework can be used with any GR waveform approximant, with arbitrary parametrized deformations, to serve as model waveforms. In order to illustrate the workings of the method, we perform a range of numerical experiments in which simulated gravitational waves modeled in the restricted post-Newtonian, stationary phase approximation are added to Gaussian and stationary noise that follows the expected Advanced LIGO/Virgo noise curves.

DOI: [10.1103/PhysRevD.85.082003](https://doi.org/10.1103/PhysRevD.85.082003)

PACS numbers: 04.80.Cc

I. INTRODUCTION

General relativity (GR) is a nonlinear, dynamical theory of gravity. Until the 1970s, all of its tests involved the weak-field, stationary regime; these are the standard Solar System tests that are discussed in most textbooks; see, e.g., [1]. GR passed them with impressive accuracy. Nevertheless, the more interesting part of any field theory resides in its dynamics, and this is especially true of GR [2,3]. A first test of the latter came from the Hulse-Taylor binary and a handful of similar tight neutron star binaries [4–6], whose orbital elements are changing in close agreement with GR under the assumption that energy and angular momentum are carried away by gravitational waves (GW). Thus, these discoveries led to the very first, albeit indirect, evidence for GW. However, even the most relativistic of these binaries, PSR J0737-3039 [5,6], is still in the relatively slowly varying, weak-field regime from a GR point of view, with a compactness of $GM/(c^2 R) \simeq 4.4 \times 10^{-6}$, with M the total mass, R the orbital separation, and a typical orbital speed $v/c \simeq 2 \times 10^{-3}$. By contrast, for an inspiraling compact binary, in the limit of a test particle around a nonspinning black hole, the last stable

orbit occurs at a separation of $R = 6 GM/c^2$, where $GM/(c^2 R) = 1/6$ and $v/c = 1/\sqrt{6}$. This constitutes the genuinely strong-field, dynamical regime of general relativity, which in the foreseeable future will only be empirically accessible by means of gravitational-wave detectors.

Several large gravitational-wave observatories have been operational for some years now: the two 4 km arm length LIGO interferometers in the U.S. [7], the 3 km arm length Virgo detector in Italy [8,9], and the 600 m arm length GEO600 [10]. By around 2015, LIGO and Virgo will have been upgraded to their so-called advanced configurations [11–14], and shortly afterwards, up to tens of detections per year are expected [15]. Another planned GW observatory is the Japanese LCGT [16], and the construction of a large interferometer in India is under consideration [17]. Among the most promising sources are the inspiral and merger of compact binaries composed of two neutron stars (BNS), a neutron star and a black hole (NSBH), or two black holes (BBH).

Within GR, the inspiral part of the coalescence process, in particular, has been modeled in great detail using the post-Newtonian (PN) formalism (see [18] and references therein), in which quantities such as the conserved energy

and flux are found as expansions in v/c , where v is a characteristic speed. During inspiral, the GW signals will carry a detailed imprint of the orbital motion. Indeed, the contribution at leading order in amplitude has a phase that is simply $2\Phi(t)$, with $\Phi(t)$ being the orbital phase. Thus, the angular motion of the binary is directly encoded in the waveform's phase, and assuming quasicircular inspiral, the radial motion follows from the instantaneous angular frequency $\omega(t) = \dot{\Phi}(t)$ through the relativistic version of Kepler's third law. If there are deviations from GR, the different emission mechanism and/or differences in the orbital motion will be encoded in the phase of the signal waveform, allowing us to probe the strong-field dynamics of gravity. In this regard we note that, with binary pulsars, one can only constrain the *conservative* sector of the dynamics to 1PN order [i.e. $(v/c)^2$ beyond leading order], and the dissipative sector to leading order; see, e.g., the discussion in [19] and references therein. Hence, when it comes to $\Phi(t)$, these observations do not fully constrain the 1PN contribution. Yet, several of the more interesting dynamical effects occur starting at 1.5PN order. These include “tail effects” [20,21] and spin-orbit interactions; spin-spin effects first appear at 2PN order [22]. As indicated by the Fisher matrix results of [23], Advanced LIGO/Virgo should be able to put significant constraints on the 1.5PN contribution to the phase, and possibly also higher-order contributions.

Possible deviations from GR that have been considered in the past in the context of compact binary coalescence include scalar-tensor theories [24–29]; a varying Newton constant [30]; modified dispersion relation theories, usually referred to in the literature as “massive graviton” models [28,29,31–37]; violations of the no-hair theorem [40–45]; violations of cosmic censorship [45,46]; and parity violating theories [47–50]. The (rather few) *specific* alternative theories of gravity that have been considered in the context of ground-based gravitational-wave detectors—essentially scalar-tensor and massive graviton theories—happen to be hard to constrain much further with GW observations, and we will not consider them in this paper. However, general relativity may be violated in some other manner, including a way that is yet to be envisaged. This makes it imperative to develop methods that can search for *generic* deviations from GR.

Over the past several years, many proposals have been put forward to test GR using compact binary coalescence:

- (1) One can search directly for the imprint of specific alternative theories, such as the so-called massive graviton models and scalar-tensor theories [24–29,31–34,36]. For the massive graviton case, a full Bayesian analysis was recently performed by Del Pozzo, Veitch, and Vecchio [35].
- (2) A method due to Arun *et al.* exploits the fact that, at least for binaries where neither component has spin, all coefficients ψ_i in the PN expansion [see Eq. (2)

below for their definition] of the inspiral phase depend only on the two component masses, m_1 and m_2 [23,51,52]. In that case, only two of the ψ_i are independent, so that a comparison of any three of them allows for a test of GR. Such a method would be very general, in that one does not have to look for any particular way in which gravity might deviate from GR; instead, it allows generic failures of GR to be searched for. However, so far its viability was only explored using Fisher matrix calculations.

- (3) In the so-called parametrized post-Einsteinian (ppE) formalism of Yunes and Pretorius, gravitational waveforms are parametrized so as to include effects from a variety of alternative theories of gravity [53,54]. A Bayesian analysis was performed in [55].
- (4) Recently, a test of the no-hair theorem was presented, also in a Bayesian setting [56].

The first method presupposes that GR will be violated in a particular way. As for the third and fourth points, the particular Bayesian implementation that was used involves comparing the GR waveform with a waveform model that includes parametrized deformations away from GR, thus introducing further free parameters.

Now, imagine one introduces free parameters p_i , $i = 1, 2, \dots, N_T$ in such a way that $p_i = 0$ for all i corresponds to GR being correct. Then one can compare a waveform model in which all the p_i are allowed to vary with the GR waveform model in which all the p_i are zero. As we will explain in this paper, this amounts to asking the following question: Do *all* the p_i differ from zero at the same time? Let us call the associated hypothesis $H_{12\dots N_T}$, which is to be compared with the GR hypothesis $\mathcal{H}_{\text{GR}}$.

A more general question would be as follows: Do *one or more* of the p_i differ from zero, without specifying which? As we shall see, this question is more difficult to cast into the language of model selection. Below we will call the associated hypothesis $\mathcal{H}_{\text{modGR}}$, to be compared with the GR hypothesis $\mathcal{H}_{\text{GR}}$. What we will show is that, although there is no single waveform model that corresponds to $\mathcal{H}_{\text{modGR}}$, testing the latter amounts to testing $2^{N_T} - 1$ hypotheses $H_{i_1 i_2 \dots i_k}$ corresponding to all subsets $\{p_{i_1}, p_{i_2}, \dots, p_{i_k}\}$ of the full set $\{p_1, p_2, \dots, p_{N_T}\}$. Each of the hypotheses $H_{i_1 i_2 \dots i_k}$ is tested by a waveform model in which $p_{i_1}, p_{i_2}, \dots, p_{i_k}$ are free, but all the other p_j are fixed to zero. The Bayes factors against GR for all of these tests can then be combined into a *single* odds ratio which compares the full hypothesis $\mathcal{H}_{\text{modGR}}$ with $\mathcal{H}_{\text{GR}}$.

In a scenario with low signal-to-noise ratio (SNR), as will be the case with advanced ground-based detectors, parameter estimation will degrade significantly when trying to estimate too many parameters at once, and so will model selection if the alternative model to GR has too many additional degrees of freedom [35]. This could be problematic if one only tests the “all-inclusive” hypothesis

$H_{12\dots N_T}$ against GR, i.e., if the question one asks is, “do all the p_i differ from zero at the same time?” The question we want to ask instead, namely, do one or more of the p_i differ from zero, is more general, and most of the subhypotheses $H_{i_1 i_2 \dots i_k}$ involve a smaller number of free parameters, making the corresponding test more powerful in a low-SNR situation. And, as in most Bayesian frameworks, information from multiple sources can be combined in a straightforward way.

Our framework can be used with any family of waveforms with parametrized deformations, including the ppE family. In order to illustrate the method, following [23,51,52] we make the simplest choice by adopting model waveforms in which the deviations away from GR take the form of shifts in a subset of the post-Newtonian inspiral phase coefficients. To establish the validity of the framework, we perform simulations with simple analytic frequency domain waveforms in stationary Gaussian noise that follows the expected Advanced LIGO/Virgo noise curves. In the future, for actual tests of GR, one may want to use time domain waveforms and introduce deviations in e.g. a Hamiltonian used to numerically evolve the motion of the binary.

As noted in [35,53] and also illustrated here, if one is only interested in *detection*, then it might suffice to only search with template waveforms predicted by GR, but parameter estimation can be badly off; this is what is called “fundamental bias.” Indeed, even if there is a deviation from GR in the signal, then a model waveform with completely different values for the masses and other parameters could still be a good fit to the data, with minimal loss of signal-to-noise ratio. We note that, given model waveforms that feature a certain family of deformations away from GR, fundamental bias can still occur if the signal has a deviation that does not belong to the particular class of deformations allowed by the models. However, in that case one expects the non-GR model waveforms to still be preferred over GR ones in the sense of model selection, even if parameter estimation may be deceptive in interpreting the nature of the deviation. We will show an explicit example of this, and will argue that *generic* deviations from GR can be picked up, of course subject to the limitations imposed by the detectors, as is the case with any kind of measurement. In particular, we expect that a GR violation will generally be visible, on the condition that its effect on the phase at frequencies where the detector is the most sensitive is comparable to the effect of a few percent shift in one of the lower-order phase coefficients.

This paper is structured as follows. We first introduce our waveform model for compact binary inspiral and discuss its sensitivity to changes in phase parameters (Sec. II). In Sec. III, we explain the basic method for single and multiple sources. In Sec. IV, we construct different simulated catalogs of sources and evaluate the level at which

deviations from GR can be found. We end with a summary and a discussion of future steps to be taken.

Unless stated otherwise, we will take $G = c = 1$.

II. WAVEFORM MODEL AND ITS SENSITIVITY TO CHANGES IN PHASE COEFFICIENTS

We now introduce our waveform model. Since we are concerned with testing the strong-field dynamics of gravity, eventually all of the effects which we expect to see with compact binary coalescence should be represented in the waveform. This includes, but is not limited to, precession due to spin-orbit and spin-spin interactions [22], subdominant signal harmonics [57], and merger/ringdown [58]. In due time these will indeed need to be taken into account. However, in this paper we first and foremost wish to demonstrate the validity of a particular method, for which it will not be necessary to use very sophisticated waveforms. Also, as suggested by Fisher matrix calculations such as those of Mishra *et al.* [23], methods based on measuring phase coefficients will be the most accurate at low total mass. In this paper we limit ourselves to BNS sources, for which spin will be negligible, as well as subdominant signal harmonics [46,59]. Since we will assume a network of Advanced LIGO and Virgo detectors, the merger and ringdown signals will not have a large impact either [60]. Thus, for a first analysis we will focus on the inspiral part of the coalescence process, modeling the waveform in the frequency domain using the stationary phase approximation (SPA) [61,62]. In particular, we use the so-called restricted TaylorF2 waveforms [63,64] up to 2PN order in phase.

Since the way we illustrate our method here is based on allowing for deviations in phase coefficients, we need to know how sensitive our waveform model is to minor changes in the values of these coefficients. To get a sense of this, one could use the results of [23] as a guide, but since these are based on the Fisher matrix, they necessarily assume that the signal and template are from the same waveform family. Before explaining our method for testing GR, we will first look at what happens both to detectability and parameter estimation when the signal contains a deviation from GR but is being searched for with a bank of GR templates.

A. Model waveform(s) and detector configuration

We start from the way TaylorF2 is implemented in the LIGO Algorithms Library [63]:

$$h(f) = \frac{1}{D} \frac{\mathcal{A}(\theta, \phi, \iota, \psi, \mathcal{M}, \eta)}{\sqrt{\dot{F}(\mathcal{M}, \eta; f)}} f^{2/3} e^{i\Psi(t_c, \phi_c, \mathcal{M}, \eta; f)}, \quad (1)$$

where D is the luminosity distance to the source, (θ, ϕ) specify the sky position, (ι, ψ) give the orientation of the inspiral plane with respect to the line of sight, $\mathcal{M}$ is the chirp mass, and η is the symmetric mass ratio.

In terms of the component masses (m_1, m_2), one has $\eta = m_1 m_2 / (m_1 + m_2)^2$ and $\mathcal{M} = (m_1 + m_2) \eta^{3/5}$. t_c and ϕ_c are the time and phase at coalescence, respectively. The “frequency sweep” $\dot{F}(\mathcal{M}, \eta; f)$ is an expansion in powers of the frequency f with mass-dependent coefficients, and

$$\Psi(t_c, \phi_c, \mathcal{M}, \eta; f) = 2\pi f t_c - \phi_c - \pi/4 + \sum_{i=0}^7 [\psi_i + \psi_i^{(l)} \ln f] f^{(i-5)/3}. \quad (2)$$

In the case of GR, the functional dependence of the coefficients ψ_i and $\psi_i^{(l)}$ on $(\mathcal{M}, \eta)$ can be found in [23]. However, here we will not assume that those relationships necessarily hold, except in the case of ψ_0 , which has been tested using binary pulsars. With minor abuse of notation, let us relabel the remaining coefficients as ψ_i , $i = 1, \dots, M$.

We note that $\dot{F}$ is related to the phase Ψ , and, in principle, we should also leave open the possibility that its expansion coefficients deviate from their GR values. However, with the Advanced LIGO and Virgo network and for stellar mass binaries, we do not expect to be very sensitive to subdominant contributions to the amplitude [46,59].

Let us focus on the phase (2). One way of testing GR would be to use a model waveform in which all the ψ_i are considered to be free parameters, measure these together with $\mathcal{M}$ and η , and check whether one obtains agreement with the functional relations $\psi_i(\mathcal{M}, \eta)$ predicted by GR. However, the events we expect in Advanced LIGO/Virgo will probably not have sufficient SNR for this to be directly feasible [23].

Below, we instead suggest a scheme where a large number of tests are done, in each of which a specific, *limited* subset of the phase coefficients is left free while the others have the dependence on masses as in GR. The results from all of these tests can then be combined into the odds ratio for a general deviation from GR versus GR being correct.

In this study we will assume a network of two Advanced LIGO detectors, one in Hanford, WA, and the other one in Livingston, LA, together with the Advanced Virgo detector in Cascina, Italy. We take the Advanced LIGO noise curve to be the one with zero detuning of the signal recycling mirror and high laser power [65]. With these assumptions, the curves in Fig. 1 represent the incoherent sums of the principal noise sources, as they are currently understood; however, there may be unexpected, additional sources of noise. The high-power, zero-detuning option gives most of the desired sensitivity with the fewest technical difficulties. Advanced Virgo can also be optimized for BNS sources by an appropriate choice of the signal recycling detuning and the signal recycling mirror transmittance [13], and this is what we assume here.




FIG. 1. The high-power, zero-detuning noise curve for Advanced LIGO, and the BNS-optimized Advanced Virgo noise curve.

B. Changes in phase coefficients and detectability

It is important to investigate to what extent signals that may deviate from general relativity could be detected in the first place. The fitting factor (FF) is a measure of the adequateness of a template family to fit the signal; $1 - FF$ is the reduction in signal-to-noise ratio that occurs from using a model waveform which differs from the exact signal waveform when searching the data. Let $h_e(\vec{\lambda})$ be the “exact” waveform of the signal and $h_m(\vec{\theta})$ the model used for detection; the exact and model waveforms are dependent on sets of parameters $\vec{\lambda}$ and $\vec{\theta}$, respectively. The fitting factor is then defined as [66]

$$FF \equiv \max_{\vec{\theta}} \left(\frac{\langle h_e(\vec{\lambda}) | h_m(\vec{\theta}) \rangle}{\sqrt{\langle h_m(\vec{\theta}) | h_m(\vec{\theta}) \rangle \langle h_e(\vec{\lambda}) | h_e(\vec{\lambda}) \rangle}} \right), \quad (3)$$

where $\langle a | b \rangle$ denotes the usual noise weighted inner product,

$$\langle a | b \rangle = 2 \int_{f_{\min}}^{f_{\max}} \frac{a^*(f)b(f) + a(f)b^*(f)}{S_n(f)} df, \quad (4)$$

and $S_n(f)$ is the one-sided noise spectral density [19]. $f_{\min}$ is the detectors’ lower cutoff frequency, and in our case, $f_{\max}$ is the frequency at the last stable orbit, $f_{\max} = (6^{3/2} \pi \mathcal{M} \eta^{-3/5})^{-1}$. We note that the detection rate scales like the cube of the signal-to-noise ratio, so that the fractional reduction in the event rate is $1 - FF^3$ [66]. In this case the waveform is not “exact” in the sense of numerical relativity; instead, we use the fitting factor as a measure of how similar a modified TaylorF2 waveform and a GR version are.

A sample deviation is tested using a modified waveform, which is different only in the 1.5PN order phase coefficient: $\psi_3^{\text{GR}}(\mathcal{M}, \eta) \rightarrow \psi_3^{\text{GR}}(\mathcal{M}, \eta)[1 + \delta\chi_3]$. The Advanced LIGO noise curve is used. The signal is “detected” with a template bank of standard GR waveforms that is regularly spaced in the parameters present in the phase, ϕ_c , t_c , $\mathcal{M}$, and η . We study a $(1.4, 1.4)M_\odot$




FIG. 2. The fitting factors for a range of $\mathcal{M}$, once the other parameters are maximized over. Here a deviation of ψ_3 from 2.5% to 17.5% is used for a $(1.4, 1.4)M_\odot$ system. The vertical dashed line represents the true value of $\mathcal{M}$, while the maximum is offset to compensate for the modification in the phase.




FIG. 3. The fitting factors for a range of η , once the other parameters are maximized over. Here a deviation of ψ_3 from 2.5% to 17.5% is used for a $(1.4, 1.4)M_\odot$ system. The vertical dashed line represents the true value of η , while the maximum is offset to compensate for the modification in the phase.

binary with $\delta\chi_3$ ranging from 0.025 to 0.175. As seen in Figs. 2 and 3, the mass parameters can absorb the change in the phase due to the modified phase coefficient while providing a fitting factor of over 95%.

Thus, with a template bank of GR waveforms, it is possible to detect a signal containing a large deviation from GR without significant loss in the signal-to-noise ratio, but recovered with intrinsic parameters that deviate significantly from the true values. We now describe a method which will be able to, nevertheless, recognize a deviation from GR when one is present.

III. METHOD

We will first recall some basic facts about Bayesian inference. Next we outline our method for finding a possible violation of GR using inspiral signals by allowing for deviations in phase coefficients. For simplicity, we start with an example where only two phase coefficients are

taken into account; then we go on to the general case. Finally, we explain how to combine information from a catalog of sources.

A. Bayesian inference and nested sampling

We now give a brief overview of Bayesian inference, as well as the method of nested sampling due to Skilling [67], which was first introduced into the ground-based GW data analysis by Veitch and Vecchio [68–70] and which we will adopt here as well.

Let us consider hypotheses $\mathcal{H}_i, \mathcal{H}_j$. Here $\mathcal{H}_i$ could be the hypothesis that there is a deviation from GR, while $\mathcal{H}_j$ is the hypothesis that GR is correct; or $\mathcal{H}_i$ could simply be the hypothesis that a signal of a particular form is present in the data, while $\mathcal{H}_j$ is the hypothesis that there is only noise. The statements that we can make about any hypothesis are based on a data set d (observations) and all the relevant prior information “I” that we hold.

Within the framework of Bayesian inference, the key quantity that one needs to compute is the posterior probability of a hypothesis $\mathcal{H}_i$. Applying Bayes’ theorem we obtain

$$P(\mathcal{H}_i|d, \text{I}) = \frac{P(\mathcal{H}_i|\text{I})P(d|\mathcal{H}_i, \text{I})}{P(d|\text{I})}, \quad (5)$$

where $P(\mathcal{H}_i|d, \text{I})$ is the posterior probability of the hypothesis $\mathcal{H}_i$ given the data, $P(\mathcal{H}_i|\text{I})$ is the prior probability of the hypothesis, and $P(d|\mathcal{H}_i, \text{I})$ is the marginal likelihood or evidence for $\mathcal{H}_i$, which can be written as

$$P(d|\mathcal{H}_i, \text{I}) = \mathcal{L}(\mathcal{H}_i) = \int d\vec{\theta} p(\vec{\theta}|\mathcal{H}_i, \text{I}) p(d|\vec{\theta}, \mathcal{H}_i, \text{I}). \quad (6)$$

In the above expression, $p(\vec{\theta}|\mathcal{H}_i, \text{I})$ is the prior probability density of the unknown parameter vector $\vec{\theta}$ within the model corresponding to $\mathcal{H}_i$, and $p(d|\vec{\theta}, \mathcal{H}_i, \text{I})$ is the likelihood function of the observation d , assuming a given value of the parameters $\vec{\theta}$ and the model $\mathcal{H}_i$.

If we want to compare different hypotheses, $\mathcal{H}_i$ and $\mathcal{H}_j$, in light of the observations made, we can compute the ratio of posterior probabilities, which is known as the odds ratio:

$$O_j^i = \frac{P(\mathcal{H}_i|d, \text{I})}{P(\mathcal{H}_j|d, \text{I})} = \frac{P(\mathcal{H}_i|\text{I})}{P(\mathcal{H}_j|\text{I})} \frac{P(d|\mathcal{H}_i, \text{I})}{P(d|\mathcal{H}_j, \text{I})} = \frac{P(\mathcal{H}_i|\text{I})}{P(\mathcal{H}_j|\text{I})} B_j^i, \quad (7)$$

where $P(\mathcal{H}_j|\text{I})/P(\mathcal{H}_i|\text{I})$ is the prior odds of the two hypotheses, the relative confidence we assign to the models before any observation, and B_j^i is the Bayes factor.

In addition to computing the relative probabilities of different models, one usually wants to make inference on the unknown parameters, and therefore one needs to compute the joint posterior probability density function (PDF)

$$p(\vec{\theta}|d, \mathcal{H}_i, \mathcal{I}) = \frac{p(\vec{\theta}|\mathcal{H}_i, \mathcal{I})p(d|\vec{\theta}, \mathcal{H}_i, \mathcal{I})}{p(d|\mathcal{H}_i, \mathcal{I})}. \quad (8)$$

From the previous expression it is simple to compute the marginalized PDF on any given parameter, say θ_1 within a given model $\mathcal{H}_i$:

$$p(\theta_1|d, \mathcal{H}_i, \mathcal{I}) = \int d\theta_2 \dots \int d\theta_N p(\vec{\theta}|d, \mathcal{H}_i, \mathcal{I}). \quad (9)$$

The key quantities for Bayesian inference in Eqs. (7)–(9) can be efficiently computed using e.g. a nested sampling algorithm [67]. The basic idea of nested sampling is to use a collection of n objects, called live points, randomly sampled from the prior distributions, but subject to a constraint over the value of their likelihood. A live point $\vec{\xi}$ is a point in the multidimensional parameter space. At each iteration, the live point $\vec{\xi}^*$ having the lowest likelihood $\mathcal{L}(\vec{\xi}^*)$ is replaced with a new point $\vec{\xi}$ sampled from the prior distribution. To be accepted, the new point must obey the condition

$$\mathcal{L}(\vec{\xi}) > \mathcal{L}(\vec{\xi}^*). \quad (10)$$

The above condition ensures that regions of progressively increasing likelihood are explored, and the evidence integral, Eq. (6), is calculated using those points as the computation progresses.

In this paper we use a specific implementation of this technique that was developed for ground-based observations of coalescing binaries by Veitch and Vecchio; we point the interested reader to [68–72] for technical details. To select a new live point, a Metropolis-Hastings Markov chain Monte Carlo method is used with p steps. The uncertainty in the evidence computation for a given number of live points n and Markov chain Monte Carlo steps p was quantified in [70]. For the calculations in this paper we took $n = 1000$ and $p = 100$, in which case the standard deviation on log Bayes factors is $\mathcal{O}(1)$.

B. Basic method for a single source

Given that we have no knowledge of which coefficient (s) might deviate from the GR values, we want to test the hypothesis that at least one of the known coefficients $\{\psi_0, \psi_1, \psi_2, \dots, \psi_M\}$ is different. Computational limitations and possible lack of sensitivity to changes in the higher-order coefficients will induce us to only look for deviations in a set of testing coefficients $\{\psi_1, \dots, \psi_{N_T}\} \subset \{\psi_0, \psi_1, \dots, \psi_M\}$, with $N_T < M$. For the examples of Sec. IV, we will choose $N_T = 3$ due to computational constraints, but a larger number could be used.

Let us introduce some notation. We define hypotheses $H_{i_1 i_2 \dots i_k}$ as follows:

$H_{i_1 i_2 \dots i_k}$ is the hypothesis that the phasing coefficients $\psi_{i_1}, \dots, \psi_{i_k}$ do not have the functional dependence on $(\mathcal{M}, \eta)$ as predicted by general relativity, but all other

coefficients $\psi_j, j \notin \{i_1, i_2, \dots, i_k\}$ do have the dependence as in GR.

Thus, for example, H_{12} is the hypothesis that ψ_1 and ψ_2 deviate from their GR values, with all other coefficients being as in GR. With each of the hypotheses above, we can associate a waveform model that can be used to test it. Let $\vec{\theta} = \{\mathcal{M}, \eta, \dots\}$ be the parameters occurring in the GR waveform. Then $H_{i_1 i_2 \dots i_k}$ is tested by a waveform in which the independent parameters are

$$\{\vec{\theta}, \psi_{i_1}, \psi_{i_2}, \dots, \psi_{i_k}\}; \quad (11)$$

i.e. the coefficients $\{\psi_{i_1}, \psi_{i_2}, \dots, \psi_{i_k}\}$ are allowed to vary freely. In practice, these will need to be subsets of a limited set of coefficients; in Sec. IV, where we present results, we will consider all subsets of the set $\{\psi_1, \psi_2, \psi_3\}$. This choice will already suffice to illustrate the method without being overly computationally costly, but in the future one may want to use a larger set.

Now, the hypothesis we would really like to test is that one or more of the ψ_i differ from GR, without specifying which. This corresponds to a logical “or” of the above hypotheses:

$$\mathcal{H}_{\text{modGR}} = \bigvee_{i_1 < i_2 < \dots < i_k} H_{i_1 i_2 \dots i_k}. \quad (12)$$

Our aim is to compute the following odds ratio:

$$O_{\text{GR}}^{\text{modGR}} \equiv \frac{P(\mathcal{H}_{\text{modGR}}|d, \mathcal{I})}{P(\mathcal{H}_{\text{GR}}|d, \mathcal{I})}. \quad (13)$$

However, the hypothesis (12) is not what model waveforms with one or more free coefficients ψ_i will test; rather, such waveforms test the hypotheses $H_{i_1 i_2 \dots i_k}$ themselves. What we will need to do is break up the logical “or” in $\mathcal{H}_{\text{modGR}}$ into the component hypotheses $H_{i_1 i_2 \dots i_k}$. Fortunately, this is trivial.

Before treating the problem more generally, let us consider a simple example. Imagine that only two coefficients ψ_1 and ψ_2 are being used for the testing of GR. Then

$$\mathcal{H}_{\text{modGR}} = H_1 \vee H_2 \vee H_{12}. \quad (14)$$

In this example, the odds ratio of interest can then be written as

$$^{(2)}O_{\text{GR}}^{\text{modGR}} \equiv \frac{P(H_1 \vee H_2 \vee H_{12}|d, \mathcal{I})}{P(\mathcal{H}_{\text{GR}}|d, \mathcal{I})}, \quad (15)$$

where the superscript (2) reminds us that only two of the parameters are being used for testing.

An important observation is that the hypotheses H_1, H_2, H_{12} are logically disjoint: the “and” of any two of them is always false. Indeed, in H_1 , ψ_2 takes the value predicted by GR, but in H_2 , it differs from the GR value, as it does in H_{12} . Similarly, in H_2 , ψ_1 takes the GR value, but in H_1 , it differs from the GR value, and the same in H_{12} . More generally, any two hypotheses $H_{i_1 i_2 \dots i_k}$ and $H_{j_1 j_2 \dots j_l}$ with

$\{i_1, i_2, \dots, i_k\} \neq \{j_1, j_2, \dots, j_l\}$ are logically disjoint. This means that the odds ratio is simply

$$^{(2)}O_{\text{GR}}^{\text{modGR}} = \frac{P(H_1|d, \mathbf{I})}{P(\mathcal{H}_{\text{GR}}|d, \mathbf{I})} + \frac{P(H_2|d, \mathbf{I})}{P(\mathcal{H}_{\text{GR}}|d, \mathbf{I})} + \frac{P(H_{12}|d, \mathbf{I})}{P(\mathcal{H}_{\text{GR}}|d, \mathbf{I})}. \quad (16)$$

Using Bayes' theorem, this can be written as

$$^{(2)}O_{\text{GR}}^{\text{modGR}} = \frac{P(H_1|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} B_{\text{GR}}^1 + \frac{P(H_2|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} B_{\text{GR}}^2 + \frac{P(H_{12}|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} B_{\text{GR}}^{12}. \quad (17)$$

Here B_{GR}^1 , B_{GR}^2 , B_{GR}^{12} are the Bayes factors

$$B_{\text{GR}}^1 = \frac{P(d|H_1, \mathbf{I})}{P(d|\mathcal{H}_{\text{GR}}, \mathbf{I})}, \quad B_{\text{GR}}^2 = \frac{P(d|H_2, \mathbf{I})}{P(d|\mathcal{H}_{\text{GR}}, \mathbf{I})}, \quad (18)$$

$$B_{\text{GR}}^{12} = \frac{P(d|H_{12}, \mathbf{I})}{P(d|\mathcal{H}_{\text{GR}}, \mathbf{I})},$$

and $P(H_1|\mathbf{I})/P(\mathcal{H}_{\text{GR}}|\mathbf{I})$, $P(H_2|\mathbf{I})/P(\mathcal{H}_{\text{GR}}|\mathbf{I})$, $P(H_{12}|\mathbf{I})/P(\mathcal{H}_{\text{GR}}|\mathbf{I})$ are ratios of prior odds.

In practice, we will write the testing coefficients as

$$\psi_i = \psi_i^{\text{GR}}(\mathcal{M}, \eta)[1 + \delta\chi_i], \quad (19)$$

with $\psi_i^{\text{GR}}(\mathcal{M}, \eta)$ the functional form of the dependence of ψ_i on $(\mathcal{M}, \eta)$ according to GR, and the dimensionless $\delta\chi_i$ is a fractional shift in ψ_i . Note that in GR, the 0.5PN contribution is identically zero; it will be treated separately, as explained in Sec. IV. In the example of this subsection, one can assume that ψ_1 , ψ_2 are any of the PN coefficients *other* than the 0.5PN one. With the above notation, the Bayes factors (18) are

$$B_{\text{GR}}^1 = \frac{\int d\vec{\theta} d\delta\chi_1^{\{1\}} \pi(\delta\chi_1) \pi(\vec{\theta}) p(d|\vec{\theta}, \delta\chi_1, H_1, \mathbf{I})}{\int d\vec{\theta} \pi(\vec{\theta}) p(d|\vec{\theta}, \mathcal{H}_{\text{GR}}, \mathbf{I})}, \quad (20)$$

$$B_{\text{GR}}^2 = \frac{\int d\vec{\theta} d\delta\chi_2^{\{2\}} \pi(\delta\chi_2) \pi(\vec{\theta}) p(d|\vec{\theta}, \delta\chi_2, H_2, \mathbf{I})}{\int d\vec{\theta} \pi(\vec{\theta}) p(d|\vec{\theta}, \mathcal{H}_{\text{GR}}, \mathbf{I})}, \quad (21)$$

$$B_{\text{GR}}^{12} = \frac{\int d\vec{\theta} d\delta\chi_1 d\delta\chi_2^{\{12\}} \pi(\delta\chi_1, \delta\chi_2) \pi(\vec{\theta}) p(d|\vec{\theta}, \delta\chi_1, \delta\chi_2, H_{12}, \mathbf{I})}{\int d\vec{\theta} \pi(\vec{\theta}) p(d|\vec{\theta}, \mathcal{H}_{\text{GR}}, \mathbf{I})}. \quad (22)$$

Here $\{1\}\pi(\delta\chi_1)$, $\{2\}\pi(\delta\chi_2)$, $\{12\}\pi(\delta\chi_1, \delta\chi_2)$ are priors for, respectively, $\delta\chi_1$, $\delta\chi_2$, and the pair $(\delta\chi_1, \delta\chi_2)$. We choose these to be constant functions in the relevant parameter or pair of parameters, having support within a large interval or square centered on the origin, and normalized to 1. For the other parameters, $\vec{\theta}$, we use the same functional form and limits as [70], with the exception of the distance being allowed to vary between 1 and 1000 Mpc. Specifically, for the sky location and the orientation of the orbital plane, we choose uniform priors on the corresponding unit spheres. For the phase at coalescence ϕ_c , we choose a flat prior with $\phi_c \in [0, 2\pi]$, and the time of coalescence t_c is in a time interval of 100 ms. The prior on η is flat on the interval $[0, 0.25]$. For the chirp mass we use an approximation to the Jeffreys prior which gives $p(\mathcal{M}|\mathbf{I}) \propto \mathcal{M}^{-11/6}$; see [70] for motivation. In addition, component masses are restricted to the interval $m_1, m_2 \in [1, 34]M_\odot$.

At this point it is worth commenting on the mutual relationships of the hypotheses $H_{i_1 i_2 \dots i_k}$ amongst each other, and with $\mathcal{H}_{\text{GR}}$, and on the waveform models used to test them. As an example, let us discuss the relationship between H_1 and $\mathcal{H}_{\text{GR}}$. Consider the numerator of the Bayes factor B_{GR}^1 in Eq. (20):

$$\int d\vec{\theta} d\delta\chi_1^{\{1\}} \pi(\delta\chi_1) \pi(\vec{\theta}) p(d|\vec{\theta}, \delta\chi_1, H_1, \mathbf{I}). \quad (23)$$

The parameter space of the GR waveforms, $\{\vec{\theta}\}$, has a natural embedding into the parameter space $\{\vec{\theta}, \delta\chi_1\}$ of

the waveforms used to test H_1 : it can be identified with the hypersurface $\delta\chi_1 = 0$. We could have explicitly excluded this hypersurface from $\{\vec{\theta}, \delta\chi_1\}$ by setting a prior on $\delta\chi_1$ of the form $\{1\}\pi_0(\delta\chi_1) = 0$ if $\delta\chi_1 = 0$, and $\{1\}\pi_0(\delta\chi_1) = \text{const}$ otherwise. However, this would not have made a difference in the integral above; indeed, with respect to the integration measure induced by the prior probability density on $\{\vec{\theta}, \delta\chi_1\}$, the surface $\delta\chi_1 = 0$ constitutes a set of measure zero. Now look at the denominator in the expression for B_{GR}^1 , which is the evidence for the GR hypothesis:

$$\int d\vec{\theta} \pi(\vec{\theta}) p(d|\vec{\theta}, \mathcal{H}_{\text{GR}}, \mathbf{I}). \quad (24)$$

Despite the fact that the GR waveforms form a set of measure zero within the set of waveforms used for testing H_1 , the above integral is clearly not zero. It is the evidence for a qualitatively different hypothesis, whose parameter space $\{\vec{\theta}\}$ carries a different integration measure with respect to which the marginalization of the likelihood is carried out. In particular, $\mathcal{H}_{\text{GR}}$ is not “included” in the H_1 model in any sense that is meaningful for model selection.

Let us consider, in particular, the most inclusive hypothesis, in the present example H_{12} . The numerator of the Bayes factor B_{GR}^{12} in Eq. (22) is

$$\int d\vec{\theta} d\delta\chi_1 d\delta\chi_2^{\{12\}} \pi(\delta\chi_1, \delta\chi_2) \quad (25)$$

$$\times \pi(\vec{\theta}) p(d|\vec{\theta}, \delta\chi_1, \delta\chi_2, H_{12}, \mathcal{I}).$$

The parameter space of the GR waveforms has a natural embedding into the parameter space $\{\vec{\theta}, \delta\chi_1, \delta\chi_2\}$ of the waveforms used to test H_{12} , by identification with the hypersurface $\delta\chi_1 = \delta\chi_2 = 0$. Similarly, the parameter spaces of the waveforms that test H_1 and H_2 can be identified with the hypersurface $\delta\chi_2 = 0$ and the hypersurface $\delta\chi_1 = 0$, respectively. Thus, to make sure that the waveform models with, respectively, free parameters $\{\vec{\theta}, \delta\chi_1\}$, $\{\vec{\theta}, \delta\chi_2\}$, and $\{\vec{\theta}, \delta\chi_1, \delta\chi_2\}$ really correspond to, respectively, the disjoint hypotheses H_1 , H_2 , and H_{12} , one may wish to choose a prior $^{\{12\}}\pi_0(\delta\chi_1, \delta\chi_2)$ such that $^{\{12\}}\pi_0(\delta\chi_1, \delta\chi_2) = 0$ if $\delta\chi_1 = 0$ or $\delta\chi_2 = 0$ or both, and $^{\{12\}}\pi_0(\delta\chi_1, \delta\chi_2) = \text{const}$ otherwise. However, this does not make any difference for the integral above, and hence has no bearing on evidence and Bayes factors. On the other hand, despite the fact that the parameter spaces $\{\vec{\theta}, \delta\chi_2\}$ and $\{\vec{\theta}, \delta\chi_1\}$ constitute sets of measure zero within $\{\vec{\theta}, \delta\chi_1, \delta\chi_2\}$, the numerators of B_{GR}^1 and B_{GR}^2 in Eqs. (20) and (21) are not zero. With our choices of prior probability densities on the different parameter spaces, and the associated waveform models, there is no meaningful sense in which H_1 or H_2 is “included” in H_{12} . Thus, we are indeed testing the disjoint hypotheses H_1 , H_2 , and H_{12} . The latter is the hypothesis that *both* $\delta\chi_1$ and $\delta\chi_2$ differ from zero: the prior density $^{\{12\}}\pi_0(\delta\chi_1, \delta\chi_2)$ gives the surfaces $\delta\chi_1 = 0$ and $\delta\chi_2 = 0$ zero prior mass.

The waveform model that leaves $\{\vec{\theta}, \delta\chi_1, \delta\chi_2\}$ free corresponds to the hypothesis H_{12} . When computing the Bayes factor B_{GR}^{12} , the question one addresses is, do ψ_1 and ψ_2 both differ from what GR predicts? This is analogous to what has been done in recent Bayesian work on testing GR [35,55,56]. As we are in the process of showing, it is possible to address a more general question, namely, do ψ_1 , or ψ_2 , or both at the same time, differ from their GR values?

For completeness, we note that what nested sampling will give us directly are not the Bayes factors of Eqs. (20)–(22), but rather the Bayes factors for the various hypotheses against the noise-only hypothesis $\mathcal{H}_{\text{noise}}$:

$$B_{\text{noise}}^1 = \frac{P(d|H_1, \mathcal{I})}{P(d|\mathcal{H}_{\text{noise}}, \mathcal{I})}, \quad B_{\text{noise}}^2 = \frac{P(d|H_2, \mathcal{I})}{P(d|\mathcal{H}_{\text{noise}}, \mathcal{I})},$$

$$B_{\text{noise}}^{12} = \frac{P(d|H_{12}, \mathcal{I})}{P(d|\mathcal{H}_{\text{noise}}, \mathcal{I})}, \quad B_{\text{noise}}^{\text{GR}} = \frac{P(d|\mathcal{H}_{\text{GR}}, \mathcal{I})}{P(d|\mathcal{H}_{\text{noise}}, \mathcal{I})}. \quad (26)$$

These can trivially be combined to obtain the Bayes factors of Eqs. (20)–(22):

$$B_{\text{GR}}^1 = \frac{B_{\text{noise}}^1}{B_{\text{noise}}^{\text{GR}}}, \quad B_{\text{GR}}^2 = \frac{B_{\text{noise}}^2}{B_{\text{noise}}^{\text{GR}}}, \quad B_{\text{GR}}^{12} = \frac{B_{\text{noise}}^{12}}{B_{\text{noise}}^{\text{GR}}}. \quad (27)$$

Now, upon calculating the Bayes factors for each model, we would like to combine these measurements into an overall odds ratio between the GR model and *any* of the competing hypotheses [Eq. (17)]. In order to do this, we must specify the prior odds for each model against GR, $P(H_1|\mathcal{I})/P(\mathcal{H}_{\text{GR}}|\mathcal{I})$, $P(H_2|\mathcal{I})/P(\mathcal{H}_{\text{GR}}|\mathcal{I})$, etc. Here one might want to let oneself be guided by, e.g., the expectation that a violation of GR will likely occur at higher post-Newtonian order, and give more weight to H_2 and H_{12} . Or, if one expects a deviation to happen only in a particular phase coefficient (such as ψ_2 in the case of “massive gravity”), one may want to downweight the most inclusive hypothesis, in this example H_{12} . In reality, we will not know beforehand what form a violation will take; in particular, it could affect all of the PN coefficients. For the purposes of this analysis, we invoke the principle of indifference among the alternative hypotheses, taking none of them to be preferable to any other. This imposes the condition that the prior odds of each against GR are equal. We explicitly note that this is a choice of the authors for computing final results. This choice results in our effectively taking the average of the Bayes factors for the alternative hypotheses when we compute the odds ratio versus GR.

When combining the Bayes factors into the odds ratio, we therefore assume

$$\frac{P(H_1|\mathcal{I})}{P(\mathcal{H}_{\text{GR}}|\mathcal{I})} = \frac{P(H_2|\mathcal{I})}{P(\mathcal{H}_{\text{GR}}|\mathcal{I})} = \frac{P(H_{12}|\mathcal{I})}{P(\mathcal{H}_{\text{GR}}|\mathcal{I})}. \quad (28)$$

Furthermore, we let

$$\frac{P(\mathcal{H}_{\text{modGR}}|\mathcal{I})}{P(\mathcal{H}_{\text{GR}}|\mathcal{I})} = \frac{P(H_1 \vee H_2 \vee H_{12}|\mathcal{I})}{P(\mathcal{H}_{\text{GR}}|\mathcal{I})} = \alpha, \quad (29)$$

where we do not specify α ; it will end up being an overall scaling of the odds ratio. This, together with (28) and the logical disjointness of the hypotheses H_1 , H_2 , H_{12} , implies

$$\frac{P(H_1|\mathcal{I})}{P(\mathcal{H}_{\text{GR}}|\mathcal{I})} = \frac{P(H_2|\mathcal{I})}{P(\mathcal{H}_{\text{GR}}|\mathcal{I})} = \frac{P(H_{12}|\mathcal{I})}{P(\mathcal{H}_{\text{GR}}|\mathcal{I})} = \frac{\alpha}{3}. \quad (30)$$

The final expression for the odds ratio for a modification of GR versus GR, in the case where up to two coefficients are used for testing, is then

$$^{(2)}O_{\text{GR}}^{\text{modGR}} = \frac{\alpha}{3} [B_{\text{GR}}^1 + B_{\text{GR}}^2 + B_{\text{GR}}^{12}]. \quad (31)$$

Before continuing to the case of more than two testing parameters, let us compare and contrast what is proposed here with what was done in previous Bayesian work, e.g. [35,55,56]. There, one introduced free parameters p_i in the waveform (not necessarily corresponding to shifts in phase coefficients; they could, e.g., be shifts in ringdown frequencies and damping times), which are zero in GR. Next, one constructed a Bayes factor comparing a model waveform in which all of the p_i are allowed to vary freely with the GR model in which all of the p_i are fixed to zero. In our language and in the case of two testing parameters,

this corresponds to only comparing the hypothesis H_{12} with the GR hypothesis $\mathcal{H}_{\text{GR}}$. Hence, what was effectively done in previous work was to address the question, do *all* of the additional free parameters differ from zero at the same time? A more interesting question to ask is, do *one or more* of the extra parameters differ from zero? This corresponds to testing our hypothesis $\mathcal{H}_{\text{modGR}}$. There is no waveform model that can be used to test the latter hypothesis directly, but as we were able to show, $\mathcal{H}_{\text{modGR}}$ can be broken up into subhypotheses, H_1 , H_2 , and H_{12} , in the case of two testing parameters. Each of these *can* be associated to a waveform model; hence they can be tested against the GR hypothesis $\mathcal{H}_{\text{GR}}$. The resulting Bayes factors can be combined into an odds ratio as in Eq. (31), which *does* compare the more general hypothesis $\mathcal{H}_{\text{modGR}}$ with $\mathcal{H}_{\text{GR}}$.

C. The general case

So far we have assumed just two testing coefficients, but we may want to use more. In practice, it makes sense to pick $\{\psi_1, \dots, \psi_{N_T}\}$, $N_T \leq M$. The number of coefficients used, N_T , will be dictated mostly by computational cost; in Sec. IV we will pick $N_T = 3$, but a larger number could be chosen. We then define

$$\mathcal{H}_{\text{modGR}} = \bigvee_{i_1 < i_2 < \dots < i_k; k \leq N_T} H_{i_1 i_2 \dots i_k}. \quad (32)$$

When using this set of testing coefficients, the odds ratio for “modification to GR” versus GR becomes

$$\begin{aligned} (N_T)O_{\text{GR}}^{\text{modGR}} &= \frac{P(\mathcal{H}_{\text{modGR}}|d, \mathbf{I})}{P(\mathcal{H}_{\text{GR}}|d, \mathbf{I})} \\ &= \frac{P(\bigvee_{i_1 < i_2 < \dots < i_k; k \leq N_T} H_{i_1 i_2 \dots i_k} | d, \mathbf{I})}{P(\mathcal{H}_{\text{GR}}|d, \mathbf{I})}, \end{aligned} \quad (33)$$

where, as before, $H_{i_1 i_2 \dots i_k}$ is the hypothesis that $\{\psi_{i_1}, \psi_{i_2}, \dots, \psi_{i_k}\}$ do not have the functional dependence on $(\mathcal{M}, \eta)$ as predicted by GR, but all of the remaining coefficients do. Thus, we are considering the odds ratio for one or more of the phase coefficients $\psi_1, \dots, \psi_{N_T}$ deviating from GR, versus all of them having the functional dependence on masses as in GR.

Using the logical disjointness of the $H_{i_1 i_2 \dots i_k}$ for different subsets $\{i_1, i_2, \dots, i_k\}$ as well as Bayes’ theorem, one can write

$$(N_T)O_{\text{GR}}^{\text{modGR}} = \sum_{k=1}^{N_T} \sum_{i_1 < i_2 < \dots < i_k} \frac{P(H_{i_1 i_2 \dots i_k}|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} B_{\text{GR}}^{i_1 i_2 \dots i_k}, \quad (34)$$

where

$$B_{\text{GR}}^{i_1 i_2 \dots i_k} = \frac{P(d|H_{i_1 i_2 \dots i_k}, \mathbf{I})}{P(d|\mathcal{H}_{\text{GR}}, \mathbf{I})}. \quad (35)$$

Again, one computes the $2^{N_T} - 1$ individual Bayes factors $B_{\text{GR}}^{i_1 i_2 \dots i_k}$ of each of the alternative hypotheses versus

GR. The evaluation of the odds ratio requires that we use specific values for the prior odds ratios. We will set them equal to each other, as we did in Eq. (28):

$$\frac{P(H_{i_1 i_2 \dots i_k}|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} = \frac{P(H_{j_1 j_2 \dots j_l}|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} \quad \text{for any } k, l \leq N_T, \quad (36)$$

in which case the odds ratio $(N_T)O_{\text{GR}}^{\text{modGR}}$ will be proportional to a straightforward average of the Bayes factors. We note once again that other choices could, in principle, be made. If one expects a violation to be mainly visible in a *particular* phase coefficient (see, e.g., Table 1 in [55]), then one might want to downweight the more inclusive hypotheses. Or, one may argue that a deviation will most likely affect *all* of the coefficients, starting from some PN order, but then it would not necessarily be sensible to give more inclusive hypotheses a lower weight. Hence, we assign equal prior odds to all of the subhypotheses $H_{i_1 i_2 \dots i_k}$.

Also, as before, we let

$$\frac{P(\mathcal{H}_{\text{modGR}}|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} = \alpha, \quad (37)$$

where we do not specify α ; it will end up being an overall prefactor in the odds ratio. The equality (37), together with (36) and the logical disjointness of the $2^{N_T} - 1$ hypotheses $H_{i_1 \dots i_k}$, implies

$$\frac{P(H_{i_1 i_2 \dots i_k}|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} = \frac{\alpha}{2^{N_T} - 1}. \quad (38)$$

In terms of the $H_{i_1 i_2 \dots i_k}$, the odds ratio can then be written as

$$(N_T)O_{\text{GR}}^{\text{modGR}} = \frac{\alpha}{2^{N_T} - 1} \sum_{k=1}^{N_T} \sum_{i_1 < i_2 < \dots < i_k} B_{\text{GR}}^{i_1 i_2 \dots i_k}. \quad (39)$$

Analogously to (20)–(22), we will use priors on the parameters $\delta\chi_i$,

$$\{i_1 i_2 \dots i_k\} \pi(\delta\chi_{i_1}, \delta\chi_{i_2}, \dots, \delta\chi_{i_k}), \quad (40)$$

which are constant within some large box centered on the origin.

D. Combining information from multiple sources

Although the detection rate for binary neutron stars is still rather uncertain, we expect advanced instruments to detect several events per year [15]. It is therefore important to take advantage of multiple detections to provide tighter constraints on the validity of GR. Consider a set of $\mathcal{N}$ independent GW events, corresponding to $\mathcal{N}$ independent data sets d_A . We do not assume that deviations from GR are necessarily consistent between events, but rather that they can vary from source to source. We assume there is a common underlying theory of gravity that describes emission of gravitational waves from the sources that are observed, but that shifts in the values of the parameters can vary from one source to another, over and above the

dependence of the parameters on the masses. One can write down a combined odds ratio for the catalog of sources:

$$\begin{aligned} (N_T)O_{\text{GR}}^{\text{modGR}} &= \frac{P(\mathcal{H}_{\text{modGR}}|d_1, \dots, d_{\mathcal{N}}, \mathbf{I})}{P(\mathcal{H}_{\text{GR}}|d_1, \dots, d_{\mathcal{N}}, \mathbf{I})} \\ &= \frac{\sum_{k=1}^{N_T} \sum_{i_1 < i_2 < \dots < i_k} P(H_{i_1 i_2 \dots i_k}|d_1, \dots, d_{\mathcal{N}}, \mathbf{I})}{P(\mathcal{H}_{\text{GR}}|d_1, \dots, d_{\mathcal{N}}, \mathbf{I})} \\ &= \sum_{k=1}^{N_T} \sum_{i_1 < i_2 < \dots < i_k} \frac{P(H_{i_1 i_2 \dots i_k}|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} B_{\text{GR}}^{i_1 i_2 \dots i_k}, \end{aligned} \quad (41)$$

where

$$(\text{cat})B_{\text{GR}}^{i_1 i_2 \dots i_k} = \frac{P(d_1, \dots, d_{\mathcal{N}}|H_{i_1 i_2 \dots i_k}, \mathbf{I})}{P(d_1, \dots, d_{\mathcal{N}}|\mathcal{H}_{\text{GR}}, \mathbf{I})}. \quad (42)$$

Since the events $d_1, \dots, d_{\mathcal{N}}$ are all independent, one has

$$\begin{aligned} P(d_1, \dots, d_{\mathcal{N}}|H_{i_1 i_2 \dots i_k}, \mathbf{I}) &= \prod_{A=1}^{\mathcal{N}} P(d_A|H_{i_1 i_2 \dots i_k}, \mathbf{I}), \\ P(d_1, \dots, d_{\mathcal{N}}|\mathcal{H}_{\text{GR}}, \mathbf{I}) &= \prod_{A=1}^{\mathcal{N}} P(d_A|\mathcal{H}_{\text{GR}}, \mathbf{I}). \end{aligned} \quad (43)$$

Thus,

$$(\text{cat})B_{\text{GR}}^{i_1 i_2 \dots i_k} = \prod_{A=1}^{\mathcal{N}} (A)B_{\text{GR}}^{i_1 i_2 \dots i_k}, \quad (44)$$

with

$$(A)B_{\text{GR}}^{i_1 i_2 \dots i_k} = \frac{P(d_A|H_{i_1 i_2 \dots i_k}, \mathbf{I})}{P(d_A|\mathcal{H}_{\text{GR}}, \mathbf{I})}. \quad (45)$$

To evaluate the combined odds ratio of the catalog, we choose to invoke indifference [as in Eqs. (28) and (36)] and set the individual prior odds ratios equal to each other, so that

$$\frac{P(H_{i_1 i_2 \dots i_k}|\mathbf{I})}{P(\mathcal{H}_{\text{GR}}|\mathbf{I})} = \frac{\alpha}{2^{N_T} - 1}. \quad (46)$$

Together with Eqs. (41) and (44), this leads to

$$(N_T)O_{\text{GR}}^{\text{modGR}} = \frac{\alpha}{2^{N_T} - 1} \sum_{k=1}^{N_T} \sum_{i_1 < i_2 < \dots < i_k} \prod_{A=1}^{\mathcal{N}} (A)B_{\text{GR}}^{i_1 i_2 \dots i_k}, \quad (47)$$

which, up to an overall prefactor, amounts to taking the average of the cumulative Bayes factors (44).

Alternatively, one may prefer not to make any assumptions for the prior odds ratios $P(H_{i_1 i_2 \dots i_k}|\mathbf{I})/P(\mathcal{H}_{\text{GR}}|\mathbf{I})$ at all, and focus on the cumulative Bayes factors $(\text{cat})B_{\text{GR}}^{i_1 i_2 \dots i_k}$ separately and individually. It will also be of interest to look at the cumulative Bayes factors for the various hypotheses $H_{i_1 i_2 \dots i_k}$ and $\mathcal{H}_{\text{GR}}$ against the noise-only hypothesis $\mathcal{H}_{\text{noise}}$:

$$\begin{aligned} \prod_A (A)B_{\text{noise}}^{i_1 i_2 \dots i_k} &= \prod_A \frac{P(d_A|H_{i_1 i_2 \dots i_k}, \mathbf{I})}{P(d_A|\mathcal{H}_{\text{noise}}, \mathbf{I})}, \\ \prod_A (A)B_{\text{noise}}^{\text{GR}} &= \prod_A \frac{P(d_A|\mathcal{H}_{\text{GR}}, \mathbf{I})}{P(d_A|\mathcal{H}_{\text{noise}}, \mathbf{I})}, \end{aligned} \quad (48)$$

and we note that

$$(\text{cat})B_{\text{GR}}^{i_1 i_2 \dots i_k} = \frac{\prod_A (A)B_{\text{noise}}^{i_1 i_2 \dots i_k}}{\prod_{A'} (A')B_{\text{noise}}^{\text{GR}}}. \quad (49)$$

Finally, we will also look at the individual contributions $(A)B_{\text{noise}}^{i_1 i_2 \dots i_k}$ and $(A)B_{\text{noise}}^{\text{GR}}$, for $A = 1, 2, \dots, \mathcal{N}$.

IV. RESULTS

To illustrate the method, we construct catalogs of 15 binary neutron star sources, distributed uniformly in volume, with random sky positions and orientations, and whose total number is taken to be on the conservative side of the “realistic” estimates in [15] for the number of detectable sources in a one-year time span. We take the individual neutron star masses to lie between 1 and $2M_{\odot}$. The distance interval is between 100 Mpc and 400 Mpc; the former number is the radius within which one would expect ~ 0.5 BNS inspirals per year, and 400 Mpc is the approximate horizon distance in Advanced LIGO. The corresponding signals are added coherently to stationary, Gaussian simulated data for the Advanced Virgo interferometer and the two Advanced LIGOs. A lower cutoff of 8 is implemented on the optimal network SNR, defined as the quadrature sum of individual detector SNRs. The analysis of the surviving signals is performed with an appropriately modified version of the nested sampling code available in LAL[63].

For simplicity, we took the phase up to only 2PN order ($M = 4$), both in the signal and in the model waveforms. As an example, we considered the case of $N_T = 3$ testing coefficients ψ_1, ψ_2, ψ_3 . Nested sampling then gives $2^3 - 1 = 7$ Bayes factors $B_{\text{GR}}^1, B_{\text{GR}}^2, B_{\text{GR}}^3, B_{\text{GR}}^{12}, B_{\text{GR}}^{13}, B_{\text{GR}}^{23}$, and B_{GR}^{123} , which for a single source can be combined to form the odds ratio (39). Note that up to and including 2PN order, the post-Newtonian coefficients take the general form [23]

$$\psi_i(\mathcal{M}, \eta) = \frac{3}{128\eta} g_i(\eta) (\pi \mathcal{M} \eta^{-3/5})^{(i-5)/3}, \quad (50)$$

where the $g_i(\eta)$ are polynomials in η ; for the lowest three subleading PN orders, one has

$$\begin{aligned} g_1(\eta) &\equiv 0, & g_2(\eta) &= \frac{20}{9} \left(\frac{743}{336} + \frac{11}{4} \eta \right), \\ g_3(\eta) &= -16\pi. \end{aligned} \quad (51)$$

Accordingly, in the model waveforms we allow for deformations

$$\begin{aligned}
\psi_1^{\text{GR}}(\mathcal{M}, \eta) &= 0 \rightarrow \frac{3}{128\eta} (\pi\mathcal{M}\eta^{-3/5})^{-4/3} \delta\chi_1, \\
\psi_2^{\text{GR}}(\mathcal{M}, \eta) &\rightarrow \frac{3}{128\eta} g_2(\eta) (\pi\mathcal{M}\eta^{-3/5})^{-1} [1 + \delta\chi_2], \\
\psi_3^{\text{GR}}(\mathcal{M}, \eta) &\rightarrow \frac{3}{128\eta} g_3(\eta) (\pi\mathcal{M}\eta^{-3/5})^{-2/3} [1 + \delta\chi_3].
\end{aligned} \tag{52}$$

We take the prior on $\delta\chi_i$ to be flat and centered on zero, with a total width of 0.5. This will be much larger than the deviations we will use in simulated signals and hence suffices to illustrate the method; for real measurements one may want to choose a still wider prior.

For the purposes of showing results, the factor α in Eqs. (39) and (47) will be set to 1.

A. Measurability of a deviation in a post-Newtonian coefficient

In order to gauge the sensitivity of our method to deviations from GR, we first consider a large number of simulated signals with a constant relative offset $\delta\chi_3$. If GR is violated, we do not expect the deviation to be this simple, but these examples will serve to illustrate both the workings and the effectiveness of the technique for low-SNR sources of the kind we expect in Advanced LIGO and Virgo. We note that ψ_3 is the lowest-order coefficient which incorporates the nonlinearity of general relativity through so-called tail effects [20,21], and is therefore of particular interest.

1. Signals with constant relative deviation $\delta\chi_3 = >0.1$

We start with signals that have $\delta\chi_3 = 0.1$. We first compute the odds ratios for individual sources, $^{(N_T)}O_{\text{GR}}^{\text{modGR}}$, according to (39), with $N_T = 3$. Next we divide these up *randomly* into catalogs of 15 sources each and compute the combined odds ratios $^{(N_T)}O_{\text{GR}}^{\text{modGR}}$ as in (47). We do the same thing for injections that are pure GR, i.e. $\delta\chi_i = 0$ for all i , and again compute the quantities $^{(N_T)}O_{\text{GR}}^{\text{modGR}}$ for individual sources, and $^{(N_T)}O_{\text{GR}}^{\text{modGR}}$ for catalogs of 15 sources each.

Before considering catalogs, let us look at the (log) odds ratios for individual sources as a function of SNR. This is shown in Fig. 4. The overwhelming majority of signals have SNR between 8 and 15, consistent with our SNR cut and the placement of sources uniform in volume up to 400 Mpc. Even for low-SNR sources, there is separation between the GR injections and the injections with modified ψ_3 . As one would expect, the separation becomes much clearer with increasing SNR.

Figure 5 shows normalized distributions of the log odds ratios, both for individual sources and for catalogs of 15 sources each. Combining the odds ratios for sources within a catalog will strongly boost our confidence in a violation of GR if one is present at the given level. For this particular




FIG. 4 (color online). The log odds ratios for individual sources. The blue crosses represent signals with standard GR waveforms, the red circles are signals with a constant 10% relative offset in ψ_3 . A separation between the two is visible for $\text{SNR} \geq 10$ and becomes more pronounced as the SNR increases.




FIG. 5 (color online). Top panel: The normalized distribution $P(\ln O_{\text{GR}}^{\text{modGR}})$ of log odds ratios for individual sources, where the injections are either GR or have $\delta\chi_3 = 0.1$. Bottom panel: The normalized distribution $P(\ln O_{\text{GR}}^{\text{modGR}})$ of logs of the combined odds ratios for GR injections and injections with $\delta\chi_3 = 0.1$, for catalogs of 15 sources each. The effectiveness of the catalog approach to testing for deviations from GR comes from the combination of multiple sources, each source contributing to the overall result in proportion to its own Bayes factors.

choice of $\delta\chi_3$ in the injected non-GR signals, the separation from the GR injections is complete.

It is useful to look at which of the Bayes factors of the component hypotheses tend to give the largest contribution to the odds ratio. What is computed *directly* by the nested sampling code is not $B_{\text{GR}}^1, \dots, B_{\text{GR}}^3, B_{\text{GR}}^{12}, \dots, B_{\text{GR}}^{23}, B_{\text{GR}}^{123}$, but rather the Bayes factors for each of the hypotheses against the noise-only hypothesis $\mathcal{H}_{\text{noise}}$:

$$B_{\text{noise}}^{i_1 \dots i_k} = \frac{P(d|H_{i_1 \dots i_k}, \mathbf{I})}{P(d|\mathcal{H}_{\text{noise}}, \mathbf{I})}, \quad B_{\text{noise}}^{\text{GR}} = \frac{P(d|\mathcal{H}_{\text{GR}}, \mathbf{I})}{P(d|\mathcal{H}_{\text{noise}}, \mathbf{I})}, \quad (53)$$

and one has

$$B_{\text{GR}}^{i_1 \dots i_k} = \frac{B_{\text{noise}}^{i_1 \dots i_k}}{B_{\text{noise}}^{\text{GR}}}. \quad (54)$$

In Fig. 6, we show the cumulative number of times that a particular $B_{\text{noise}}^{i_1 \dots i_k}$ is the largest, against SNR, for the case where the injections have $\delta\chi_3 = 0.1$. The results are




FIG. 6 (color online). Top panel. Curves and left vertical axis: For a given SNR, the cumulative number of times that the Bayes factor against noise for a particular component hypothesis is the largest for injections with that SNR or below, for $\delta\chi_3 = 0.1$. All 1471 simulated sources were used. As expected, B_{noise}^3 dominates, and Bayes factors for hypotheses that let ψ_3 be non-GR tend to outperform those that do not. The GR model outperforms the one with the largest number of free parameters. Histogram and right vertical axis: The number of sources per SNR bin. Bottom panel: The same as above, but restricting to sources with $\text{SNR} < 12$. A similar behavior as for the full set of 1471 sources is observed. Note that, already at SNR close to threshold, the GR hypothesis is more likely to be disfavored.

entirely as expected, considering that the injected waveform has a shift in ψ_3 only:

- (i) The Bayes factor B_{noise}^3 corresponding to the hypothesis H_3 dominates.
- (ii) The Bayes factors $B_{\text{noise}}^{i_1 \dots i_k}$ corresponding to hypotheses that involve ψ_3 being non-GR tend to outperform those that do not.
- (iii) The Bayes factors for the non-GR hypotheses deviate from the GR one already at low SNR, showing that our method will perform well in the low-SNR scenario.

Because of the first two points, one may be tempted to assign different prior odds to the various hypotheses instead of setting them all equal to each other. For instance, one might consider downweighting the most inclusive hypothesis, H_{123} , by invoking Occam's razor. However, the violation of GR we assume here is of a rather special form. In reality, one will not know beforehand what the nature of the deviation will be; in particular, its effect may not be restricted to a single phase coefficient. It is possible that *all* coefficients are affected, in which case one would not want to *a priori* deprecate H_{123} . As explained in Sec. III C, our hypothesis $\mathcal{H}_{\text{modGR}}$ corresponds to the question of whether *one or more* of the phasing coefficients $\{\psi_1, \psi_2, \psi_3\}$ differ from their GR values; one may want to ask a different question, but this is the one that is the most general within our framework. To retain full generality, all subhypotheses $H_{i_1 i_2 \dots i_k}$ need to be taken into account and given equal weight.

2. Signals with constant relative deviation $\delta\chi_3 = 0.025$

It is clear that, if signals arriving at the Advanced Virgo-LIGO network have a (constant) fractional deviation in ψ_3 as large as 10%, then at least under the assumption of Gaussian noise, we would have no trouble in discerning this violation of GR even if only 15 events were ever recorded. Now let us look at a smaller deviation in ψ_3 , say, 2.5%.

In Fig. 7 we plot the log odds ratios for individual sources against SNR, both for signals with GR waveforms and signals with $\delta\chi_3 = 0.025$. This time, the two distributions largely coincide, although there are some outliers which could boost the *combined* odds ratio when they are present in a catalog of sources.

In Fig. 8 we show normalized distributions of the log odds ratio for individual sources, as well as for catalogs with 15 sources each. For individual sources, the distributions are more or less on top of each other. The picture is somewhat different for the catalogs. If a catalog with $\delta\chi_3 = 0.025$ happens to contain one of the outliers visible in Fig. 7, then it can boost the combined odds ratio for the catalog.

It is instructive to look at a representative catalog with $\ln^{(3)} \mathcal{O}_{\text{GR}}^{\text{modGR}} > 0$. In Fig. 9 we show the buildup of the log Bayes factors for the various subhypotheses against GR, as




FIG. 7 (color online). The log odds ratios for individual sources. The blue crosses are for signals with standard GR waveforms, and the red circles are for signals with a constant 2.5% relative offset in ψ_3 . This time, there is little separation between the two, although there are outliers for $\text{SNR} \approx 15$.




FIG. 9 (color online). The buildup of cumulative Bayes factors against GR for individual hypotheses, and the odds ratio, for a typical catalog with $\delta\chi_3 = 0.025$. Note that on the basis of the Bayes factor against GR of the most inclusive hypothesis H_{123} alone, one can conclude that the GR model is in fact the favored one. Even the log Bayes factor for H_{23} ends up being negative. Additionally, the hypothesis with the largest Bayes factor is not H_3 but H_1 . This illustrates that it is necessary to include as many hypotheses as possible in the analysis.




FIG. 8 (color online). Top panel: The normalized distribution $P(\ln O_{GR}^{\text{modGR}})$ of log odds ratios for individual sources, where the injections are either GR or have $\delta\chi_3 = 0.025$. Bottom panel: The normalized distribution $P(\ln O_{GR}^{\text{modGR}})$ of logs of the combined odds ratios for GR injections and injections with $\delta\chi_3 = 0.025$, for catalogs of 15 sources each. For individual sources, the two distributions essentially lie on top of each other. However, when sources are combined into catalogs, it is possible for an outlier to boost the odds ratio of the entire catalog.

well as the odds ratio itself. In a scenario where the evidence for a GR violation is marginal, it is imperative to include as many hypotheses as possible in the analysis. Indeed, in this example, the hypothesis H_3 is not the most favored one; instead, it is H_1 . Note also that if we had only tested the most inclusive hypothesis H_{123} against GR (as one might do if one expects a deviation in all of the PN parameters), we would have concluded that the GR hypothesis is the favored one. The same is true if we had only tested H_2 , as one would do when specifically looking for a massive graviton. Even the log Bayes factor for H_{23} ends up being negative. We also remind the reader that we will not know beforehand what the precise nature of the GR violation will be.

By using the distribution of log odds ratios for simulated catalogs of GR sources, one can establish a *threshold* which the odds ratio of a given catalog must overcome in order for a violation of GR to become credible. This would be the analog of what was done in [70] (see Fig. 7 of that paper), where the distribution of the log Bayes factor $\ln B_{S,N}$ for the presence of a signal versus noise only was computed for many realizations of the noise, in the absence of a signal. Consider the distribution $P(\ln^{(N_T)} O_{GR}^{\text{modGR}} | \kappa, \mathcal{H}_{GR}, I)$ of the log odds ratio for the collection κ of simulated catalogs of signals that are in accordance with GR. Given a “false alarm probability” β , a threshold $\ln \mathcal{O}_\beta$ for the odds ratio can be set as follows:

$$\beta = \int_{\ln \mathcal{O}_\beta}^{\infty} P(\ln \mathcal{O} | \kappa, \mathcal{H}_{GR}, I) d \ln \mathcal{O}. \quad (55)$$

Now suppose we also have a distribution $P(\ln^{(N_T)} O_{GR}^{\text{modGR}} | \kappa', \mathcal{H}_{\text{alt}}, I)$ of the log odds ratio for a

collection κ' of simulated catalogs of signals which follow some alternative theory (in this example, one which leads to a shift $\delta\chi_3 = 0.025$). Then we can quantify the chance that a deviation from GR of this particular kind will be detected with a false alarm probability smaller than the given β , by means of an *efficiency* ζ , defined as

$$\zeta = \int_{\ln\mathcal{O}_\beta}^{\infty} P(\ln\mathcal{O}|\kappa', \mathcal{H}_{\text{alt}}, \mathcal{I}) d\ln\mathcal{O}. \quad (56)$$

We note that with these definitions, the efficiency is independent of the overall prior odds of $\mathcal{H}_{\text{modGR}}$ versus $\mathcal{H}_{\text{GR}}$, as the factor α in Eqs. (37) and (47) will just cause both the distributions $P(\ln^{(N_T)}\mathcal{O}_{\text{GR}}^{\text{modGR}}|\kappa, \mathcal{H}_{\text{GR}}, \mathcal{I})$ and $P(\ln^{(N_T)}\mathcal{O}_{\text{GR}}^{\text{modGR}}|\kappa', \mathcal{H}_{\text{alt}}, \mathcal{I})$, and the threshold $\ln\mathcal{O}_\beta$, to be shifted by $\ln\alpha$.

In the present example, with $\delta\chi_3 = 0.025$ and catalogs of 15 sources each, for $\beta = 0.05$ one has $\zeta = 0.22$. Hence, by this standard and for the given number of sources in a catalog, a GR violation of this kind is just borderline detectable.

However, the number of sources per catalog, 15 in the examples shown so far, and the chosen false alarm probability, $\beta = 0.05$, are somewhat arbitrary. In reality, the size of the catalog will depend on the number of detected sources, and the false alarm probability is set according to the required confidence. It is therefore of interest to investigate the effects of both of these factors. In Fig. 10 we show the behavior of the efficiency as a function of the catalog size and the false alarm probability. To account for the arbitrariness in which the sources are combined to form a catalog, we show the median and 68% confidence interval of the efficiencies from 5000 random source orderings as the central curves and the error bars, respectively. Results are shown for $\beta \in \{0.05, 0.01\}$.

As is evident from Fig. 10, the efficiency rises as a function of the catalog size. This highlights the importance of combining all available sources in the advanced detector era, when one looks for deviations from GR. The maximum catalog sizes shown are comparable to the “realistic” estimates of the number of detections of binary neutron star inspirals in the span of a year [15].

We see that $\delta\chi_3 = 0.025$ is a borderline case in terms of the discernability of a GR violation. Later on, when we show posterior PDFs, it will become evident that, indeed, $\delta\chi_3$ can typically be measured with an accuracy of this order.

B. Effect of number of hypotheses used

It is of interest to see what would have happened if we had used a smaller number of testing coefficients, say, $\{\psi_1, \psi_2\}$, so that the hypotheses to be tested are H_1 , H_2 , and H_{12} . In the example with $\delta\chi_3 = 0.025$ as in the previous subsection, the PN order where the deviation occurs, namely, 1.5PN, would then be higher than the PN




FIG. 10 (color online). The efficiency of detecting a GR violation for sources with $\delta\chi_3 = 0.025$, as a function of catalog size for false alarm probabilities $\beta \in \{0.05, 0.01\}$. The median and the 68% confidence interval from 5000 random catalog orderings are shown as the central curve and the error bars, respectively. The efficiency increases as a function of catalog size, once again underscoring the benefit of combining all available data.

orders associated with our testing coefficients, which are 0.5PN and 1PN.

In Fig. 11, the following are shown:

- (i) In the case where only $\{\psi_1, \psi_2\}$ are testing coefficients, we compute the thresholds $\ln^{(2)}\mathcal{O}_\beta$ corresponding to false alarm probabilities (FAPs) $\beta \in \{0.05, 0.01\}$. Next, we recalculate the false alarm probabilities *for the same thresholds*, but now for the case where there are three testing parameters, $\{\psi_1, \psi_2, \psi_3\}$, and show the difference in false alarm probabilities.
- (ii) On the other hand, one can compare the efficiencies $^{(2)}\zeta$ and $^{(3)}\zeta$ for the two- and three-parameter cases, *for fixed false alarm probabilities* $\beta \in \{0.05, 0.01\}$. This is shown in the bottom panel.

As expected, in the first case (fixed thresholds for the odds ratios), the false alarm probabilities increase in going from two to three testing parameters, but only moderately so. On the other hand, for fixed false alarm probabilities, there is no appreciable change in efficiency. Indeed, the spread in the GR “background” will increase with an increase in hypotheses to be tested against GR; yet, having more hypotheses does not really hurt us in terms of our ability to detect a deviation from GR.

Figure 11 indicates the typical behavior for catalogs with a *specific* deviation from GR, in this case $\delta\chi_3 = 0.025$. It is worth repeating, however, that especially when there is only marginal evidence for a GR violation, it is important to use as many hypotheses as computationally feasible; see Fig. 9 (and also Fig. 16 below). Also, we will obviously not know beforehand what the *nature* of the GR violation is.

One may nevertheless wonder how our three-parameter test would compare with a “targeted search” that only




FIG. 11 (color online). Top panel: The change in FAPs in going from two to three testing parameters, but keeping the odds ratio thresholds fixed. Bottom panel: The change in efficiencies when keeping the false alarm probabilities fixed. The plots shown are for the case where the signals have $\delta\chi_3 = 0.025$. We see that increasing the number of testing parameters has only a moderate effect on the FAPs, while for fixed FAPs, the efficiencies do not change appreciably. Note, however, that when the evidence for a deviation from GR is marginal, the use of as many hypotheses as possible can be pivotal in finding the violation (see Figs. 9 and 16).

looks for a deviation in ψ_3 , which in this example happens to be where the deviation actually is. With our choice of $\alpha = 1$, this corresponds to setting $\mathcal{O}_{\text{GR}}^{\text{modGR}} = {}^{(\text{cat})}B_{\text{GR}}^3$. Figure 12 shows the change in false alarm probabilities in going from testing only H_3 to the full test for fixed log odds ratio thresholds, as well as the change in efficiencies for fixed false alarm probabilities. The results are as follows:

- (i) The change in false alarm probabilities for fixed log odds ratio thresholds is minor.
- (ii) However, especially for a large number of sources per catalog, the efficiencies show a clear rise. This can be accounted for by the fact that, for a small violation of GR, it will not always be the case that the Bayes factor against GR for H_3 is the largest, but our method is able to compensate for that.

We conclude that for this *particular* example, our method with $N_T = 3$ testing parameters will tend to outperform a




FIG. 12 (color online). Top panel: The difference of false alarm probabilities, $(3)\beta - (\text{target})\beta$, for fixed log odds ratio thresholds and signals having $\delta\chi_3 = 0.025$, between our three-parameter test and a “targeted search” which only looks for a deviation in ψ_3 , i.e., only tests the hypothesis H_3 against GR. We see that the difference is minor. Bottom panel: The difference in efficiencies (which is more important), $(3)\zeta - (\text{target})\zeta$, for fixed false alarm probabilities. Especially for a large number of sources per catalog, our three-parameter test is actually more efficient than the “targeted search,” at least for this particular example. This is because the Bayes factor against GR for H_3 will not be the largest in every catalog, but our method naturally compensates for that. Of course, we do not expect our method to outperform a targeted search in the case of a more complicated deviation from GR.

“targeted search” that happens to look for the violation actually present. However, we do not expect this to be true for more complicated deviations from GR.

C. Measurability of deviations with non-PN frequency dependences

The aim of the previous subsections was to get a rough idea of the sensitivity of our method to deviations in post-Newtonian coefficients, and in order to gauge this we assumed a constant relative offset in the physically interesting parameter ψ_3 . However, we stress once again that we do not expect a violation of GR to manifest itself as a simple constant relative shift in one of the post-Newtonian coefficients. Even if modifications are confined to the PN

coefficients, the $\delta\chi_i$ in the signals can be dependent on $(\mathcal{M}, \eta)$, in addition to whatever charges and coupling constants may be present. Moreover, a deviation from GR could introduce terms in the phase with frequency dependences that do not correspond to any of the PN contributions. We now show that the method can also be sensitive to violations of that kind, even though the model waveforms we use in our analyses only have deformations of PN terms. Let us give a heuristic example where the phase of the simulated signals contains a term with an anomalous frequency dependence in between that of the 1PN and 1.5PN contributions. Specifically,

$$\frac{3}{128\eta}(\pi\mathcal{M}\eta^{-3/5})^{-5/6}\delta\chi_A f^{-5/6}, \quad (57)$$

and we note that the 1PN term goes like f^{-1} and the 1.5PN term like $f^{-2/3}$; thus, the deviation introduced here could be dubbed “1.25PN.” However, for the recovery, we will continue to use the same model waveforms as before, which can only have shifts in the phase coefficients at 0.5PN, 1PN, and 1.5PN. Our aim in this subsection is to show that they will nevertheless allow us to find a deviation in the signal of the form (57).

We now need to make a choice for $\delta\chi_A$. We aim to show that even if there is a deviation in the phase that is not represented in any of our model waveforms, it can be recovered if near the “bucket” of the noise curve (at $f \sim 150$ Hz) the amount by which it affects the phase is on par with a shift in the PN coefficients of more than a few percent.

For definiteness, let us take $\delta\chi_A$ to be constant, such that at $f = 150$ Hz and for a system of $(1.5, 1.5)M_\odot$, the contribution (57) to the phase is equal to the change caused by a shift in the 1.5PN contribution with $\delta\chi_3 = 0.1$:

$$\begin{aligned} &(\pi 3M_\odot)^{-5/6}\delta\chi_A(150 \text{ Hz})^{-5/6} \\ &= g_3(0.25)(\pi 3M_\odot)^{-2/3} \times 0.1 \times (150 \text{ Hz})^{-2/3}, \end{aligned} \quad (58)$$

leading to $\delta\chi_A = -2.2$.

As before, we first give results for the odds ratios of individual sources with increasing SNR; see Fig. 13. We see that even at small SNR there is already a good separation between the GR injections and the injections with a modification in the structure of the phase.

Next we show normalized distributions of the log odds ratios, both for individual sources and for catalogs of 15 sources each: Fig. 14. As expected, for the catalogs there is an excellent separation between the GR injections and injections with a modification in the phase.

Now let us look at the cumulative number of times that the Bayes factor against noise for a particular hypothesis is the largest, for individual sources with $\delta\chi_A = -2.2$, arranged with increasing SNR (Fig. 15). From $\text{SNR} \gtrsim 15$, the Bayes factor B_{noise}^2 starts to dominate, followed by B_{noise}^{23} and B_{noise}^3 , with the latter two crossing over between




FIG. 13 (color online). The log odds ratios for individual sources for injections with an anomalous frequency dependence in the phase. The blue crosses represent GR injections, while the red circles are for signals that have a contribution to the phase with a frequency dependence in between that of the 1PN and 1.5PN terms (1.25PN). As with the $\delta\chi_3 = 0.1$ example, a separation between the two is visible for $\text{SNR} \gtrsim 10$ and becomes more pronounced as the SNR increases.




FIG. 14 (color online). Top panel: The normalized distribution $P(\ln O_{\text{GR}}^{\text{modGR}})$ of log odds ratios for individual sources, where the injections are either GR or have the anomalous frequency dependence. Bottom panel: The normalized distribution $P(\ln O_{\text{GR}}^{\text{modGR}})$ of logs of the combined odds ratios for GR injections and injections for catalogs of 15 sources each.




FIG. 15 (color online). Top panel. Curves and left vertical axis: The cumulative number of times that the Bayes factor against noise for a particular component hypothesis is the largest, with increasing SNR, for individual sources with $\delta\chi_A = -2.2$. All 854 simulated sources were used. Histogram and right vertical axis: The number of sources per SNR bin. Bottom panel: The same as above, but now for sources with SNR < 12 . Close to threshold, no single hypothesis is the dominant one.

SNR ~ 20 and SNR ~ 25 . Already at SNR ~ 9 , all of the $B_{\text{noise}}^{i_1 i_2 \dots i_k}$ dominate the Bayes factor $B_{\text{noise}}^{\text{GR}}$ for the GR hypothesis. However, near the SNR threshold, no single hypothesis dominates clearly, which again shows that as many hypotheses as possible should be included in the analysis.

Especially in this case, it is interesting to look at the growth of cumulative Bayes factors against GR for individual hypotheses, as well as of the odds ratio, as sources with increasing SNR are being added within catalogs of 15 sources. This is shown for a few example catalogs in Fig. 16. The salient features are as follows:

- (i) Even if all 15 sources only have modest SNR, by their cumulative contributions they can cause a relatively large odds ratio for the catalog as a whole.
- (ii) In catalogs containing a source with a particularly high SNR, it is by no means a given that the contribution of this source will dominate the odds ratio compared to the cumulative contributions of the other sources.
- (iii) Which hypothesis comes out on top will vary from one catalog to another; in the examples of Fig. 16 we see H_{12} , H_{23} , or H_2 giving the largest




FIG. 16 (color online). A few examples of how cumulative Bayes factors against GR for individual hypotheses, and the odds ratio, grow as sources with increasing SNR are being added within three different catalogs of 15 sources in total. Top and middle panels: Catalogs where sources have only modest SNRs (< 20). Note the large differences in contributions from different hypotheses, and in the ordering of Bayes factors, between these two catalogs. Bottom panel: A catalog with a high SNR source; note, however, that the source with the highest SNR does not cause a particularly large “boost,” and the cumulative log odds ratio ends up being considerably lower than in the top and middle examples.

contribution, respectively, but there are examples where any of the other four subhypotheses contributes the most. In this respect, we note that the odds ratio for a catalog is proportional to the average of the cumulative Bayes factors themselves, not of their logarithms. If one were to *a priori* favor particular (subsets of) hypotheses, the log odds ratio could be lowered by as much as 100. This could have a large effect on the false alarm

probability; see Fig. 14. These are again arguments for using as many subhypotheses as possible, and giving them equal relative prior odds.

Note that, in principle, we could have extended our model waveforms with more free parameters so as to be more sensitive to deviations of this kind, e.g. by including a term in the phase of the form $\mathcal{A}f^a$, similar to what one has in the so-called ppE waveforms [53]. However, the point of our method is to search for *generic* deviations from GR. In the future, we will want to search with more sophisticated (time domain) waveforms whose use is computationally more demanding in a Bayesian setting, and we will not be able to allow for an arbitrarily large number of deformations in the model waveforms. However, here we have an example where the injections have a contribution to the phase that is not present in any of the model waveforms, yet is clearly observable. Recall that the overall magnitude of the anomalous phase contribution was chosen so that, at $f \sim 150$ Hz, it changes the phase by a similar amount as a shift in ψ_3 of 10%. Our results make it plausible that generically, when there is a deviation in the phase which, at frequencies where the detectors are the most sensitive, causes a phase change on par with a change in one of the (low-order) PN coefficients of more than a few percent, it will be detectable.

D. Parameter estimation

Finally, let us look at some posterior PDFs for the $\delta\chi_i$. We stress that unlike Bayes factors, the PDFs cannot be combined across sources since we should not expect the $\delta\chi_i$ to be independent of the component masses; they can differ from source to source. Even looking at the PDFs for a single source may then be misleading: even if the deviation is exactly in one or more of the PN coefficients, a given source will have values for the $\delta\chi_i$ that are representative just for the $(\mathcal{M}, \eta)$ of that source, and possibly also the values of additional charges that may appear in an alternative theory of gravity. In a given catalog, there may be only one source with sufficient SNR to allow for accurate parameter estimation, in which case the posteriors will not tell us much even if they are strongly peaked. More generally, the deviation from GR may manifest itself by the appearance of terms in the phase that do not have the frequency dependence of any of the PN contributions. However, in the event that the log odds ratio and Bayes factors strongly favor the GR hypothesis, posterior PDFs will allow us to constrain deviations in the PN coefficients, thereby adding further support that GR is the correct theory. Hence, we start with an analysis of pure GR injections.

1. Example: A GR injection

Let us first look at a GR source with $(\mathcal{M}, \eta, D) = (1.31M_\odot, 0.243, 131 \text{ Mpc})$, and a LIGO-Virgo network

SNR of 23.0. The Bayes factors for the various component hypotheses are

$$\begin{aligned} \ln B_{\text{GR}}^1 &= -2, & \ln B_{\text{GR}}^2 &= -2, & \ln B_{\text{GR}}^3 &= -2, \\ \ln B_{\text{GR}}^{12} &= -3, & \ln B_{\text{GR}}^{13} &= -1, & \ln B_{\text{GR}}^{23} &= -1, \\ \ln B_{\text{GR}}^{123} &= -2. \end{aligned} \quad (59)$$

We recall that these are the Bayes factors for a particular deviation from GR versus GR. The GR hypothesis is favored in all cases. We can also look at the Bayes factors for all of the hypotheses against noise:

$$\begin{aligned} \ln B_{\text{noise}}^{\text{GR}} &= 211, & \ln B_{\text{noise}}^1 &= 209, & \ln B_{\text{noise}}^2 &= 209, \\ \ln B_{\text{noise}}^3 &= 209, & \ln B_{\text{noise}}^{12} &= 208, & \ln B_{\text{noise}}^{13} &= 210, \\ \ln B_{\text{noise}}^{23} &= 210, & \ln B_{\text{noise}}^{123} &= 209. \end{aligned} \quad (60)$$

Hence, the signal is picked up very well by the waveforms of all of the hypotheses, with the GR waveform doing slightly better.

Let us now look at some posterior PDFs. In Fig. 17, we show the PDFs for $\delta\chi_1$, $\delta\chi_2$, and $\delta\chi_3$, respectively, for the waveforms that have free parameters $\{\vec{\theta}, \delta\chi_1\}$, $\{\vec{\theta}, \delta\chi_2\}$, and $\{\vec{\theta}, \delta\chi_3\}$, with $\vec{\theta}$ the parameters of the GR waveform. We see that the distributions are all narrowly peaked around the correct value of zero.

2. Example: A signal with $\delta\chi_3 = 0.1$

We now consider an example with $(\mathcal{M}, \eta, D) = (1.18M_\odot, 0.244, 196 \text{ Mpc})$, with a nonzero relative shift in ψ_3 of $\delta\chi_3 = 0.1$, and network SNR of 23.2. The Bayes factors are

$$\begin{aligned} \ln B_{\text{GR}}^1 &= 117, & \ln B_{\text{GR}}^2 &= 124, & \ln B_{\text{GR}}^3 &= 124, \\ \ln B_{\text{GR}}^{12} &= 123, & \ln B_{\text{GR}}^{13} &= 124, & \ln B_{\text{GR}}^{23} &= 125, \\ \ln B_{\text{GR}}^{123} &= 114. \end{aligned} \quad (61)$$

This time, the GR hypothesis is very much disfavored. However, we note that the Bayes factor for the hypothesis that only ψ_3 differs from its GR value is not the largest. In fact, all the Bayes factors except for B_{GR}^1 and B_{GR}^{123} are rather similar in magnitude, and no clear conclusions can be drawn from them regarding the underlying nature of the deviation from GR.

When looking at the Bayes factors against noise, we see that the signal is clearly detected for all hypotheses:

$$\begin{aligned} \ln B_{\text{noise}}^{\text{GR}} &= 128, & \ln B_{\text{noise}}^1 &= 245, & \ln B_{\text{noise}}^2 &= 252, \\ \ln B_{\text{noise}}^3 &= 252, & \ln B_{\text{noise}}^{12} &= 251, & \ln B_{\text{noise}}^{13} &= 252, \\ \ln B_{\text{noise}}^{23} &= 253, & \ln B_{\text{noise}}^{123} &= 242. \end{aligned} \quad (62)$$

Now let us consider posterior PDFs. We expect the PDF of $\delta\chi_3$ for the hypothesis H_3 , where only $\{\vec{\theta}, \delta\chi_3\}$ are allowed to vary, to be peaked at the injected value of 0.1,




FIG. 17. Posterior PDFs for a single GR injection with network SNR of 23.0. Top panel: $\delta\chi_1$ measured with a waveform that has $\{\hat{\theta}, \delta\chi_1\}$ as free parameters. Middle panel: $\delta\chi_2$ measured with a waveform that has $\{\hat{\theta}, \delta\chi_2\}$ free. Bottom panel: $\delta\chi_3$ measured with a waveform that has $\{\hat{\theta}, \delta\chi_3\}$ free. In each case the distribution is tightly centered on zero, with standard deviations of 0.014, 0.015, and 0.019, respectively.

and this is the case with very good accuracy, as shown in the bottom panel of Fig. 18.

In the upper and middle panels of the same figure, the PDFs of $\delta\chi_1$ for the hypothesis H_1 , and of $\delta\chi_2$ for the hypothesis H_2 are shown. In these cases the parameter in the signal that has the shift is now not represented; in the first case only $\delta\chi_1$ is allowed to vary on top of the parameters $\hat{\theta}$ of GR, and in the second case only $\delta\chi_2$. In the nested sampling process, the waveform will still try to adapt itself to the deformation in the signal. The result is that $\delta\chi_1$ and $\delta\chi_2$ are strongly peaked, but away from the correct values $\delta\chi_1 = \delta\chi_2 = 0$. Thus, if one were to study the data only using waveforms from a *specific* alternative




FIG. 18. Posterior PDFs for a single injection with $\delta\chi_3 = 0.1$ and network SNR of 23.2. Top panel: $\delta\chi_1$ measured with a waveform that has $\{\hat{\theta}, \delta\chi_1\}$ free. Middle panel: $\delta\chi_2$ measured with a waveform that has $\{\hat{\theta}, \delta\chi_2\}$ free. Bottom panel: $\delta\chi_3$ measured with a waveform that has $\{\hat{\theta}, \delta\chi_3\}$ free. As expected, the bottom PDF is sharply peaked at the correct value of $\delta\chi_3 = 0.1$, with a standard deviation of 0.012. For the top and middle PDFs, the one test parameter that is used differs from the parameter in the signal that has the shift. The parameters in the waveform will rearrange themselves such as to best accommodate the properties of the signal. Both $\delta\chi_1$ and $\delta\chi_2$ end up being sharply peaked, but not at the correct value of zero.

theory of gravity (e.g. a massive graviton model with a deviation in ψ_2 only), one might find a violation of GR but draw the wrong conclusions about the nature of the deviation.

We can also look at the PDF for the hypothesis H_{123} , where the waveforms have $\delta\chi_1, \delta\chi_2, \delta\chi_3$ free; see Fig. 19. Once again, the peak is more or less at the correct value of $\delta\chi_3$, but we now have a much bigger spread. This, too, is as




FIG. 19. The posterior PDF for the injection with $\delta\chi_3 = 0.1$ as in the previous figure, but recovered with waveforms where $\{\vec{\theta}, \delta\chi_1, \delta\chi_2, \delta\chi_3\}$ are all free. The peak is near the correct value of $\delta\chi_3$ (with a median of 0.083), but this time the spread is considerably larger (with a standard deviation of 0.055), as we are trying to measure more parameters at the same time.

expected; parameter estimation degrades if one tries to measure too many parameters at once.

Finally, we look at the two-dimensional PDF for $\{\delta\chi_2, \delta\chi_3\}$, in the case where the waveform is the one that tests the hypothesis H_{123} ; see Fig. 20. Here, too, there is little to learn about the underlying nature of the deviation.

3. Example: A signal with $\delta\chi_3 = 0.025$

Let us consider an example with $(\mathcal{M}, \eta, D) = (1.14M_\odot, 0.242, 216 \text{ Mpc})$, $\delta\chi_3 = 0.025$, and a network SNR of 20.6. As expected, the Bayes factors for the component hypotheses against GR are considerably smaller than in the case of $\delta\chi_3 = 0.1$, but GR is still disfavored:

$$\begin{aligned} \ln B_{\text{GR}}^1 &= 11, & \ln B_{\text{GR}}^2 &= 12, & \ln B_{\text{GR}}^3 &= 12, \\ \ln B_{\text{GR}}^{12} &= 10, & \ln B_{\text{GR}}^{13} &= 11, & \ln B_{\text{GR}}^{23} &= 11, \\ \ln B_{\text{GR}}^{123} &= 11. \end{aligned} \quad (63)$$




FIG. 20 (color online). The 68% and 95% confidence contours of the two-dimensional PDF for $\{\delta\chi_2, \delta\chi_3\}$, still for an injection with $\delta\chi_3 = 0.1$, and with a waveform that has $\{\vec{\theta}, \delta\chi_1, \delta\chi_2, \delta\chi_3\}$ as free parameters. The maximum of the PDF is given by the black dot, and the injection values are represented by the red cross.

Also as expected, the signal is easily found by all of the model waveforms:

$$\begin{aligned} \ln B_{\text{noise}}^{\text{GR}} &= 186, & \ln B_{\text{noise}}^1 &= 197, & \ln B_{\text{noise}}^2 &= 198, \\ \ln B_{\text{noise}}^3 &= 198, & \ln B_{\text{noise}}^{12} &= 196, & \ln B_{\text{noise}}^{13} &= 197, \\ \ln B_{\text{noise}}^{23} &= 197, & \ln B_{\text{noise}}^{123} &= 197. \end{aligned} \quad (64)$$

As before, we look at the posterior PDF of $\delta\chi_3$ for the hypothesis H_3 , where only $\{\vec{\theta}, \delta\chi_3\}$ are allowed to vary: see the bottom plot in Fig. 21. The distribution is peaked near the correct value and stays away from zero; however,




FIG. 21. The posterior PDFs for $\delta\chi_1$ (top panel), $\delta\chi_2$ (middle panel), and $\delta\chi_3$ (bottom panel) for a single injection with $\delta\chi_3 = 0.025$ and network SNR of 20.6, recovered with waveforms where, respectively, $\{\vec{\theta}, \delta\chi_1\}$, $\{\vec{\theta}, \delta\chi_2\}$, and $\{\vec{\theta}, \delta\chi_3\}$ are free. Again, $\delta\chi_3$ is peaked at close to the correct value, with a median of 0.027 and a standard deviation of 0.0092, but both $\delta\chi_1$ and $\delta\chi_2$ are strongly peaked at incorrect values.

one should not expect the same to happen for lower-SNR sources.

Let us also look at the PDF for $\delta\chi_1$ when $\{\vec{\theta}, \delta\chi_1\}$ are free parameters, and of $\delta\chi_2$ when $\{\vec{\theta}, \delta\chi_2\}$ are free; see the top and middle plots of Fig. 21. As before, $\delta\chi_1$ and $\delta\chi_2$ are not peaked at the right values of $\delta\chi_1 = \delta\chi_2 = 0$.

4. Example: A signal with non-PN frequency dependence in the phasing

We now look at a signal with a nonstandard contribution to the phase, with a frequency dependence between 1PN and 1.5PN, as in Eq. (57). In the example we use here, $(\mathcal{M}, \eta, D) = (1.29M_\odot, 0.250, 208 \text{ Mpc})$, with a network SNR of 22.4. The Bayes factors for the component hypotheses against GR are

$$\begin{aligned} \ln B_{\text{GR}}^1 &= 91, & \ln B_{\text{GR}}^2 &= 93, & \ln B_{\text{GR}}^3 &= 89, \\ \ln B_{\text{GR}}^{12} &= 92, & \ln B_{\text{GR}}^{13} &= 91, & \ln B_{\text{GR}}^{23} &= 92, \\ \ln B_{\text{GR}}^{123} &= 91. \end{aligned} \quad (65)$$

Thus, also in this case, the GR hypothesis is very much disfavored, despite the fact that none of our model waveforms contains the anomalous frequency dependence which is present in the phase of the signal. We can also look at the Bayes factors against noise:

$$\begin{aligned} \ln B_{\text{noise}}^{\text{GR}} &= 148, & \ln B_{\text{noise}}^1 &= 239, & \ln B_{\text{noise}}^2 &= 241, \\ \ln B_{\text{noise}}^3 &= 237, & \ln B_{\text{noise}}^{12} &= 240, & \ln B_{\text{noise}}^{13} &= 239, \\ \ln B_{\text{noise}}^{23} &= 239, & \ln B_{\text{noise}}^{123} &= 239. \end{aligned} \quad (66)$$

It is interesting to look at the posterior PDF of $\delta\chi_3$ for the case where $\{\vec{\theta}, \delta\chi_3\}$ are allowed to vary (bottom panel of Fig. 22). The distribution looks uncannily like the analogous one for a signal with $\delta\chi_3 = 0.1$; see the bottom panel of Fig. 18. We can also look at the PDF of $\delta\chi_1$ in the case where $\{\vec{\theta}, \delta\chi_1\}$ are free parameters, and the PDF of $\delta\chi_2$ when $\{\vec{\theta}, \delta\chi_2\}$ are free; see the top and middle panels of Fig. 22. Here, too, there is an interesting resemblance to the analogous panels in Fig. 18, for an injection with $\delta\chi_3 = 0.1$.

We summarize below.

- (i) Here, the Bayes factors against GR clearly disfavor the GR hypothesis, despite the fact that none of our model waveforms has the kind of non-PN contribution to the phase that the signal contains.
- (ii) As before, the Bayes factors against GR are quite close to each other, and one cannot conclude much from them about the nature of the underlying deviation from GR.
- (iii) The Bayes factors against noise indicate that the signal will not be missed.
- (iv) The posteriors are quite similar to the ones where the deviation from GR is purely in the 1.5PN coefficient, with $\delta\chi_3 = 0.1$.




FIG. 22. The posterior PDFs for $\delta\chi_1$ (top panel), $\delta\chi_2$ (middle panel), and $\delta\chi_3$ (bottom panel) for a single injection with $\delta\chi_A = -2.2$ and network SNR of 22.4, recovered with waveforms where, respectively, $\{\vec{\theta}, \delta\chi_1\}$, $\{\vec{\theta}, \delta\chi_2\}$, and $\{\vec{\theta}, \delta\chi_3\}$ are free. The distribution of $\delta\chi_3$ has its median at 0.11 and a standard deviation of 0.017. Note the remarkable resemblance with Fig. 18, where the signal had $\delta\chi_3 = 0.1$. Also for $\delta\chi_1$ and $\delta\chi_2$, the distributions are very similar to the ones for a signal with $\delta\chi_3 = 0.1$.

Finally, let us look at the two-dimensional PDF for $\{\delta\chi_2, \delta\chi_3\}$ in the case where the waveform is the one that tests H_{123} ; see Fig. 23. Unlike the one-dimensional PDFs, here there is not much resemblance to the two-dimensional PDF for $\delta\chi_3 = 0.1$ (Fig. 20). Still, nothing much can be learned about the actual nature of the violation.

This shows, once again, that we will also be able to discern violations of GR of a kind that has no analog in the model waveforms. In the posteriors, these can even




FIG. 23. The 68% and 95% confidence contours of the two-dimensional PDF for $\{\delta\chi_2, \delta\chi_3\}$, for an injection with $\delta\chi_A = -2.2$, and with a waveform that has $\{\tilde{\theta}, \delta\chi_1, \delta\chi_2, \delta\chi_3\}$ as free parameters. The maximum of the PDF is given by the black dot. Here the distribution is quite different from the case with $\delta\chi_3 = 0.1$ (Fig. 20).

“masquerade” as deviations in one of the post-Newtonian coefficients, if, for example, one is only looking for a massive graviton with a waveform model that has $\{\tilde{\theta}, \delta\chi_2\}$ as free parameters.

E. A note on parameter estimation and multiple sources

We end this section with some cautionary remarks on the use of parameter estimation in testing GR. As we have seen in Sec. III D, model selection allows for combining the information from multiple sources to compute a single odds ratio, but for parameter estimation the situation is quite different.

In the examples presented in this paper, the deviations $\delta\chi_3$ and $\delta\chi_A$ were taken to be the same for all sources. Hence, in principle, we could have combined PDFs from multiple sources, simply by multiplying them, to arrive at more accurate measurements. However, we reiterate that, in reality, one cannot expect the deviations to be constant in this fashion; rather, they might vary from source to source, with dependence on the masses as well as whatever additional charges may be present in the correct theory of gravity. If it so happens that deviations in individual sources can go both ways, making positive or negative contributions to the overall phase depending on the parameters of the source, then a combined PDF may not show any significant deviation at all. Moreover, since most sources will have SNRs near threshold, it could be that only a few sources will allow for accurate parameter estimation; here we only showed PDFs for single, relatively “loud” sources with $\text{SNR} \gtrsim 20$. However, if there is significant dependence of the GR deviation on source parameters (unlike in the heuristic examples shown here), the particular deviation exhibited by the loudest source may not be representative. Hence, in contrast to model selection, parameter estimation alone does not provide a solid foundation to look for deviations from GR.

V. CONCLUSIONS AND FUTURE DIRECTIONS

As we showed at the beginning, it is possible for a signal to contain very significant deviations from general relativity while still being detectable with a template family of GR waveforms; this is the “fundamental bias” discussed in [35,53]. A violation of GR can, in principle, take any form, and the question which then presents itself is how to search for generic deviations.

We have developed a general method to search for deviations from general relativity using signals from compact binary coalescence events. To this end, we constructed an odds ratio $\mathcal{O}_{\text{GR}}^{\text{modGR}}$ for modifications to GR against GR, which is the posterior probability that there is a deviation from GR, versus GR being correct. This odds ratio can be written as a linear combination of Bayes factors $B_{\text{GR}}^{i_1 i_2 \dots i_k}$ for hypotheses $H_{i_1 i_2 \dots i_k}$, in each of which one or more of the phase parameters ψ_i is assumed to deviate from the GR value, without actually assuming any specific dependence on the frequency and/or physical parameters pertinent to a given theory. Since this includes hypotheses where only a single one of the ψ_i is non-GR, our method will be particularly well suited in low-SNR scenarios, which we expect to be in with the upcoming advanced detector network. Finally, information from multiple sources can easily be combined to arrive at an odds ratio $\mathcal{O}_{\text{GR}}^{\text{modGR}}$ for the “catalog” of all observed events.

The method we developed, applied to phase coefficients in inspiral waveforms, essentially addresses the question, do *one or more* of the coefficients differ from the values predicted by GR? This is in contrast with previous Bayesian analyses such as [35,55,56], where, effectively, the question being asked was limited to, do *all* of the additional free parameters introduced differ from their GR values? In addition to being better adapted to a low-SNR environment, the test proposed here is far more general.

In order to gauge how large a deviation might be detectable, we first considered signals with a constant fractional deviation in the 1.5PN coefficient ψ_3 . This coefficient is of particular interest, since it incorporates the so-called “tail effects” [20,21] (as well as spin-orbit coupling [22], although we did not consider spin here), which are not empirically accessible with binary pulsar observations and can only be studied through direct detection of gravitational waves. When considering catalogs of only 15 binary neutron star sources, we saw that a deviation in ψ_3 at the 10% level would easily be detectable. In fact, even a deviation at the few percent level can be discernible. This is confirmed by posterior PDFs for ψ_3 in the case where this is the only parameter that is assumed to deviate from its GR value.

We also considered a deviation in the phase with a frequency dependence that does not match any of the post-Newtonian terms, and hence is not present in any of the recovery waveforms that we used. More precisely,

we looked at signals whose phases have an additional contribution, with a frequency dependence in between that of the 1PN and 1.5PN terms (1.25PN). The magnitude of the deviation was chosen such that near $f \sim 150$ Hz, where the detectors are the most sensitive, the change in phase is roughly the same as the change caused by a 10% shift in the 1.5PN coefficient. The deviation was clearly detectable in the log odds ratios, the Bayes factors, and the posterior PDFs. We expect this to be an instance of a more general fact. Namely, even if there is a deviation in the phase which the model waveforms technically do not allow for, it will typically be observable, on the condition that near the “bucket” it causes a change in the phase that is on par with the effect of a shift in the (low-order) PN phase coefficients of more than a few percent.

In order to establish the basic validity and usefulness of our method for testing GR, we considered constant fractional deviations in ψ_3 , and a non-PN frequency dependence in the signal. However, even if there is a deviation in one or more of the PN phase contributions only, these may depend on $(\mathcal{M}, \eta)$ as well as whatever additional charges and coupling constants might be present. It would be of great interest to study the effects of more general deviations from GR on the odds ratio $\mathcal{O}_{\text{GR}}^{\text{modGR}}$, the cumulative Bayes factors $\prod_A^{(A)} B_{\text{GR}}^{i_1 i_2 \dots i_k}$ for the component hypotheses $H_{i_1 i_2 \dots i_k}$ against $\mathcal{H}_{\text{GR}}$, and the cumulative Bayes factors $\prod_A^{(A)} B_{\text{noise}}^{i_1 i_2 \dots i_k}$ against noise. Additionally, in the future one should consider deviations in the amplitude as well, and use model waveforms which have such freedom. *A priori* there is no reason not to use an arbitrary number of free coefficients in both the phase and the amplitude of recovery waveforms, the only limiting factor being computational time.

Should no deviation from GR be convincingly found through model selection, one will still be interested in *constraining* the theory. Advanced gravitational-wave detectors will then take us considerably beyond the binary pulsar observations. We recall that even the 1PN phase coefficient ψ_2 is not fully constrained by the latter, since the dissipative dynamics is only probed to leading PN order. As we have seen from PDFs, with a single compact binary coalescence event in Advanced LIGO/Virgo at SNR ~ 20 , the coefficient ψ_2 can be constrained to better than 2%, and the same is true of the 1.5PN coefficient ψ_3 . Mainly for computational reasons, we did not study constraints on the 2PN and higher-order terms, but it would clearly be of interest to see how well one can pin down the corresponding coefficients. We also draw attention to the results for the 0.5PN contribution, which in general relativity is identically zero. For simplicity, we restricted the mass range of our simulated sources to that of binary neutron stars. Given the very encouraging results, in future work the BHNS and BBH ranges should also be investigated.

The preliminary results presented here motivate the construction of a full data analysis pipeline for testing or constraining general relativity. To have a real chance of finding a deviation, much more sophisticated GR waveforms will have to be used, with the inclusion of merger and ringdown, higher harmonics both in the inspiral and ringdown parts, dynamical spins, and residual eccentricity. The development of such waveforms, with input from numerical simulations, is currently a subject of intense investigations [58,73–82]. We note that the method we presented is not tied to any particular waveform approximant. Moreover, the deformations need not be in the phase. Indeed, in the future one would presumably want to use time domain waveforms, for which it may be more convenient (and physically more appropriate) to introduce parametrized deformations directly in coefficients appearing in, e.g., a Hamiltonian used to evolve the inspiral part of the waveform. Irrespective of the parametrization, one would still be able to associate with it an exhaustive set of logically disjoint hypotheses $H_{i_1 i_2 \dots i_k}$.

Once sufficiently accurate waveforms are available, a test of GR on Advanced LIGO/Virgo could go as follows:

- (i) Starting from the best available GR waveforms, introduce parametrized deformations, leading to disjoint hypotheses like our $H_{i_1 i_2 \dots i_k}$, which together form $\mathcal{H}_{\text{modGR}}$.
- (ii) Use many injections of GR waveforms in real or realistic data, arranged into simulated “catalogs,” to investigate the distributions of the cumulative odds ratio $\mathcal{O}_{\text{GR}}^{\text{modGR}}$ as well as of the cumulative Bayes factors when GR is correct. Use these to set *thresholds* for the *measured* odds ratio and Bayes factors to overcome.
- (iii) Apply our method to the catalog of sources actually found by the detectors. If the measured cumulative odds ratio $\mathcal{O}_{\text{GR}}^{\text{modGR}}$ is below threshold, then there is no real reason to believe that a deviation from GR is present. The posterior PDFs for the free phase and amplitude coefficients in the model waveforms, taken from the highest-SNR sources, will provide (potentially very strong) constraints on these parameters.
- (iv) If the measured odds ratio is above threshold, then a violation of GR is likely. As we have seen, Bayes factors and PDFs can be misleading in trying to find out what the precise nature of the deviation may be. However, one may be able to follow up on the violation by again using our method, this time with waveforms with more complicated deformations and a larger number of free parameters, inspired by particular alternative theories of gravity, similar to what is done in pPE [53].

Thus, although much work remains to be done, we have the basics of a very general method for testing general relativity using compact binary coalescence events to be

detected by the upcoming Advanced LIGO and Advanced Virgo observatories.

ACKNOWLEDGMENTS

T. G. F. L., W. D. P., S. V., C. V. D. B., and M. A. are supported by the research programme of the Foundation for Fundamental Research on Matter (FOM), which is partially supported by the Netherlands Organisation for Scientific Research (NWO). J. V.'s research was funded by the Science and Technology Facilities Council (STFC), United Kingdom, Grant No. ST/J000345/1. K. G., T. S., and A. V. are supported by the Science and Technology Facilities Council (STFC), United Kingdom, Grant No. ST/H002006/1. The work of R. S. is supported

by the EGO Consortium through the VESF Grant No. EGO-DIR-41-2010. It is a pleasure to thank N. Cornish, B. R. Iyer, B. S. Sathyaprakash, and N. Yunes for their valuable comments and suggestions. We also thank the anonymous referee, whose comments helped us to greatly improve the presentation of the paper. The authors would like to acknowledge the LIGO Data Grid clusters, without which the simulations could not have been performed. Specifically, these include the computing resources supported by National Science Foundation Grants No. PHY-0923409 and No. PHY-0600953 to UW-Milwaukee. Also, we thank the Albert Einstein Institute in Hannover, supported by the Max-Planck-Gesellschaft, for use of the Atlas high-performance computing cluster.

-
- [1] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (W. H. Freeman and Company, New York, 1973).
 - [2] C. M. Will, *Living Rev. Relativity* **9**, 3 (2006), <http://www.livingreviews.org/lrr-2006-3>.
 - [3] B. S. Sathyaprakash and B. F. Schutz, *Living Rev. Relativity* **12**, 2 (2009).
 - [4] R. A. Hulse and J. H. Taylor, *Astrophys. J.* **195**, L51 (1975).
 - [5] M. Burgay, N. D'Amico, A. Possenti, R. N. Manchester, A. G. Lyne, B. C. Joshi, M. A. McLaughlin, M. Kramer, J. M. Sarkissian, F. Camilo, V. Kalogera, C. Kim., and D. R. Lorimer, *Nature (London)* **426**, 531 (2003).
 - [6] M. Kramer, I. H. Stairs, R. N. Manchester, M. A. McLaughlin, A. G. Lyne, R. D. Ferdman, M. Burgay, D. R. Lorimer, A. Possenti, N. D'Amico, J. M. Sarkissian, G. B. Hobbs, J. E. Reynolds, P. C. C. Freire, and F. Camilo, *Science* **314**, 97 (2006).
 - [7] B. Abbott *et al.*, *Rep. Prog. Phys.* **72**, 076901 (2009).
 - [8] F. Acernese *et al.*, *Classical Quantum Gravity* **21**, S385 (2004).
 - [9] F. Acernese *et al.*, *Classical Quantum Gravity* **25**, 184001 (2008); T. Accadia *et al.*, *Classical Quantum Gravity* **28**, 114002 (2011).
 - [10] H. Grote, for the LIGO Scientific Collaboration, *Classical Quantum Gravity* **25**, 114043 (2008); *Classical Quantum Gravity* **27**, 084003 (2010).
 - [11] G. M. Harry, for the LIGO Scientific Collaboration, *Classical Quantum Gravity* **27**, 084006 (2010).
 - [12] <http://www.ligo.caltech.edu/advLIGO/>.
 - [13] Advanced Virgo Baseline Design (Virgo Collaboration, Note No. VIR027A09, 2009, <https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0027A-09.pdf>).
 - [14] <http://www.cascina.virgo.infn.it/advirgo/>.
 - [15] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Classical Quantum Gravity* **27**, 173001 (2010).
 - [16] K. Kuroda and the LCGT Collaboration, *Classical Quantum Gravity* **27**, 084004 (2010).
 - [17] B. S. Sathyaprakash, for LIGO Scientific Collaboration, Scientific Benefits of LIGO-India, LSC Internal Report No. G1100991, 2011.
 - [18] L. Blanchet, *Living Rev. Relativity* **5**, 3 (2002), <http://relativity.livingreviews.org/Articles/lrr-2006-4/>.
 - [19] M. Maggiore, *Gravitational Waves. Volume 1: Theory and Experiments* (Oxford University Press, Oxford, 2008).
 - [20] L. Blanchet and B. S. Sathyaprakash, *Classical Quantum Gravity* **11**, 2807 (1994).
 - [21] L. Blanchet and B. S. Sathyaprakash, *Phys. Rev. Lett.* **74**, 1067 (1995).
 - [22] L. Blanchet, T. Damour, B. Iyer, C. M. Will, and A. G. Wiseman, *Phys. Rev. Lett.* **74**, 3515 (1995).
 - [23] C. K. Mishra, K. G. Arun, B. R. Iyer, and B. S. Sathyaprakash, *Phys. Rev. D* **82**, 064010 (2010).
 - [24] C. M. Will, *Phys. Rev. D* **50**, 6058 (1994).
 - [25] T. Damour and G. Esposito-Farèse, *Phys. Rev. D* **58**, 042001 (1998).
 - [26] P. D. Scharre and C. M. Will, *Phys. Rev. D* **65**, 042002 (2002).
 - [27] C. M. Will and N. Yunes, *Classical Quantum Gravity* **21**, 4367 (2004).
 - [28] E. Berti, A. Buonanno, and C. M. Will, *Phys. Rev. D* **71**, 084025 (2005).
 - [29] E. Berti, A. Buonanno, and C. M. Will, *Classical Quantum Gravity* **22**, S943 (2005).
 - [30] N. Yunes, F. Pretorius, and D. Spergel, *Phys. Rev. D* **81**, 064018 (2010).
 - [31] C. M. Will, *Phys. Rev. D* **57**, 2061 (1998).
 - [32] K. G. Arun and C. M. Will, *Classical Quantum Gravity* **26**, 155002 (2009).
 - [33] A. Stavridis and C. M. Will, *Phys. Rev. D* **80**, 044002 (2009).
 - [34] D. Keppel and P. Ajith, *Phys. Rev. D* **82**, 122001 (2010).
 - [35] W. Del Pozzo, J. Veitch, and A. Vecchio, *Phys. Rev. D* **83**, 082002 (2011).
 - [36] E. Berti, J. Gair, and A. Sesana, *Phys. Rev. D* **84**, 101501 (2011).

- [37] The designations “massive gravity” and “massive graviton” originate from [22], where actually the effect of a modified dispersion relation, or a wavelength-dependent propagation velocity, has been taken into account. While it is attractive to ascribe such a modification to a graviton mass, a modification of the dispersion relation can be a more general effect. Moreover, endowing the graviton with a mass introduces additional deviations from GR than a mere modified dispersion relation. See e.g. the original work [38] and the recent work [39] for a thorough discussion of the issues related to massive gravity models.
- [38] H. van Dam and M. J. G. Veltman, *Nucl. Phys.* **B22**, 397 (1970); V. I. Zhakarov, *JETP Lett.* **12**, 312 (1970).
- [39] C. de Rham, G. Gabadadze, and A. J. Tolley, *Phys. Rev. Lett.* **106**, 231101 (2011).
- [40] F. D. Ryan, *Phys. Rev. D* **56**, 1845 (1997).
- [41] E. Berti, V. Cardoso, and C. M. Will, *Phys. Rev. D* **73**, 064030 (2006).
- [42] S. A. Hughes, *AIP Conf. Proc.* **873**, 233 (2006).
- [43] L. Barack and C. Cutler, *Phys. Rev. D* **75**, 042003 (2007).
- [44] J. R. Gair, C. Li, and I. Mandel, *Phys. Rev. D* **77**, 024035 (2008).
- [45] I. Kamaretsos, M. Hannam, S. Husa, and B. S. Sathyaprakash, *Phys. Rev. D* **85**, 024018 (2012).
- [46] C. Van Den Broeck and A. S. Sengupta, *Classical Quantum Gravity* **24**, 1089 (2007).
- [47] S. Alexander, L. S. Finn, and N. Yunes, *Phys. Rev. D* **78**, 066005 (2008).
- [48] N. Yunes and L. S. Finn, *J. Phys. Conf. Ser.* **154**, 012041 (2009).
- [49] C. Sopuerta and N. Yunes, *Phys. Rev. D* **80**, 064006 (2009).
- [50] N. Yunes, R. O’Shaughnessy, B. J. Owen, and S. Alexander, *Phys. Rev. D* **82**, 064017 (2010).
- [51] K. G. Arun, B. R. Iyer, M. S. S. Qusailah, and B. S. Sathyaprakash, *Classical Quantum Gravity* **23**, L37 (2006).
- [52] K. G. Arun, B. R. Iyer, M. S. S. Qusailah, and B. S. Sathyaprakash, *Phys. Rev. D* **74**, 024006 (2006).
- [53] N. Yunes and F. Pretorius, *Phys. Rev. D* **80**, 122003 (2009).
- [54] N. Yunes and S. A. Hughes, *Phys. Rev. D* **82**, 082002 (2010).
- [55] N. Cornish, L. Sampson, N. Yunes, and F. Pretorius, *Phys. Rev. D* **84**, 062003 (2011).
- [56] S. Gossan, J. Veitch, and B. S. Sathyaprakash, *arXiv:1111.5819*.
- [57] L. Blanchet, G. Faye, B. R. Iyer, and S. Sinha, *Classical Quantum Gravity* **25**, 165003 (2008).
- [58] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **84**, 124052 (2011).
- [59] C. Van Den Broeck and A. S. Sengupta, *Classical Quantum Gravity* **24**, 155 (2007).
- [60] S. Vitale and M. Zanolin, *Phys. Rev. D* **82**, 124065 (2010).
- [61] K. S. Thorne, in *300 Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, England, 1987).
- [62] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
- [63] <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [64] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009).
- [65] D. Shoemaker (LSC), 2009, <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=2974>.
- [66] T. A. Apostolatos, *Phys. Rev. D* **52**, 605 (1995).
- [67] J. Skilling, *AIP Conf. Proc.* **735**, 395 (2004).
- [68] J. Veitch and A. Vecchio, *Phys. Rev. D* **78**, 022001 (2008).
- [69] J. Veitch and A. Vecchio, *Classical Quantum Gravity* **25**, 184010 (2008).
- [70] J. Veitch and A. Vecchio, *Phys. Rev. D* **81**, 062003 (2010).
- [71] B. Aylott *et al.*, *Classical Quantum Gravity* **26**, 165008 (2009).
- [72] B. Aylott, J. Veitch, and A. Vecchio, *Classical Quantum Gravity* **26**, 114011 (2009).
- [73] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, J. T. Whelan, B. Brügmann, P. Diener, J. Gonzalez, M. Hannam, S. Husa, M. Koppitz, D. Pollney, L. Rezzolla, L. Santamaría, A. M. Sintes, U. Sperhake, and J. Thornburg, *Classical Quantum Gravity* **24**, S689 (2007).
- [74] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaría, and J. Seiler, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [75] L. Santamaría, F. Ohme, P. Ajith, B. Brügmann, N. Dorband, M. Hannam, S. Husa, P. Mösta, D. Pollney, C. Reisswig, E. L. Robinson, J. Seiler, and B. Krishnan, *Phys. Rev. D* **82**, 064016 (2010).
- [76] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999).
- [77] A. Buonanno and T. Damour, *Phys. Rev. D* **62**, 064015 (2000).
- [78] T. Damour, P. Jaranowski, and G. Schäfer, *Phys. Rev. D* **62**, 084011 (2000).
- [79] T. Damour, *Phys. Rev. D* **64**, 124013 (2001).
- [80] A. Buonanno, Y. Chen, and T. Damour, *Phys. Rev. D* **74**, 104005 (2006).
- [81] R. Sturani, S. Fischetti, L. Cadonati, G. M. Guidi, J. Healy, and D. Shoemaker, *J. Phys. Conf. Ser.* **243**, 012007 (2010).
- [82] R. Sturani, S. Fischetti, L. Cadonati, G. M. Guidi, J. Healy, and D. Shoemaker, *arXiv:1012.5172*.