Improved binary pulsar constraints on the parametrized post-Einsteinian framework

Remya Nair[1,*] and Nicolás Yunes[1,2,†]

[1]*eXtreme Gravity Institute, Department of Physics, Montana State University,
Bozeman, Montana 59717, USA*
[2]*Department of Physics, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*

The parametrized post-Einsteinian formalism was developed to search for generic deviations from general relativity with gravitational waves. We here present constraints on this framework using Bayesian analysis of a set of binary pulsar observations. In particular, we use measurements of the Keplerian and post-Keplerian parameters of six different binary pulsar systems, and Markov-Chain Monte-Carlo exploration to calculate posteriors on the parametrized post-Einsteinian parameters and derive robust constraints. We find improvements of 1–2 orders of magnitude in the strength of constraints when combining all six observations, relative to what one can achieve when using only the double binary pulsar. We also find that the constraints are robust to any correlation with the binary's component masses. The bounds on the parametrized post-Einsteinian framework derived here could be used as a prior in future Bayesian tests of general relativity with gravitational wave observations.

I. INTRODUCTION

Einstein's theory of gravity, general relativity (GR), provides us with the essential tools to study the large scale structure of our Universe. Since the advent of GR, the field of cosmology has steadily been on the path to becoming a high precision science. With the spectacular measurements of the cosmic microwave background anisotropy by WMAP [1] and later Planck [2], we are indeed living in an era of precision cosmology. But, our standard theoretical paradigm for the Universe, the inflationary big bang model, is still incomplete, as the physics behind the late time acceleration of the Universe remains elusive and the initial conditions that led to inflation (accelerated expansion in the early Universe) are still unknown. The standard flat-ΛCDM model is also currently under scrutiny from the latest Planck findings that prefer a positive curvature at more than 99% confidence level [3]. All of this has suggested to some that a more robust description of the accelerating Universe is perhaps conceivable through a modification of GR.

Einstein's theory, however, has passed a plethora of tests with flying colors, from observations in the Solar System that explore the quasistationary weak field regime [4], to binary pulsar observations that probe gravity in the (quasistationary) strong field regime [5–7]. The latter are particularly constraining since they can be sensitive to gravity modifications that are suppressed in the Solar System. In particular, binary pulsar observations are affected by the backreaction of gravitational waves (GWs) on the orbital dynamics of binary systems [8], which is simply not accessible in the Solar System. The observation of such orbital decay requires extremely precise timing of the radio "pulses" that originate from the magnetic poles of the pulsars [9]. Given a set of measurements of arrival times, radio astronomers can then fit a timing model that can extract the (Keplerian) orbital parameters of the binary system. If the observed pulsar is sufficiently stable (as is the case for millisecond pulsars), then the timing model can typically fit every pulse over a time span of many months or even several years, allowing for the extraction of post-Keplerian parameters, like the orbital period decay.

The discovery of binary pulsars provided the perfect ground to test the strong field regime of the gravitational interaction. In a series of papers, Damour and Deruelle developed the "parametrized post-Keplerian" (ppK) framework to study the quasistationary regime of binary pulsars, and they parametrized all the observables obtained from pulsar timing [10,11]; this was later further developed by Damour and Taylor in the context of tests of modified gravity theories [12]. Given a theory of gravity, one can relate the ppK parameters to the Keplerian orbital parameters and the component masses of the binary system. Since the component masses are the only unknowns in the system that are not observed directly, it follows that a measurement of any two ppK parameters uniquely determines the two

---
[*]remya.nair@montana.edu
[†]nyunes@illinois.edu

masses, and the values for other ppK parameters can be predicted. If any other ppK parameter is measured, it provides a consistency test for the underlying theory of gravitation. Hence, the ppK framework allows for theory-independent tests of gravity using binary pulsar observations.

Let us consider the case of the double-pulsar J0737-3039A/B to appreciate the power of pulsar timing. This binary system was discovered in 2003 and soon enough a total of six ppK parameters were measured for the system. The relativistic precession of the orbit was known to better than 0.004%, back in 2006 [13], and consistent measurements would further decrease the uncertainties. All the ppK parameters are consistent with GR and this double pulsar provides the best test for the GR quadrupole formalism for GW generation [7]. Other pulsar systems have been used to test the validity of various assumptions of GR, one of which, for example, is the absence of dipolar emission. Binary pulsar systems with significantly different compactness (the ratio of their mass to their radius) can be used for testing dipolar radiation. In the double-pulsar system, the compactness is similar, since both objects are NSs, and hence dipolar radiation (if it exists) is strongly suppressed. On the other hand, PSR J1738 + 0333 is composed of a low mass white dwarf and a pulsar companion, so the compactness is wildly different and this system can be used to stringently constrain modified gravity theories that predict the existence of dipolar emission [14].

Binary pulsar observations, however, are no longer alone as probes of GR in the strong field. The latest additions to gravity probes are GWs from coalescing binaries, which help us inspect the highly dynamical, strongly curved spacetimes around massive objects. As with binary pulsars, a generic framework to perform model-independent tests of GR with GWs was developed in the last decade [15]. The parametrized post-Einsteinian (ppE) formalism proposes to augment the GR predictions for the GW models with additional (non-GR) parameters. GR and other modified theory predictions, such as scalar tensor theories [16], dynamical Chern-Simons gravity [17], Einstein-dilaton Gauss-Bonnet gravity [18], among others [19], can then be recovered by specific choices of these non-GR parameters. The LIGO-VIRGO Collaboration has observed GWs from multiple compact binary mergers and put some constraints on the allowed ppE deviations from GR [20–23]. These ppE constraints can be mapped to constraints on specific modified gravity theories, as was done in [24,25]. As the number of these events increases, we can expect the constraints to become more and more stringent [23].

The ppE modifications to the GWs emitted by binary systems do not just affect the observables of GW detectors, but they also modify ppK parameters. For example, ppE modifications to the GW amplitude or to the GW phase affect the rate of decay of the orbital period of any binary

system, since the latter depends on the energy flux carried away from the system by GWs. In [26] (from here on YH10), the authors showed explicitly how ppE modifications affect the rate of orbital period decay, and then they derived relational constraints on the ppE parameters as a function of the post-Newtonian (PN) order[1] at which they appear, given binary pulsar observations. These relational constraints were then evaluated using measurements of the orbital decay of the double binary pulsar PSR J0737-3039. To obtain these constraints, the component masses were fixed to their best-fit values, obtained by radio astronomers from the observations of Keplerian and post-Keplerian parameters of the binary while assuming GR is valid.

In this paper, we extend the study of YH10 by carrying out a Bayesian analysis on six binary pulsar observations to derive posteriors on the ppE parameters, marginalized over the component masses. One of our main results is shown in Fig. 1, where we plot the 95 percentile upper limit on the amplitude and phase ppE corrections, as a function of the PN order at which these corrections appear [see Eq. (1)]. These limits are obtained using Gaussian priors on the component masses (see p1 in Sec. III C), informed from the GR estimates which in turn are obtained from fitting pulsar timing data to a GR timing model. As compared to the results of YH10, also plotted in this figure, we see that our bounds are better by 1–2 orders of magnitude at all PN orders. This is because we use more observations than in YH10, and the improvement in the constraints are in spite of carrying out a Bayesian analysis in which the components are allowed to vary. These constraints are thus more stringent and more robust than those obtained previously and should be used as informative priors for future studies with binary pulsars or GW observations.

We want to make an important note here. In combining multiple binary observations and obtaining these constraints, we have made an implicit assumption that the non-GR correction is independent of the system parameters like masses and spins. If this is not the case, then combining measurements from different binary systems with different masses/spins (and hence different values for the non-GR corrections) will not be justified. In this scenario, one has to carry out a theory specific analysis, by incorporating the exact dependence of the non-GR corrections on the system parameters.

The remainder of this paper explains the details involved in obtaining these results and it is organized as follows. In Sec. II, we present a brief overview of parametrized tests of GR and how they can be used to obtain evidence or constrain deviations from GR in a model-independent manner. In Sec. III, we discuss our Bayesian scheme

---

[1]The PN approximation is one in which the field equations are solved as an expansion about small velocities and weak fields. An expansion of NPN order is one that is proportional to v^{2N} relative to its leading order, controlling factor [27].
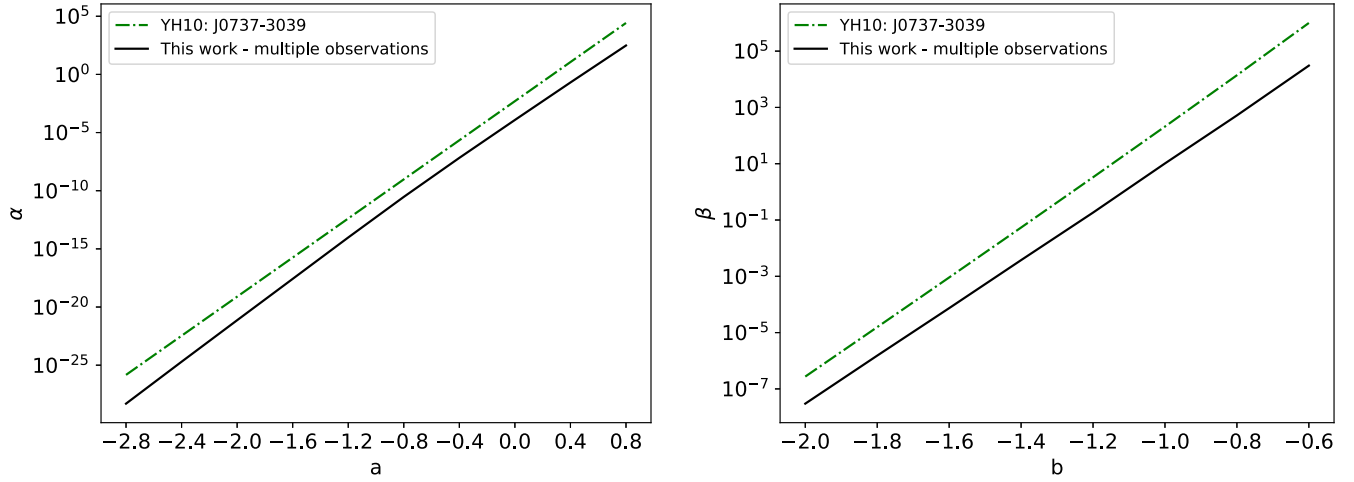
FIG. 1.   Joint 95 percentile upper limits on the ppE amplitude parameter α (left) and ppE phase parameter β (right) as a function of the ppE amplitude exponent a (left) and b (right), after marginalizing over the component masses, from the observation of six binary pulsars: J2222-0137, J1012-5307, J0348-0432, J0737-3039, J1909-3744, and J1738-0333 [13,28–36]. These limits are obtained using Gaussian mass priors informed from GR estimates (see p1 in Sec. III C). The region above the (black) solid line is ruled out by these observations. For comparison, we also include here the relational constraints derived in [26], which are 1–2 orders of magnitude weaker than those obtained here.

and how we obtain our upper limits on the ppE parameters. In Sec. IV, we conclude with a discussion of our results and implications. Henceforth, we used geometric units in which $G = 1 = c$ unless otherwise stated.

II. PARAMETRIZED TESTS OF GR

In this section, we introduce the ppE and the ppK frameworks and show how they lead to modifications to binary pulsar observables. The interested reader is referred to [15,19] and [10–12] for more details on these frameworks.

A. ppE framework

The ppE framework was introduced by Yunes and Pretorius to study deviations from Einstein equations, in a systematic and model-independent manner [15]. In this sense, the ppE framework is similar to the parametrized post-Newtonian formalism, which was built as an expansion about Minkowski space to describe weak field interactions in the Solar System, or the ppK framework mentioned in the previous section. The advantage of the ppE framework is that one can measure or constrain generic deviations from GR predictions instead of looking for some specific kind of deviation that may be predicted by a particular (alternative) theory.

The authors of [15] focused on GWs emitted by the quasicircular inspiral, merger, and ringdown of binary black holes. The deviations were encoded in the GW model of the detector response to an impinging GW through corrections in both the GW amplitude and the GW phase in the frequency domain. The standard ppE framework proposes modifications to the frequency-domain waveform

corresponding to the two GR polarizations, but since then, the framework has been expanded to allow for the presence of additional polarizations [37]. The frequency-domain corrections arise due to modifications in the binding energy of the binary system or its rate of change. The proposed enhancements correspond to deviations in the waveform amplitude and phase, which can be written as

$$\tilde{h}(f) = \tilde{h}_{\mathrm{GR}}(1 + \alpha u^a)\exp^{i\beta u^b}, \qquad u = (\pi \mathcal{M} f)^{1/3}, \qquad (1)$$

where $\tilde{h}(f)$ is the ppE modified (Fourier-domain) waveform, $\tilde{h}_{\mathrm{GR}}$ is the GR waveform, $\mathcal{M} = (m_1 m_2)^{3/5}/(m_1 + m_2)^{1/5}$ is the chirp mass, with component masses $m_{1,2}$, and f is the GW frequency. We see that in addition to the system parameters, like the masses, the modified waveform also depends on ppE parameters. The ppE parameters α and a are amplitude parameters, while β and b are ppE phase parameters.

The ppE parameters α and β control the magnitude of the GR deviations, and thus, they can be constrained from observations, while a and b determine the PN order at which the modification enters the waveform, and hence they characterize the type of physical modification. Note that $b = (k - 5)/3$ corresponds to correction appearing at the $(k/2)$PN order. A correction with $b = -5/3$ corresponds to a modification at leading PN order (0PN), and hence, corrections entering with $b < -5/3$ and $b > -5/3$ are understood as negative and positive PN corrections respectively. In general, α and β can depend on the system parameters, in addition to any fundamental (coupling) constants introduced by the alternative theory. GR predictions

can be recovered in this framework by simply setting $(\alpha, \beta) = (0, 0)$.

The ppE waveforms can be thought of as a result to generic modifications to the orbital chirping rate of the binary, $\dot{F} (= \dot{E}(dE_b/dF)^{-1})$. This rate can be modified if one modifies the conservative sector, i.e., the orbital binding energy E_b, or if one modifies the dissipative sector, i.e., the amount of energy lost from the binary systems due to emission of GW (or any other propagating degree of freedom), $\dot{E}$. Since both modifications lead to changes in the GW frequency and phase, the detection of a deviation from GR does not allow one to determine whether this deviation was a result of modifications to the conservative or the dissipative sector.

B. ppK framework

As discussed briefly in Sec. I, in order to study the strong field regime of binary pulsar observations, and the back reaction on the orbital dynamics due to GW, one can use the ppK parametrization. This provides a way to obtain theory-independent information by fitting the Keplerian and post-Keplerian parameters in a timing model. There are five Keplerian parameters that are usually employed to describe the orbital dynamics: the orbital period P_b, the orbital eccentricity e, the projected semimajor axis x, the longitude of periastron ω, and the time of periastron passage T_0. For any gravity theory, the post-Keplerian parameters can be written as functions of one or more of these Keplerian parameters, the component masses, and any extra parameters (new fundamental constant) that the theory may introduce.

Consider the post-Keplerian parameter $\dot{P}_b$ which is related to the orbital period damping due to GW emission. In GR, one can relate this post-Keplerian parameter to the Keplerian parameters P_b and e via [38]

$$\dot{P}_b^{\mathrm{GR}} = -\frac{192\pi}{5}\frac{m_p m_c}{M^2}\frac{1+73e^2/24+37e^4/96}{(1-e^2)^{7/2}}\frac{V_b^5}{c^5}, \quad (2)$$

where m_p and m_c are the pulsar and companion masses, respectively, M is the total mass, and $V_b \equiv (GMn_b)^{1/3}$, with $n_b \equiv 2\pi/P_b$. Observe that this post-Keplerian parameter (and in fact other post-Keplerian parameters in GR) is independent of the internal structure of the component objects, which may not be true in modified theories. For example, in mono-scalar-tensor theory, the rate of change of the orbital period depends on the difference between the effective scalar coupling α of the component objects [39,40] via

$$\dot{P}_b^{\mathrm{ST}} = -2\pi\frac{m_p m_c}{M^2}\frac{1+e^2/2}{(1-e^2)^{5/2}}\frac{\mathcal{V}_b^3}{c^3}\frac{(\alpha_p - \alpha_c)^2}{1+\alpha_p\alpha_c} + \mathcal{O}(\mathcal{V}_b^5/c^5), \quad (3)$$

where the subscripts p and c stand for the pulsar and its companion, and $\mathcal{V}_b = (G_*(1+\alpha_p\alpha_c)Mn_b)^{1/3}$.

Other post-Keplerian parameters include the change in the longitude of periastron $\dot{\omega}$, change in the orbital eccentricity, and the projected semimajor axes, $\dot{e}$ and $\dot{x}$. Besides these, there are post-Keplerian parameters which are important due to their relativistic effect that alters the timing of the arrival of the radio pulses: Einstein delay γ and the Shapiro time delay r, s. In this work, we focus only on $\dot{P}_b$, since this is the only parameter that has been related to the ppE corrections.

C. Mapping between ppK and ppE frameworks

As in the $\dot{P}_b$ case, other post-Keplerian parameters can also be related to Keplerian parameters, and then fitted with binary pulsar measurements, but for the purpose of this paper, we will focus on the orbital period decay $\dot{P}_b$. As we will show next, we can relate the post-Keplerian parameter $\dot{P}_b$ to the post-Einsteinian parameters α and β at any fixed PN order determined by the post-Einsteinian parameters a and b. Such a mapping is not so simple (and in fact, has not yet been derived) for other ppK parameters, because the ppE formalism has not yet been extended to eccentric orbits.

In order to put constraints on post-Einsteinian parameters using binary pulsar measurements, we first need to map the ppK parameters, in the present case $\dot{P}_b$, to the ppE corrections we introduced in Eq. (1). If the binding energy of the binary system is the same as in GR, the gravitational wave luminosity $\dot{E}$ can be related to the orbital decay $\dot{P}_b$ as

$$\frac{\dot{P}_b}{P_b} = \frac{3}{2}\frac{\dot{E}_b}{E_b} = -\frac{3}{2}\frac{\dot{E}}{E_b}, \quad (4)$$

where $\dot{E}$ is the energy lost due to GW emission (and the emission of any additional propagating degree of freedom), E_b is the binding energy of the system, and in the second equality we have used energy balance, i.e., the amount of binding energy lost due to GW emission is equal to (negative of) the amount of energy carried away by GWs (and any other propagating field).

As discussed in Sec. II, the ppE corrections cannot distinguish between modifications due to corrections to the binding energy or to the GW luminosity (or to both). In YH10, the authors considered modifications in the dissipative sector, i.e., in the energy carried away by the GW, while keeping the binding energy the same as in GR. Under this assumption, they found that $\dot{E}$ is modified as follows:

$$\dot{E} = \dot{E}_{\mathrm{GR}}\left[1 + \pi^2\mathcal{M}^2\beta b(b-1)u^{b-2}\left(\frac{d^2\Psi_{\mathrm{GR}}}{df^2}\right)^{-1}\right], \quad (5)$$

where Ψ_{GR} is the GW phase in GR. The quantity $\dot{E}_{\mathrm{GR}}$ for a binary system in an eccentric orbit is [41]

$$\dot{E}_{\mathrm{GR}} = -\frac{32}{5}\eta^2 \frac{M^5}{r_{12}^5}(1-e^2)^{-7/2}\left(1+\frac{73}{24}e^2+\frac{37}{96}e^4\right), \quad (6)$$

where r_{12} is the length of the semimajor axis. Using these expressions, and in the light of the energy balance equation [Eq. (4)], one can write the phase corrected orbital decay as

$$\frac{\dot{P}_b}{P_b} = \left(\frac{\dot{P}_b}{P_b}\right)_{\mathrm{GR}}\left(1+\frac{48}{5}\beta b(b-1)u^{b+5/3}\right), \quad (7)$$

thus relating the post-Keplerian parameter $\dot{P}_b$ to the post-Einsteinian parameter β at fixed b. Note that one can do a similar exercise of relating the ppE phase correction to $\dot{P}_b$ assuming that the correction comes from the conservative sector (changes in the binding energy of the orbit) as shown in [42].

The GW amplitude also depends on the rate of change of the orbital frequency, $\dot{F}$, and so we can also express GW luminosity in terms of the amplitude ppE parameters. YH10 followed this logic and found that $\dot{E}$ can be expressed in terms of amplitude ppE parameters as

$$\dot{E} = \dot{E}_{\mathrm{GR}}(1+\alpha u^a)^2. \quad (8)$$

This in turn can be used to obtain the amplitude corrected orbital decay as

$$\frac{\dot{P}_b}{P_b} = \left(\frac{\dot{P}_b}{P_b}\right)_{\mathrm{GR}}(1+2\alpha u^a), \quad (9)$$

where the GR prediction is

$$\left(\frac{\dot{P}_b}{P_b}\right)_{\mathrm{GR}} = -\frac{96}{5}\frac{\eta m^3}{r_{12}^4}(1-e^2)^{-7/2}\left(1+\frac{73}{24}e^2+\frac{37}{96}e^4\right).$$
$$(10)$$

As noted in YH10, observe that there is a one-to-one mapping between the amplitude and phase correction parameters, namely, $a = b+5/3$ and $\beta = 5\alpha/(48b(b-1))$ if the waveform amplitude and phase are modified due to the same mechanism.

Equations (7) and (9) show us how to use the measurements of post-Keplerian parameters obtained from binary pulsar observations to constrain post-Einsteinian parameters. In YH10, the authors obtained *relational constraints* on the magnitude of the ppE corrections (α and β) by relating them to the binary pulsar observational error δ, which is indicative of the accuracy with which the orbital decay is measured. To be clear, the authors of YH10 assumed that the observed ppK parameter $(\dot{P}_b)_{\mathrm{o}}$ matched its predicted value in GR $(\dot{P}_b)_{\mathrm{th}}$, up to some uncertainty δ (in units of Hz), i.e.,

$$\left(\frac{\dot{P}_b}{P_b}\right)_{\mathrm{o}} = \left(\frac{\dot{P}_b}{P_b}\right)_{\mathrm{GR}} + \delta. \quad (11)$$

Then, they set this equation equal to either Eq. (7) or (9), and solved for α and β to find a constraint in terms of Keplerian parameters and the uncertainty δ. Upper limits on the post-Keplerian parameters can then be obtained for different values of a and b (corresponding to deviations arising at different PN orders) assuming a value for the Keplerian parameters, which YH10 took to be the best-fit values obtained from the fits to a GR timing model.

III. BAYESIAN ANALYSIS AND PARAMETER ESTIMATION

In this section, we describe the dataset used in this work and give a brief overview of the Bayesian formalism. In YH10, relational constraints were obtained on the ppE amplitude and phase parameters, by fixing the component masses to the values predicted by assuming the validity of GR. Although this is a good approximation on the upper limits of these parameters, a more robust method to obtain constraints is to do a joint Bayesian parameter estimation study varying over all free parameters in the model, i.e., over the post-Einsteinian parameter and the component masses simultaneously. Note that we do not allow the eccentricity to vary because for the binary pulsars considered here, the orbits are all very nearly circular. We use a Bayesian framework for parameter estimation and study the effect of various mass priors on the estimation of the ppE corrections, considering the observations of multiple binary pulsars. We expect that the inclusion of more binary pulsar observations will enhance the constraints on ppE parameters, while allowing the masses to vary may have the opposite effect. In this section, we give a brief overview of the data used, the Bayesian scheme, and the Markov Chain Monte Carlo (MCMC) sampling approach we will employ to estimate the posterior distribution of masses and ppE corrections.

A. Data

We use six binary pulsar observations of the orbital decay to put constraints on the amplitude and phase corrections appearing at various PN orders. These are J2222-0137, J1012-5307, J0348-0432, J0737-3039, J1909-3744, and J1738-0333 [13,28–36]. All of these binaries have low eccentricities ($e < 0.1$) and consist of a pulsar with either an NS or a white dwarf companion. In order to obtain constraints on the post-Einsteinian parameter α or β using Eq. (7) and (9), we work with the measurements of the Keplerian parameters $\{P_b, e\}$ and the post-Keplerian parameter $\dot{P}_b$; the best-fit measurements and the uncertainties of these quantities for the different binary pulsars we considered are presented in Table I. The post-Einsteinian

TABLE I.   Keplerian and post-Keplerian parameters for the six binary pulsar systems used in this work. The figures in parentheses represent the estimated 1σ uncertainties in the last quoted digit. The eccentricities of all of these systems are very well measured, do not contribute significantly to the error budget, and are therefore ignored in our analysis. For easy reference, we also show the estimated values and 1σ uncertainty on the variable $(\dot{P}_b/P_b)$. Interested readers can refer to the cited papers to find the values of Keplerian and post-Keplerian measurements for these binaries.

System	P_b [days]	e	$\dot{P}_b$	$\dot{P}_b/P_b$	Ref.
J2222-0137	2.44576469(13)	3.8×10^{-4}	$-0.06(9) \times 10^{-12}$	$-2.839 \pm 4.259 \times 10^{-19}$	[28,29]
J1012-5307	0.60467271355(3)	1.2×10^{-6}	$-0.15(15) \times 10^{-15}$	$-2.871 \pm 2.871 \times 10^{-19}$	[30,31]
J0348-0432	0.102424062722(7)	2.6×10^{-6}	$-2.74(45) \times 10^{-13}$	$-3.096 \pm 0.508 \times 10^{-17}$	[32]
J0737-3039	0.10225156248(5)	8.7×10^{-2}	$-1.25(17) \times 10^{-12}$	$-1.417 \pm 0.019 \times 10^{-16}$	[13]
J1909-3744	1.533449474406(13)	1.1×10^{-7}	$-6(15) \times 10^{-15}$	$-4.528 \pm 0.132 \times 10^{-20}$	[33,34]
J1738-0333	0.3547907398724(13)	3.5×10^{-7}	$-25.9(3.2) \times 10^{-15}$	$-8.449 \pm 1.043 \times 10^{-19}$	[35,36]

parameter α or β and the component masses will be free parameters of the model (more discussion to follow).

Before proceeding, let us make a clarifying point here. Recall that in GR, the measurement of the two masses in the binary requires at least two ppK measurements, such as $\dot{P}_b$ and the Einstein time delay parameter γ. What would happen if we only used one ppK measurement, such as $\dot{P}_b$? If so, there would not be enough information in the data to break the $\{m_p, m_c\}$ degeneracy, and we would not be able to independently constrain the two masses. Therefore, one should expect that any attempt to constrain ppE amplitude parameters and the masses of the binary using only $\dot{P}_b$ should fail. As we will show later, this is not quite the case: even when one assumes nothing about the masses, one can still constrain the ppE amplitude parameters, albeit less strongly than when one assumes the masses are known to some degree. The masses themselves, however, are indeed very poorly estimated when one uses only $\dot{P}_b$ and assumes no mass priors.

B. Bayesian formulation

For an introduction to Bayesian theory and inference, we refer the interested reader to [43]. Our goal is to construct a posterior distribution on the set of all the model parameters $\boldsymbol{\theta} = \{m_p^1, m_c^1, m_p^2, m_c^2, \ldots p_{\mathrm{ppE}}\}$, where m_p^i and m_c^i correspond to the mass of the pulsar and its companion (for the ith observation), respectively, and p_{ppE} corresponds to α (for ppE amplitude corrections) or β (for ppE phase corrections). According to Bayes theorem, the posterior distribution for $\boldsymbol{\theta}$ in light of measurements can be written as

$$p(\boldsymbol{\theta}|D) = \frac{\mathcal{L}(D|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(D)}, \qquad (12)$$

where $p(\boldsymbol{\theta}|D)$ is the probability density for the parameters $\boldsymbol{\theta}$ given the data D, also termed the *posterior* probability density. The quantity $\mathcal{L}(D|\boldsymbol{\theta})$ is the *likelihood* function, which represents the probability of measuring D given the set of parameters $\boldsymbol{\theta}$. Finally, the quantity $p(\boldsymbol{\theta})$ is the *prior*

probability on $\boldsymbol{\theta}$, which represents our state of knowledge about these parameters before we analyze the data. The denominator $p(D) = \int \mathcal{L}(D|\boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}$ is an overall normalization constant, which is an important term to consider in model selection studies. In most parameter estimation schemes, one works with log-probability densities,

$$\ln p(\boldsymbol{\theta}|D) \propto \ln \mathcal{L}(D|\boldsymbol{\theta}) + \ln p(\boldsymbol{\theta}), \qquad (13)$$

and the aim is to find the set of parameters that maximize $\ln p(\boldsymbol{\theta}|D)$.

We now outline how we use this framework to construct the log-posterior as expressed in Eq. (13). We work with the observable $\dot{P}_b/P_b$; this is our data D. We assume a Gaussian model for this observable, i.e.,

$$\ln \mathcal{L}(D|\boldsymbol{\theta}) \propto -\frac{1}{2} \frac{((\dot{P}_b/P_b)_{\mathrm{o}} - (\dot{P}_b/P_b)_{\mathrm{th}})^2}{\sigma^2_{(\dot{P}_b/P_b)}}. \qquad (14)$$

Here $(\dot{P}_b/P_b)_{\mathrm{o}}$ and $(\dot{P}_b/P_b)_{\mathrm{th}}$ correspond to the observed value of the data and the theoretically predicted model, respectively. The latter is given by Eq. (7) for phase correction and Eq. (9) for amplitude correction. Note that the above expressions are used for single measurement estimations. Obtaining joint measurement estimates is straightforward with our assumption of Gaussianity. For joint measurements, the above expressions are modified to

$$\mathcal{L}(D^{\mathrm{J}}|\boldsymbol{\theta}) \propto \prod_{i=1}^{N} \mathcal{L}_i(D|\boldsymbol{\theta}),$$

or

$$\ln \mathcal{L}(D^{\mathrm{J}}|\boldsymbol{\theta}) \propto -\frac{1}{2} \sum_{i=1}^{N} \frac{((\dot{P}_b/P_b)_i^{\mathrm{o}} - (\dot{P}_b/P_b)_i^{\mathrm{th}})^2}{\sigma^2_{(\dot{P}_b/P_b)_i}}, \qquad (15)$$

where the superscript J signifies *joint* analysis. The ith individual likelihood $\mathcal{L}_i(D|\boldsymbol{\theta})$ can be read from Eq. (14), where i ranges from 1 to N, N being the total number of observations used for the parameter estimation.

In the above expressions, σ represents the error on the variable $\dot{P}_b/P_b$ and it can be obtained by some simple error propagation. In the absence of any covariances, the uncertainty on a variable $f(x, y)$ which depends on observables x and y satisfies

$$\sigma_f^2 = \left(\frac{\partial f}{\partial x}\right)^2 \sigma_x^2 + \left(\frac{\partial f}{\partial y}\right)^2 \sigma_y^2,$$

where σ_x and σ_y are the observational uncertainties on x and y, respectively. Hence, for our observable, we obtain

$$\sigma_{\dot{P}_b/P_b}^2 = \left(\frac{1}{P_b^o}\right)^2 \sigma_{\dot{P}_b}^2 + \left(\frac{1}{P_b^o}\right)^2 \left(\frac{\dot{P}_b^o}{P_b^o}\right)^2 \sigma_{P_b}^2.$$

C. Priors

Now the only piece in the puzzle left to discuss is the prior appearing on the right-hand side of Eq. (13). We begin by discussing the priors on the component masses, and then proceed with a discussion of our priors on the ppE parameters.

1. Priors on masses

We will consider three cases that differ from each other based on our assumptions about the component masses. This will allow us to study the effect of the mass prior probability on our marginalized posteriors. For the different cases we studied in this work, the parameter space has different dimensions. For easy referencing, we tabulate them in Table II. Note also that, in addition to the mass priors mentioned in this subsection, we further restrict the pulsar mass and the companion mass in the ranges $m_p \in (0.5, 3)$ M_{solar}, $m_c \in (0.05, 3)$ M_{solar}.

p1—Priors based on GR estimates.—In YH10, the authors chose the component masses of the binary to be exactly equal to the best-fit values obtained by fitting binary pulsar data to a timing model to extract ppK parameters, and then using the GR expressions for the ppK parameters as a function of the Keplerian parameters and the component masses to infer the masses. We will refer to component masses obtained in this way as best-fit masses assuming GR. Such a choice is reasonable because no strong field

TABLE II. Dimensionality of the parameter space for the different cases studied in this work.

Priors	Single measurement	Joint measurements
Fixed-mass prior	1	1
Gaussian mass prior	3	13
Gaussian mass-ratio prior	3	9
Uniform mass prior	3	13

observations have so far shown any evidence against GR. To be completely theory agnostic, however, one should treat the masses as free parameters to be fixed via a parameter estimation scheme.

p1a As a starting point (and to compare our results with those in the literature), one of our prior mass choices will be to assume that the component masses are given by their best-fit values assuming GR, i.e., the prior on the masses are delta functions centered at the best-fit values assuming GR. This prior is the same as that chosen in YH10, and we will refer to it as a *fixed-mass prior*. With this prior, the parameter set reduces to $\boldsymbol{\theta} = \{p_{ppE}\}$, and, given a value of the ppE exponent a (b), we have a very simple parameter estimation problem with only one parameter α (β) to maximize the log-posterior over. This calculation can be easily performed on a single observation, or on N observations through a joint parameter estimation study.

p1b We will also find it convenient to relax this fixed-mass prior to study its effect on our constraints on GR. To do so, we choose a mass prior that is a Gaussian centered at the best-fit values assuming GR, with width given by the 1σ uncertainty in the GR estimate. With this *Gaussian mass prior*, for a single observation, the parameter set reduces to $\boldsymbol{\theta} = \{m_p, m_c, p_{ppE}\}$, and given a value of the ppE exponent a or b, we then maximize the log-posterior [Eq. (13)], over a three-dimensional parameter space. The same calculation can then be repeated for N observations through a joint analysis.

p2—Prior on mass ratio.—Instead of imposing a prior on each of the component masses around their GR values, we can make use of other measurements that constrain the masses. If the companion to the pulsar is bright enough for optical spectroscopy, the mass ratio $R = m_p/m_c$ can be determined through combining the Doppler shifts in the spectral lines with the timing observations of the pulsar. This is also true for double-pulsar systems for which the orbits of both NSs can be measured simultaneously. This is because for any Lorentz-invariant theory of gravity, the relative size of the orbits is related to the mass ratio of the system (up to first PN order): $R \equiv m_p/m_c = a_c/a_p$, and hence one can estimate the mass ratio of the binary.

One choice of prior on the masses is then to assume a Gaussian distribution on the mass ratio R, with width given by the 1σ uncertainty in the measurement. We refer to this prior as a Gaussian mass-ratio prior, and we note that out of the six observations we use as data, only four binary systems (J1012-5307, J0348-0432, J0737-3039, and J1738-0333) have separate measurements of the mass ratio, and so we only use these data for this prior. The parameter set is still $\boldsymbol{\theta} = \{m_p, m_c, p_{ppE}\}$ for a single observation, but the log-posterior now depends on the log-Gaussian prior on R per Eq. (13), so that only a certain region in the $\{m_p, m_c\}$

parameter subspace is favored when exploring the likelihood.

p3—Uniform prior.—The final case we consider is to treat the masses as completely free with uniform linear priors. This is the most ambitious case where we try to estimate $(2N + 1)$ parameters from N observations, i.e., N m_p parameters, N m_c parameters, and 1 ppE parameter, so that $\boldsymbol{\theta} = \{m_{c,1}, m_{c,2}, \ldots, m_{c,N}, m_{p,1}, m_{p,2}, \ldots, m_{p,N}, P_{\mathrm{ppE}}\}$. In this case, we do not expect to find reasonable constraints on all parameters, because a single ppK measurement per binary pulsar observation is not enough to estimate both component masses. As we will show later, however, even though the masses cannot be well estimated, the ppE amplitude parameter can still be constrained, albeit not as strongly as when one assumes strong priors on the masses.

Given these three choices of prior, which one is valid for a test of GR? As we will explain later, the deterioration of ppE constraints that occurs when we use priors (p2) and (p3) is an artifact of only using measurements of $\dot{P}_b$ as our data. In reality, each of the pulsars we consider also has measurements of other ppK parameters. If we were to include such measurements as data in our analysis, they would limit the range of allowed component masses to what one obtains with the Gaussian mass prior in (p1b). Results obtained with this latter prior are thus the ones we quote in the introduction of this paper. Of course, we could bypass this entire discussion of priors by using (p3) and including all measured ppK parameters. This, however, is not yet possible because a ppE mapping of these parameters requires the extension of the ppE framework to eccentric binaries; such an analysis must therefore be relegated to future work.

2. Priors on ppE corrections

For the amplitude and phase corrections, we consider a wide range of values for the ppE exponents a and b: $-2.8 < a < 0.8$ and $-2.0 < b < -0.6$. The priors on the amplitude and phase ppE parameters, α and β, are assumed to be *uniform* in $\log \alpha$ and $\log \beta$, respectively. The boundaries of the range are informed from YH10 estimates. We make the pessimistic assumption that the estimates in YH10 are optimistic and set our upper boundary on the ppE parameters to be slightly worse than those obtained by YH10. As an example, consider the amplitude ppE parameter when $a = -2$ for which YH10 obtains the upper limit $\alpha \lesssim 10^{-21}$. For our analysis, we would then set the prior boundaries on $\log \alpha$ to be $-40 < \log \alpha < -18$. We use similar considerations to set the boundaries for the phase correction parameter β. We then perform the optimization exercise for each discrete value of a (b) and calculate the *marginalized* 95 percentile upper limit on the ppE corrections α (β). We further checked that these upper limits do

not change by any appreciable amount if we modify the prior limits that are set on the ppE parameter $\log \alpha$ ($\log \beta$).

One may wonder how much the choice of the priors influences the constraints. As mentioned, in this study we have used uniform priors on $\log \theta_{\mathrm{ppE}}$. An alternative choice, frequently made when determining upper limits on parameters, is to use uniform priors on θ_{ppE}. Using a uniform
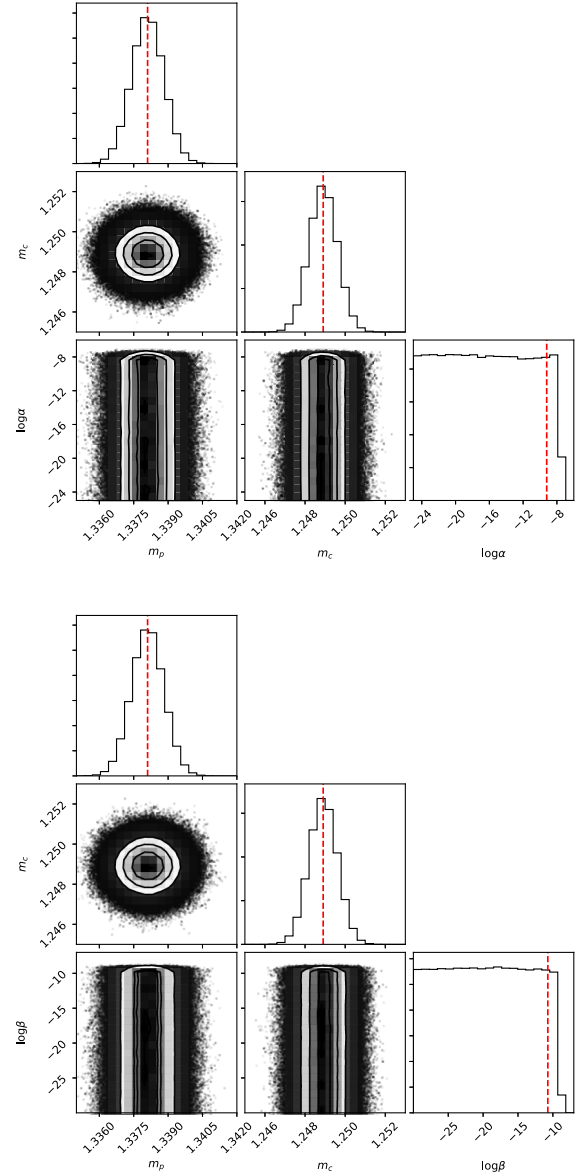

FIG. 2. Constraints on the pulsar and companion masses (m_p and m_c), and on the (log) ppE amplitude parameter $\log \alpha$ (top) and the (log) ppE phase parameter $\log \beta$ (bottom). These constraints are obtained by a Bayesian parameter estimation analysis using the measurements of PSR J0737-3039, assuming the corrections occur at -1PN order ($a = -2/3$ or $b = -7/3$). We also show the mean values ("best fits") for the component masses with a dashed (red) line. For the ppE corrections (corner right plot in both panels), the dashed (red) line corresponds to the 95 percentile upper limit.

prior on the logarithm of a parameter amounts to saying that any order of magnitude for the parameter θ_{ppE} is equally likely, whereas using a prior on the parameter itself amounts to saying that any value of the parameter θ_{ppE} is equally likely. But, in practice, a uniform prior on the logarithm of a parameter will also give more weight to lower values of the parameter when exploring the posterior distribution [Eq. (13)], since in this case $p(\theta_{\text{ppE}}) \propto 1/\theta_{\text{ppE}}$. Given that the ppE corrections are to be understood as small deformations from GR, a uniform prior on the logarithm of the ppE parameter seems to be the appropriate choice. If one insisted on using a uniform prior on the ppE parameter itself, the bound on the ppE parameter would deteriorate by roughly 1 order of magnitude, leading to the largest deterioration (by a factor of 80) when $b = -0.6$.

D. Sampling the posterior

One can reconstruct the posterior distribution from Eq. (12) in a brute-force way, by creating a simple grid and evaluating this function at each grid point to find the parameter values for which the function is maximum. It is obvious that this exercise would become computationally very expensive (almost impossible) in cases where we have a complicated likelihood function or more than a few parameters to sample. Instead, one can make use of MCMC sampling algorithms, which generate a list of samples for $\boldsymbol{\theta}$ according to the posterior distribution $p(\boldsymbol{\theta}|D)$. We use the PYTHON package *emcee* to carry out the exploration of the posterior distribution [44]. emcee is an implementation of Goodman and Weare's Affine Invariant MCMC ensemble sampler [45], which has an advantage over the standard Metropolis Hastings algorithms when the scales of covariances of the target distribution are not well known.

IV. RESULTS AND IMPLICATIONS

In this section, we present the results of our Bayesian analysis. Note that, as mentioned earlier, since we have only used one ppK measurement, $\dot{P}_b$, we can expect to find reasonable mass estimates only when we use additional mass information. This includes either using GR estimates (p1) or using measurements of the mass ratio (p2). This is true even if we do a *GR only* test case without any ppE corrections, because at least two ppK measurements are required to uniquely identify the component masses.

We begin by considering constraints from a single observation, the double binary pulsar J0737-3039, so that we can compare our results to those in YH10. We then consider constraints on ppE parameters from six different binary pulsars. We conclude by computing joint constraints from the stacking of all six binary pulsar observations.

A. J0737-3039 and comparisons to YH10

We start by first comparing our results with the work of YH10, where the authors used observations for the pulsar J0737-3039, fixed the component masses to their GR values, and obtained relational constraints on the ppE parameters. Our first step is then to follow their assumptions and to choose a fixed-mass prior (p1a), so as to obtain the upper limit on the post-Einsteinian phase parameter β, focusing on the correction appearing at -1PN order as a test case. As discussed earlier, we have only one free parameter to estimate in this case: the ppE phase correction β. We find that our 95 percentile upper limits are approximately $\beta \lesssim 10^{-10}$, ($\log \beta < -10.725$) which is roughly the same as the relational constraints obtained in YH10, namely, $\beta_{\text{YH10}} \lesssim 10^{-10}$. We find that this agreement with YH10 extends to other PN orders as well when comparing YH10 to our 95 percentile upper limits.
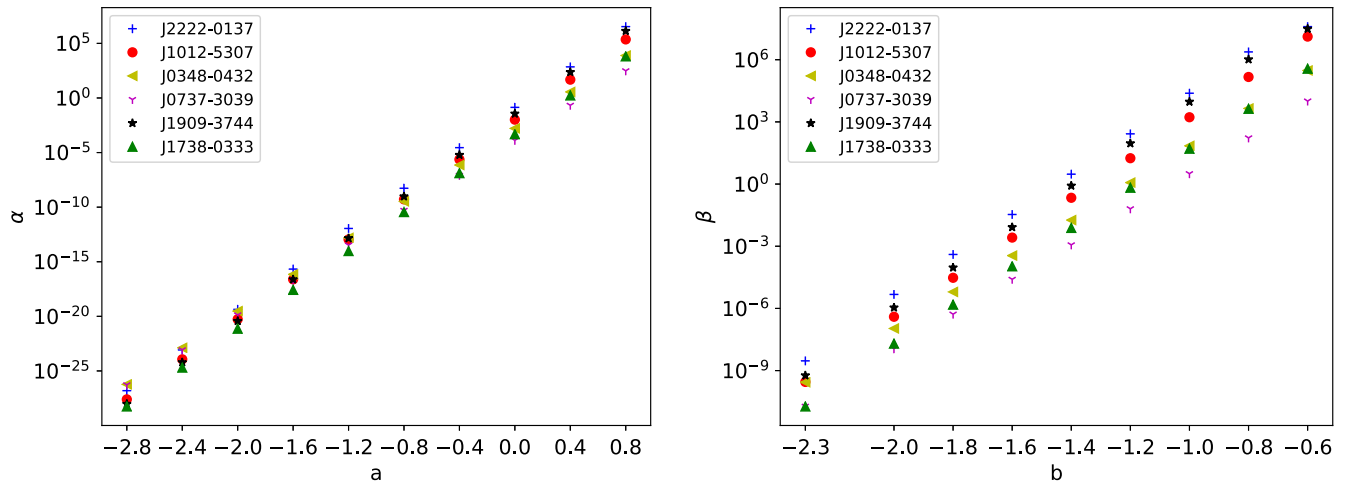

FIG. 3.   95 percentile upper limits on ppE corrections obtained from the different binary pulsar observations used in this work. The left and right panels correspond to amplitude and phase corrections, respectively. This figure shows that the best constraints come from observations of J1738-0333 and the double-pulsar J0737-3039 (for more discussion, see Sec. IV B).
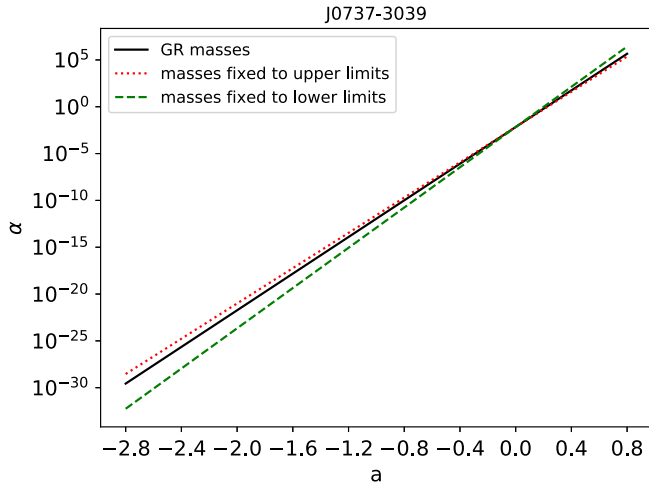
FIG. 4. Constraints on the amplitude correction α, obtained from the measurements of the binary pulsar J0737-3039. Different curves correspond to fixing the component masses to different values (see Sec. IV B for details). This figure shows that observations of lower mass systems can better constrain negative PN order modifications to GR.

Next, we relax our assumptions on the component masses and analyze the effect of using Gaussian mass priors (p1b), informed from their GR values. In this case, we need to estimate three parameters: m_p, m_c, and α or β. As an example case, we again focus on the ppE corrections appearing at -1PN order and show our constraints in Fig. 2. The 95 percentile upper limits in this case are $\alpha \lesssim 10^{-9}$ ($\log \alpha < -9.228$) and $\beta \lesssim 10^{-10}$, ($\log \beta < -10.720$), which is essentially the same as what

we find with the fixed-mass prior. As discussed earlier, pulsar timing observations are high precision measurements and they provide very precise estimates of the component masses when multiple ppK parameters are measured. Due to the extremely small error estimates, using a Gaussian prior is almost the same as using the fixed-mass prior, as is evident from the mass estimates seen in the corner plots in Fig. 2.

B. Individual measurements

We now study all the binary pulsars individually to see which measurements provide the best constraints on the ppE corrections using the fixed-mass prior case (p1a). This choice of prior is the simplest to consider when carrying out a comparative analysis, since we only have a single parameter to estimate, but as we saw above, the resulting constraints are essentially the same as what one would obtain with a Gaussian mass prior (p1b). The 95 percentile upper limits on α and β are shown in Fig. 3 as a function of a and b. The two most constraining measurements are from the millisecond pulsar J1738-0333 and the double-pulsar J0737-3039.

One can understand this by looking at the relation between the ppE parameters and the measurable quantities through Eqs. (7) and (9). This is how YH10 originally obtained their relational constraints. If we assume that the measurements of $\dot{P}_b/P_b$ are very close to their GR values (which is a very reasonable assumption), then one can roughly relate α (or β) to the accuracy, δ, with which these Keplerian and post-Keplerian observables are measured, namely,
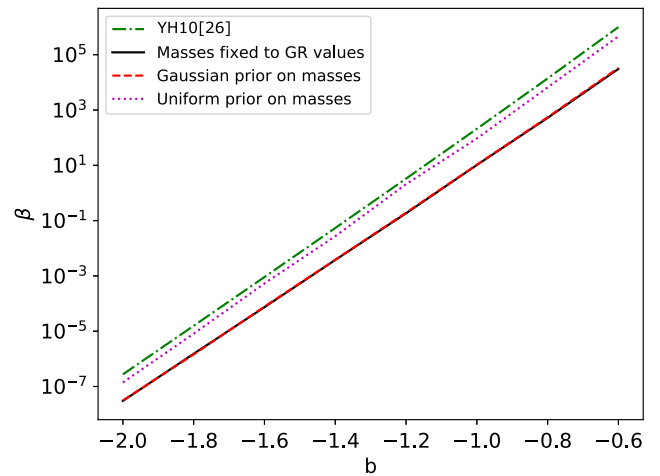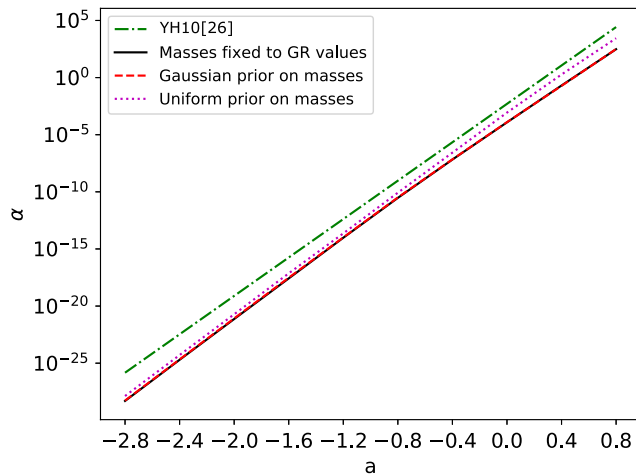

FIG. 5. Joint 95 percentile upper limits on the ppE amplitude parameter α (left) and ppE phase parameter β (right) as a function of the ppE amplitude exponent a (left) and b (right), similar to Fig. 1. In this plot, in addition to the black curve of Fig. 1, we show constraints obtained using other mass priors. The figure shows that the joint constraints are about 1 or 2 orders of magnitude stronger than those found in YH10 and independent of whether one fixes the component masses (solid black curve) or uses a Gaussian prior (dashed red curve), informed from the GR estimates. The constraints deteriorate if we use uniform priors on the component masses (dotted magenta curve) because we focus on constraints derived using a single ppK measurement ($\dot{P}_b$).

$$|\alpha| \le \left(\frac{P_b}{\dot{P}_b}\right)_{\rm GR} \left(\frac{\delta}{2}\right) u^{-a}, \qquad (16)$$

$$|\beta| \le \left(\frac{P_b}{\dot{P}_b}\right)_{\rm GR} \left(\frac{5\delta}{48|b||b-1|}\right) u^{-b-5/3}, \qquad (17)$$

where δ is the uncertainty on the observed $(\dot{P}_b/P_b)$, while $(\dot{P}_b/P_b)_{\rm GR}$ is the GR prediction. We find that these relational constraints are the strongest for J1738-0333 and J0737-3039 because the combination $\delta \times (P_b/\dot{P}_b) \propto \delta \times r_{12}^4/(\eta m^3)$ is the smallest for these binaries. That is,
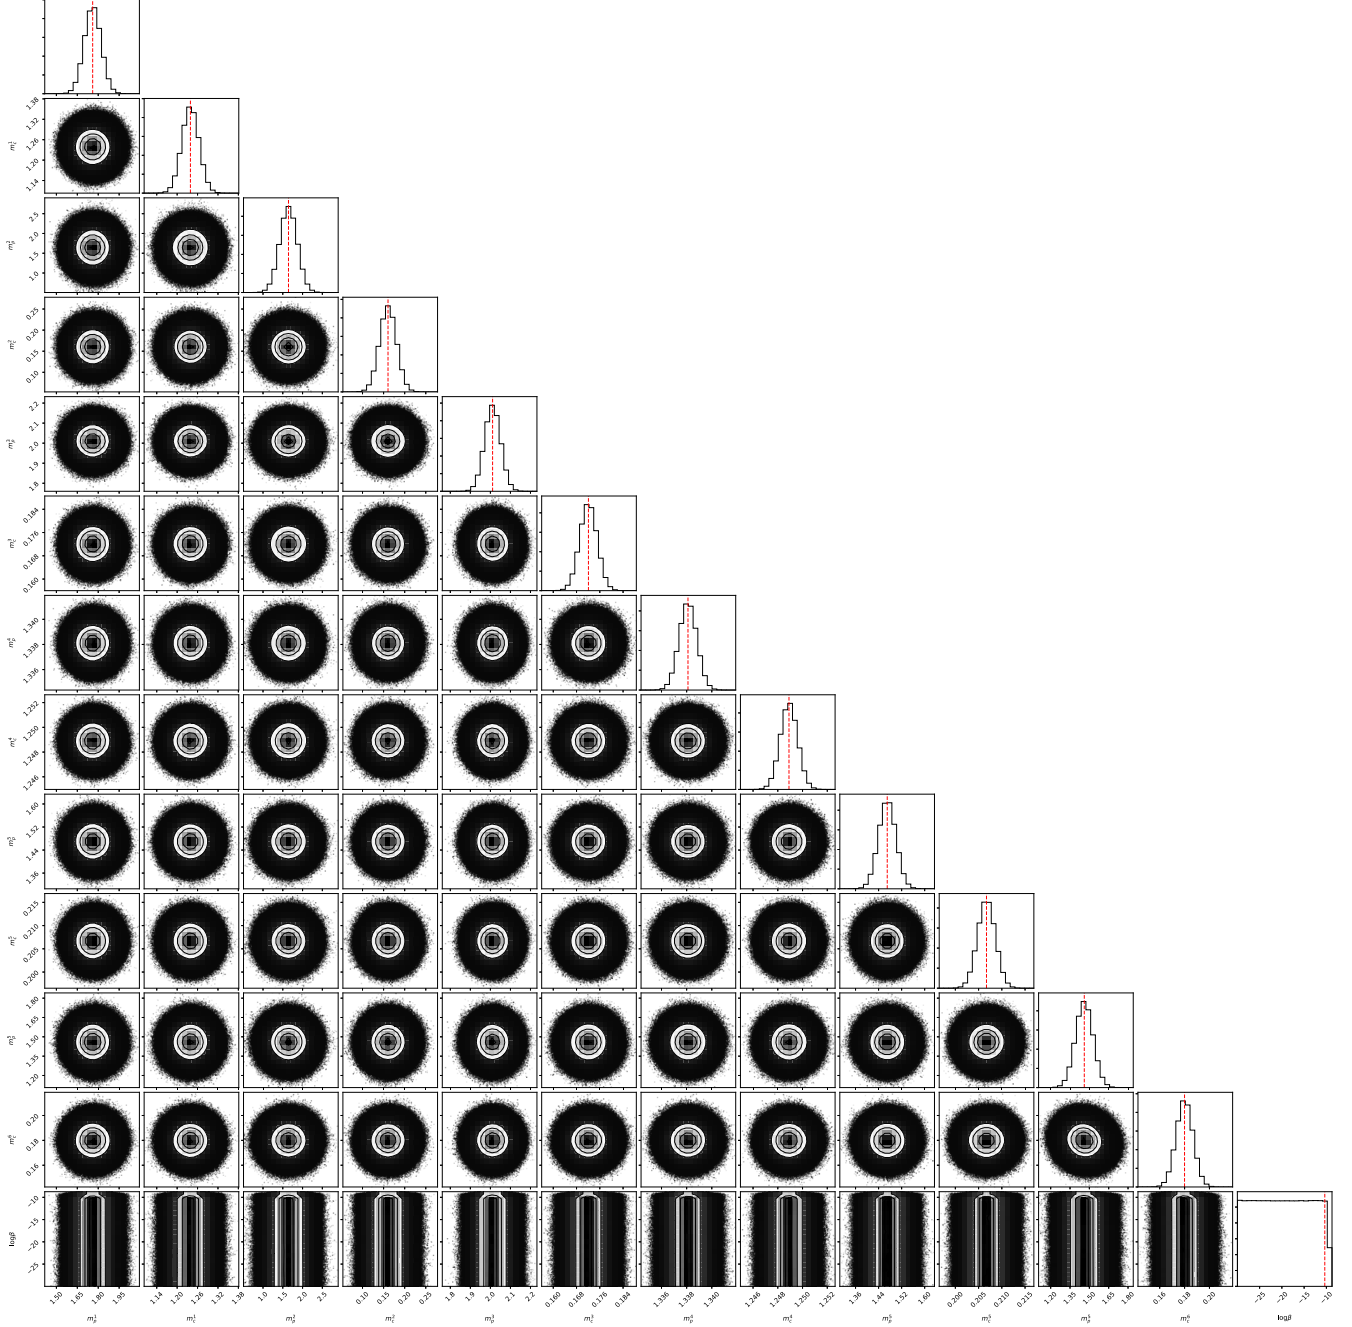

FIG. 6. Corner plot showing the joint estimation of mass parameters and the phase ppE corrections appearing at -1PN order. The priors on the masses are set using their GR values (Gaussian mass prior), and the prior on the ppE correction is uniform (see Sec. III for details). We also show the mean values (best fits) for the component masses with a dashed (red) line. For the ppE corrections (corner right plot), the dashed (red) line corresponds to the 95 percentile upper limit. Observe that the estimates on the component masses are very tight owing to the use of GR estimates as priors and the ppE correction does not appear to be correlated with any of the mass estimates.

these two binaries are not just the two most relativistic ones, but they also have very well-measured ppK parameters.

As a further analysis, let us now study how the constraints are affected by our knowledge of the component masses. For this study, we estimate constraints on the amplitude ppE parameter while fixing the component masses to the least possible/largest possible allowed values. We show the 95 percentile upper limits obtained for one of the most informative binaries, J0737-3039, in Fig. 4. From this figure, we observe that the binaries providing the best constraints on negative (positive) PN corrections have a lower (higher) total mass for fixed P_b and $\dot{P}_b$. This can be understood by looking at the nature of the amplitude correction: αu^a. For fixed frequency, lower mass values have lower values of $u\,(=(\pi\mathcal{M}f)^{1/3})$, which means that if $a < 0$ the ppE amplitude corrections are larger than in the high mass case. Hence, the low mass cases allow a larger modification and are therefore easier to constrain by these measurements. The situation reverses when $a > 0$ which is also evident in Fig. 4, although the effect is less pronounced since the positive a range is different than the negative a range in the figure. Similar conclusions can be drawn from Fig. 3.

C. Joint measurements

We now use the whole set of six binary pulsar measurements to obtain joint constraints on our ppE modifications. We discuss below the results from these studies for each mass prior case considered here.

1. Joint results using priors based on GR estimates

The 95 percentile upper limits obtained from the joint analysis using fixed-mass priors (p1a) are shown as black curves in Fig. 5. Similar results hold when we use the

Gaussian mass prior (p1b), shown by a red dashed curve in Fig. 5. As expected, the joint constraints are better than constraints with single observations and with the relational method of YH10 by roughly 1–2 orders of magnitude. At all PN orders, we see improvements on the upper limits of both the amplitude and phase corrections.

As an example case of what goes into the analysis we performed, we show in Fig. 6 the results of our Bayesian study for a ppE phase correction at -1PN order ($b = -7/3$) using a Gaussian mass prior. We do not observe any significant changes in the corner plot when we look at ppE amplitude corrections, or ppE corrections at other PN orders. The corner plot shows that the masses do not present significant correlations, with each marginalized posterior considerably Gaussian. Again, this is due to the very high precision with which the masses are estimated by the pulsar timing analysis. The construction of Fig. 5 required the calculation of corner plots like that shown in Fig. 6 at each PN order sampled.

2. Joint results using priors on the mass ratio

As an example of how the marginalized posteriors changes when we use the Gaussian mass-ratio prior, Fig. 8 shows the corner plot for a ppE phase correction at -1PN order ($b = -7/3$), for the binary J1738-0333. In this figure, we compare the Gaussian mass prior (black contours) to the Gaussian mass-ratio prior (gray contours), and we observe that the marginalized posteriors on the component masses have a larger uncertainty than when we use GR informed priors on the masses. This widening in the joint posterior is responsible, in part, for the slight deterioration of the ppE constraint. Similar results are obtained for all four binaries for which mass-ratio measurements are available (see Sec. III C).
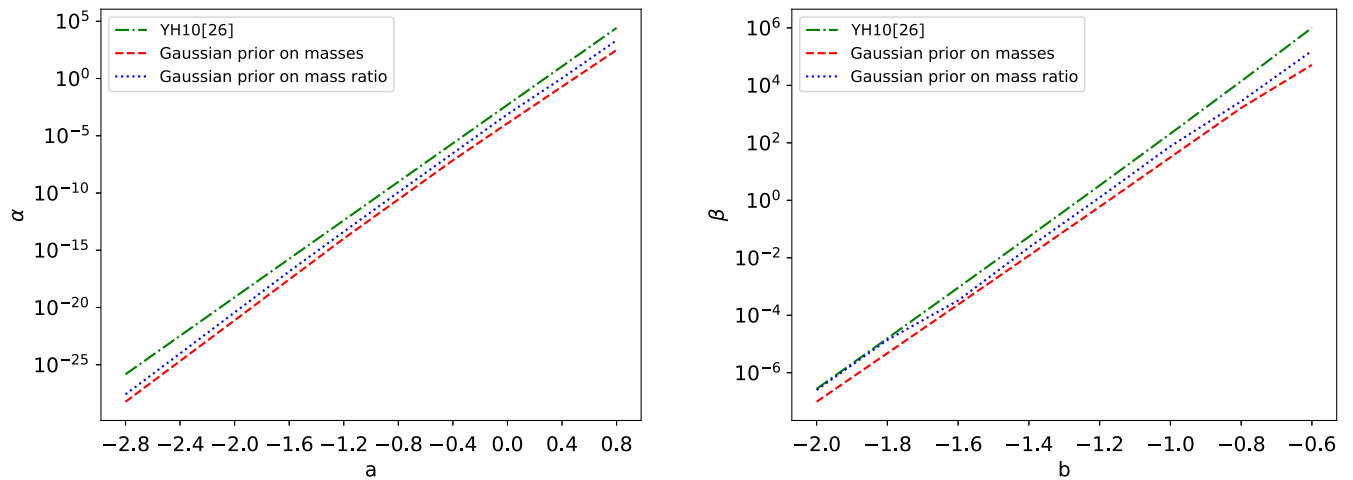

FIG. 7. Joint 95 percentile upper limits on the ppE amplitude parameter α (left) and ppE phase parameter β (right) as a function of the ppE amplitude exponent a (left) and b (right), similar to Fig. 1. In this plot, however, we compare constraints obtained using only the four binaries for which mass-ratio measurements are available. The figure shows that the joint constraints worsen if we use a prior on the mass ratio (dotted blue curve) instead of priors on the component masses derived assuming GR (dashed red curve).

Next, we do the joint analysis using all four binaries. The 95 percentile upper limits using the Gaussian mass-ratio prior (p2) are shown as dotted blue curves in Fig. 7. We also show constraints obtained by analyzing these four binary observations with Gaussian mass prior (p1b) (as opposed to six observations in Sec. IV C 1). Again, as expected, we see the constraints slightly worsen, at all PN orders, if we do not include information from the GR estimates (dashed red curve). This is because lifting the prior on the masses allows the chains to explore a wider region in the component-masses subspace, and since we are only using 1 ppK parameter here, the masses are not well constrained.

This inflates the marginalized posterior on the ppE deformation deteriorating the bound.

We refrain from plotting a curve corresponding to this case in Fig. 5 since the number of observations analyzed in this case is different from all other cases we considered (similar curves are shown in Fig. 7 instead). Nevertheless, If one were to plot it, the curve will closely trace the dotted magenta line corresponding to the uniform mass prior case.

3. Joint results using a uniform mass prior

If we use flat priors on the masses, without any additional information (p3), our overall estimation of the
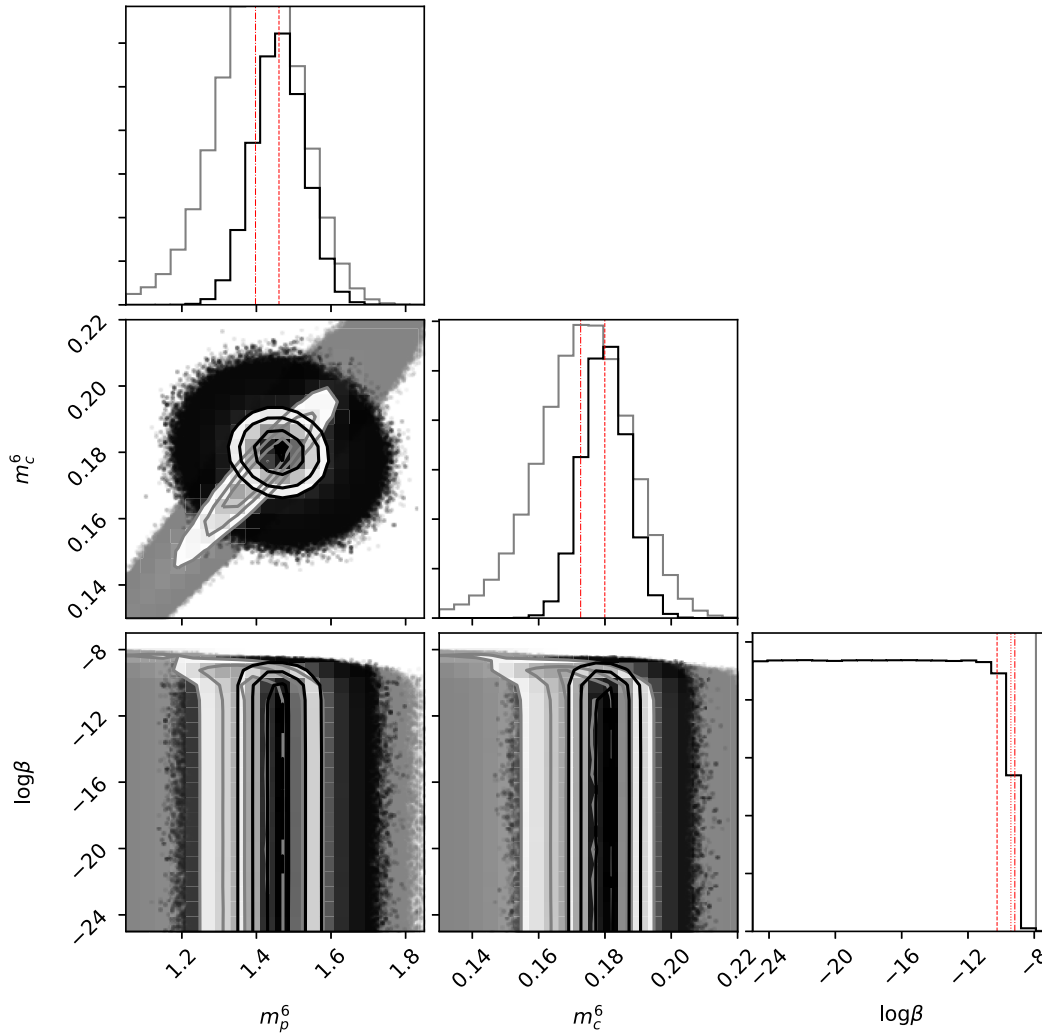

FIG. 8.   Corner plot showing the joint estimation of mass parameters and the phase ppE corrections appearing at -1PN order, for the binary J1738-0333. Black contours correspond to the case where the priors on the masses are set using their GR values (Gaussian mass prior), and the prior on the ppE correction is uniform. The gray contours, shown here for comparison, correspond to the case where the measurement of the mass ratio is used as a prior. We also show the mean values (best fits) for the component masses with a dashed (dash-dotted) line for the Gaussian mass prior (mass-ratio prior) case. For the ppE corrections (corner right plots), the dashed (dash-dotted) line corresponds to the 95 percentile upper limit for the two cases. Observe that the estimates on the component masses worsen when we use the mass-ratio prior, compared to the Gaussian mass prior case, but the effect on the estimate of the ppE correction is mild. We also show the upper limit obtained when using uniform priors on masses (dotted red line), but corresponding histograms for the masses are not shown since we only recover the pulsar priors on the masses in this case.

parameters worsens. We find that we cannot estimate the component masses for any of the six binaries with significant confidence. Again, this is due to the fact that we are using a single ppK measurement for the analysis without any additional information on the masses. Hence, we have to explore a much larger parameter space with very limited information. The 95 percentile upper limits on the amplitude and phase ppE corrections, at various PN orders, are shown as dotted magenta curves in Fig. 5. We find that the constraints on the phase corrections are now around 1–2 orders of magnitude weaker than the case in which the mass priors are chosen based on GR estimates.

To present an example of how our estimates change when we use this prior, Fig. 8 shows the 95 percentile upper limit for the ppE phase correction at −1PN order ($b = -7/3$), estimated by analyzing the binary J1738-0333. We do not show the full corner plot for this case as the marginalized histograms for the component masses will merely show that we manage to recover the priors. In spite of this, we find that, although removing prior information about the component masses has a disastrous effect on the estimation of the masses, it does not prevent us from constraining the ppE parameters. Clearly, there is a deterioration of the ppE constraints by an order of magnitude, but a constraint is still possible, even though the masses cannot be estimated. The reason for this is that a flat prior still has a boundary, so the chains are not allowed to explore *any* value of the component mass subspace. This limitation is apparently sufficient to allow for ppE constraints.

Given that the ppE constraints depend on the priors, as shown for a test case in Table III, which bound should we take seriously? The answer to this question reveals itself when we understand *why* the ppE constraints depend on the priors. As we have explained throughout this paper, this is because when one uses a single ppK parameter, the component masses of the binary pulsar cannot be measured, due to the strong degeneracy between these parameters (most evident in Fig. 8). This induces a widening of the marginalized posterior on the ppE constraints. In that sense, this variation in the ppE estimates for different mass priors is artificial and just a consequence of working with a single ppK measurement. If we instead had worked with two ppK measurements (such as $\dot{P}_b$ and the Einstein delay) or more ppK parameters, then there would be enough information in the data to constrain the component masses

much better, yielding ppE constraints that are similar to what we find when we use Gaussian mass priors. An extension of this paper to include multiple ppK parameters is left to future work, as it first necessitates the extension of the ppE framework to eccentric binaries.

Let us then close by providing an analytic fit for the constraints on the amplitude and phase ppE constraints as a function of the ppE exponents in the case of Gaussian mass priors. It is easiest to fit $\log \alpha$ ($\log \beta$) to a (b) since the relationship is very close to *linear*, as is demonstrated in Fig. 1, and also in YH10. We obtained the following fitting functions from our constraints for the fixed-mass case:

$$\log \alpha = -3.9822a + 8.5753, \tag{18}$$

$$\log \beta = 9.5777b + 8.5618, \tag{19}$$

where we have used the 95 percentile upper limits.

D. Comparison with LVC constraints

The LIGO-Virgo Collaboration (LVC) has also performed GR tests and released constraints on model-independent deviations from GR [23]. In these studies, relative shifts in the PN coefficients of the Fourier phase, $\delta \phi$ were constrained using GW events, where $\delta \phi$ is defined using (i corresponds to non-GR corrections entering at different PN orders)

$$\phi_i = \phi_i(1 + \delta \phi_i) \tag{20}$$

and then treated as an additional free parameter in the parameter estimation scheme. Note that in the LVC paper $i = n$ corresponds to $(n/2)$PN order and we follow this scheme here.

To compare binary pulsar constraints to those obtained by the LVC, we perform a study similar to that of the previous subsections, but now treating $\delta \phi_i$ as our GR correction parameter instead of using the ppE phase correction β. We use flat priors in $\delta \phi_i$ and obtain posterior samples for the corrections to compare the upper limits with LVC constraints. This exercise is straightforward for the PN corrections appearing at −1PN, −0.5PN, 0PN, and 0.5PN order, since $\delta \phi$ can be very easily related to β for these cases [24],

$$\beta_{-1\text{PN}} = \frac{3}{128}\delta\phi_{-2}\eta^{2/5},$$

$$\beta_{-0.5\text{PN}} = \frac{3}{128}\delta\phi_{-1}\eta^{1/5},$$

$$\beta_{0\text{PN}} = \frac{3}{128}\delta\phi_0,$$

$$\beta_{0.5\text{PN}} = \frac{3}{128}\delta\phi_1\eta^{-1/5}. \tag{21}$$

TABLE III. Comparison between the 95 percentile upper limits obtained on the magnitude of the phase ppE correction $\log \beta$ appearing at −1 PN order i.e., for $b = -7/3$ [Eq. (1)] for the different cases studied in this paper.

m_{GR}	$N(m^{\text{GR}}, \sigma_m^{\text{GR}})$	$N(R, \sigma_R)$	Uniform in mass
−10.3	−10.2	−9.7	−9.8

TABLE IV. Comparison of our 95 percentile upper limits on $\delta\phi$ using joint measurements of binary pulsars, with those obtained by LVC [23,46] using the measurements from GW170608.

PN order	LVC	This work
-1PN	5.4×10^{-3}	10^{-7}
0PN	8.9×10^{-2}	2.5×10^{-2}
0.5PN	1.9×10^{-1}	1.6×10^{1}

For other PN corrections, one would have to account for additional physical parameters like spin, which is beyond the scope of this paper.

Binary pulsar constraints on the -1PN, 0PN, and 0.5PN corrections can be directly compared to those obtained by the LVC. Since the LVC did not release the posterior samples for combined measurements, we use the posterior samples obtained from the most constraining GW event, GW170608, for our comparison [46]. These constraints are very close to the joint LVC constraints, as is evident from Fig. 4 in [23].

The comparison between our binary pulsar constraints and those obtained by LVC is shown in Table IV. We see that the constraints obtained from binary pulsar measurements are competitive with those obtained from GW measurements at 0PN order, and much tighter than LVC constraints at -1PN correction, as expected. But LVC constraints become tighter than binary pulsar ones at positive PN order. Hence, our constraints at lower (negative) PN orders can be used as informed priors for future LVC studies.

V. IMPLICATIONS

In this work, we have presented constraints on the ppE framework using Bayesian analysis of six binary pulsar observations. The constraints, which are 95 percentile upper limits, are 1–2 orders of magnitude tighter than those obtained through approximate (non-Bayesian) methods with the double binary pulsar in YH10. We have further found that our constraints on the ppE corrections are robust to our assumptions regarding the component masses, i.e., different mass priors.

Our most important results are shown in Fig. 5. The areas above the black solid curves are excluded by the joint measurement of six binary pulsar observations. This implies that the amplitude and phase correction magnitudes, α and β, have to be smaller than these upper limits at the different PN orders shown in the figure. The general trend in these plots, of increasing upper limits with increasing values of a or b, is as expected from theoretical considerations (see, e.g., YH10). This can be understood by

looking at the mathematical structure of the ppE corrections [see, e.g., Eq. (1)]. The ppE corrections become smaller the larger the value of a or b, since the systems we are studying are low velocity sources. We also showed that for constant P_b and $\dot{P}_b$ (and the same measurement accuracy), binary systems with lower (higher) total mass give tighter upper limits on negative (positive) PN ppE corrections (see, e.g., Fig. 4). Our constraint on the generic GR deviation parameter at -1PN order, $\delta\phi_{-2}$, is around 4 orders of magnitude tighter than the corresponding LVC constraint as of the submission of this paper.

This is the first study of ppE constraints with binary pulsar observations that analyzes assumptions regarding the component masses We analyzed different mass priors (III C) and found that the major contribution in determining these upper limits is the number of binary pulsar observations, how relativistic these binaries are, how well the ppK parameters are determined, and how many ppK measurements are used for each of these observations. We also studied how much these constraints could vary if we assumed uniform priors on the masses. Since we used a single ppK measurement for these binaries, in the absence of any additional mass information, we obtained uninformative posteriors on masses, but the effect on the estimates of the ppE corrections was considerably milder, suggesting the constraints derived here are robust.

Our ppE constraints can be used as priors when performing similar studies using other pulsar observations or future GW observations. An interesting extension to this project, which is ongoing, is to relate other post-Keplerian parameters to ppE corrections so that more data from binary pulsar observations can be utilized for testing GR. This task is complicated because it may require a deeper understanding of how to extend the ppE framework to more generic, eccentric, and possibly spinning binaries. An interesting approach would be to carry out an effective field theory treatment in which a ppE-corrected Einstein-Infeld-Hoffman Lagrangian is derived to obtain ppE-corrected (conservative) equations of motion. The latter would be important to map ppK parameters to ppE parameters.

[1] C. L. Bennett *et al.* (WMAP Collaboration), Astrophys. J. Suppl. Ser. **208,** 20 (2013).

[2] N. Aghanim *et al.* (Planck Collaboration), arXiv: 1807.06209.

[3] E. Di Valentino, A. Melchiorri, and J. Silk, Nat. Astron. **4,** 196 (2020).

[4] C. M. Will, Living Rev. Relativity **17,** 4 (2014).

[5] D. E. Lebach, B. E. Corey, I. I. Shapiro, M. I. Ratner, J. C. Webber, A. E. E. Rogers, J. L. Davis, and T. A. Herring, Phys. Rev. Lett. **75,** 1439 (1995).

[6] C. W. F. Everitt *et al.*, Phys. Rev. Lett. **106,** 221101 (2011).

[7] M. Kramer, IAU Symp. **8,** 19 (2013).

[8] R. A. Hulse and J. H. Taylor, Astrophys. J. **195,** L51 (1975).

[9] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins, Nature (London) **217,** 709 (1968).

[10] T. Damour and N. Deruelle, Ann. Inst. Henri Poincaré Phys. Théor. **43,** 107 (1985), http://www.numdam.org/item/ AIHPA_1985__43_1_107_0/.

[11] T. Damour and N. Deruelle, J. Astron. Francais **25,** 21 (1985).

[12] T. Damour and J. H. Taylor, Phys. Rev. D **45,** 1840 (1992).

[13] M. Kramer *et al.*, Science **314,** 97 (2006).

[14] J. M. Gerard and Y. Wiaux, Phys. Rev. D **66,** 024040 (2002).

[15] N. Yunes and F. Pretorius, Phys. Rev. D **80,** 122003 (2009).

[16] T. Damour and G. Esposito-Farese, Classical Quantum Gravity **9,** 2093 (1992).

[17] S. Alexander and N. Yunes, Phys. Rep. **480,** 1 (2009).

[18] P. Kanti, N. E. Mavromatos, J. Rizos, K. Tamvakis, and E. Winstanley, Phys. Rev. D **54,** 5049 (1996).

[19] N. Yunes and X. Siemens, Living Rev. Relativity **16,** 9 (2013).

[20] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **116,** 221101 (2016); **121,** 129902(E) (2018).

[21] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **118,** 221101 (2017); **121,** 129901(E) (2018).

[22] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **123,** 011102 (2019).

[23] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. D **100,** 104036 (2019).

[24] N. Yunes, K. Yagi, and F. Pretorius, Phys. Rev. D **94,** 084002 (2016).

[25] R. Nair, S. Perkins, H. O. Silva, and N. Yunes, Phys. Rev. Lett. **123,** 191101 (2019).

[26] N. Yunes and S. A. Hughes, Phys. Rev. D **82,** 082002 (2010).

[27] L. Blanchet, Living Rev. Relativity **17,** 2 (2014).

[28] D. L. Kaplan, J. Boyles, B. H. Dunlap, S. P. Tendulkar, A. T. Deller, S. M. Ransom, M. A. McLaughlin, D. R. Lorimer, and A. T. Stairs, Astrophys. J. **789,** 119 (2014).

[29] I. Cognard *et al.*, Astrophys. J. **844,** 128 (2017).

[30] K. Lazaridis *et al.*, Mon. Not. R. Astron. Soc. **400,** 805 (2009).

[31] P. J. Callanan, P. M. Garnavich, and D. Koester, Mon. Not. R. Astron. Soc. **298,** 207 (1998).

[32] J. Antoniadis *et al.*, Science **340,** 1233232 (2013).

[33] D. J. Reardon *et al.*, Mon. Not. R. Astron. Soc. **455,** 1751 (2016).

[34] B. A. Jacoby, M. Bailes, M. H. van Kerkwijk, S. Ord, A. Hotan, S. R. Kulkarni, and S. B. Anderson, Astrophys. J. **599,** L99 (2003).

[35] J. Antoniadis, M. H. van Kerkwijk, D. Koester, P. C. C. Freire, N. Wex, T. M. Tauris, M. Kramer, and C. G. Bassa, Mon. Not. R. Astron. Soc. **423,** 3316 (2012).

[36] P. C. C. Freire, N. Wex, G. Esposito-Farese, J. P. W. Verbiest, M. Bailes, B. A. Jacoby, M. Kramer, I. H. Stairs, J. Antoniadis, and G. H. Janssen, Mon. Not. R. Astron. Soc. **423,** 3328 (2012).

[37] K. Chatziioannou, N. Yunes, and N. Cornish, Phys. Rev. D **86,** 022004 (2012).

[38] P. C. Peters, Phys. Rev. **136,** B1224 (1964).

[39] C. M. Will, Astrophys. J. **214,** 826 (1977).

[40] T. Damour and G. Esposito-Farese, Phys. Rev. D **54,** 1474 (1996).

[41] P. C. Peters and J. Mathews, Phys. Rev. **131,** 435 (1963).

[42] L. Sampson, N. Yunes, and N. Cornish, Phys. Rev. D **88,** 064056 (2013).

[43] D. S. Sivia and J. Skilling, *Data Analysis–A Bayesian Tutorial*, Oxford Science Publications, 2nd ed. (Oxford University Press, Oxford, 2006).

[44] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, Publ. Astron. Soc. Pac. **125,** 306 (2013).

[45] J. Goodman and J. Weare, Commun. Appl. Math. Comput. Sci. **5,** 65 (2010).

[46] LIGO Scientific and Virgo Collaborations, Data release for testing GR with GWTC-1 (2019).