

Constraining gravitational wave amplitude birefringence with GWTC-3

Thomas C. K. Ng^{1,*}, Maximiliano Isi^{2,†}, Kaze W. K. Wong^{2,‡} and Will M. Farr^{2,3,§}

¹*Department of Physics, The Chinese University of Hong Kong, Shatin, Hong Kong*

²*Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, New York 10010, USA*

³*Department of Physics and Astronomy, Stony Brook University, Stony Brook, New York 11794, USA*



(Received 15 May 2023; accepted 4 October 2023; published 30 October 2023)

The propagation of gravitational waves can reveal fundamental features of the structure of spacetime. For instance, differences in the propagation of gravitational-wave polarizations would be a smoking gun for parity violations in the gravitational sector, as expected from birefringent theories like Chern-Simons gravity. Here we look for evidence of amplitude birefringence in the third catalog of detections by the Laser Interferometer Gravitational Wave Observatory and Virgo through the use of birefringent templates inspired by dynamical Chern-Simons gravity. From 71 binary-black-hole signals, we obtain the most precise constraints on gravitational-wave amplitude birefringence yet, measuring a birefringent attenuation of $\kappa = -0.019^{+0.038}_{-0.029}$ Gpc⁻¹ at 100 Hz with 90% credibility, equivalent to a parity-violation energy scale of $M_{\text{PV}} \gtrsim 6.8 \times 10^{-21}$ GeV.

DOI: [10.1103/PhysRevD.108.084068](https://doi.org/10.1103/PhysRevD.108.084068)

I. INTRODUCTION

Gravitational wave (GW) detections by the LIGO-Virgo-KAGRA Collaboration (LVK) [1–3] are now routinely used to test various aspects of Einstein’s theory of general relativity (GR) [4–6]. Among those, measurements of the basic properties of GWs, like their speed and polarization, can directly probe the fundamental symmetries of the underlying theory of gravity [7]. For instance, unequal propagation of GW polarization eigenstates would reveal that spacetime is birefringent, a smoking gun for parity-odd theories like Chern-Simons gravity [8–11]. Here we constrain the magnitude of possible amplitude birefringence using binary black hole (BBH) signals from the latest LVK catalog, GWTC-3 [12].

Previous studies have constrained amplitude birefringence by performing different statistical analyses. Yamada, Tanaka [13] and Wang *et al.* [14] both performed parameter estimation (PE) on the events in the first GW transient catalog [15], GWTC-1, using birefringent templates. Okounkova *et al.* [16] considered the distribution of observed inclinations of the GW events in the second GW transient catalog [17], GWTC-2, to look for signs of birefringence; Vitale *et al.* [18] carried out a related analysis. While this manuscript was finalized, Zhu *et al.* [19] reported an analysis of GWTC-3.

In this study, we use a frequency-dependent birefringence model to constrain the strength of GW amplitude

birefringence by performing PE on LVK binaries. This model is a better approximation of the birefringence effect expected from theory rather than the frequency-independent model used in Okounkova *et al.* [16]. Including the frequency dependence allows us to break the degeneracy between birefringence and source inclination, which we discuss below in Sec. II B. Compared to Refs. [13,14,16], we perform PE on more events, including events new to GWTC-3 [12], and use a phenomenology-oriented parametrization. We consider 71 binary black hole merger events with a false-alarm rate (FAR) $\leq 1/\text{yr}$, as listed in Table I of Abbott *et al.* [20]. We discuss single-event results in detail and identify degeneracies between birefringence and spin effects, in addition to the already known correlations with source orientation and distance. We use the results from individual events to place a collective population constraint on the strength of GW amplitude birefringence from GWTC-3.

In Sec. II, we briefly review the background of GW amplitude birefringence. In Sec. III, we describe the modification we apply to the baseline GR waveform, summarize our PE configuration, and outline our statistical methods for combining information from multiple events. In Sec. IV, we present our constraint on GW amplitude birefringence, discussing individual events and the catalog collectively. In Sec. V, we discuss the implications of our result, comparing to previous constraints in the literature and outlining correlation structures that appear in our measurements. Finally, we conclude in Sec. V with a summary. In the Appendix, we provide extended results and discussion for two notable events (GW170818 and GW190521).

*thomas.ng@link.cuhk.edu.hk

†misi@flatironinstitute.org

‡kwong@flatironinstitute.org

§wfarr@flatironinstitute.org

II. BACKGROUND

A. Birefringence

In GR, GWs are comprised of two independent polarization modes, usually represented in the linear basis of plus (+) and cross (×) states. In the Fourier domain, these can be combined into left-handed (L) and right-handed (R) circular states (see, e.g., [21]),

$$h_{L/R} = \frac{1}{\sqrt{2}}(h_+ \pm ih_\times), \quad (1)$$

where h is the frequency domain strain, with the plus and minus signs for L and R , respectively. These circular modes represent eigenstates of the helicity operator (helicity ± 2) and possess a definite parity. Einstein's theory, which is parity even, predicts no difference in the dynamics of these two states.

Yet, parity odd extensions of GR may make distinctions between the two circular polarizations, potentially appearing in both the generation and propagation of GWs. The latter can manifest in changes to the relative amplitude and phase of the polarizations that accrue as the wave propagates, giving us hope of detecting initially small effects that compound over long propagation distances.

In particular, *amplitude* birefringence would enhance one polarization mode over the other. Following [10], the effect of birefringence can be modeled as a frequency-dependent amplification or dampening; similar derivations can be found in [19,22–24]. To first order in the equations of motion, in theories like Chern-Simons gravity, the Fourier-domain waveform observed at a comoving distance d_C away from the source can be written as

$$h_{L/R}^{\text{br}}(f) = h_{L/R}^{\text{GR}}(f) \times \exp\left(\pm \kappa d_C \frac{f}{100 \text{ Hz}}\right), \quad (2)$$

where the emitted waveform $h_{L/R}^{\text{GR}}$ is modified by an exponential birefringent factor to yield the observed waveform $h_{L/R}^{\text{br}}$. The overall magnitude of this effect for a given frequency f is set by an attenuation coefficient κ , which encodes the intrinsic strength of the birefringence: κ^{-1} represents an “attenuation length” encoding the typical distance that yields an e -folding in the amplification or dampening of the polarizations at a fiducial signal frequency of 100 Hz at the detector. The emitted waveform for a given source (i.e., the waveform observed in the near zone, very close to the source) will generally differ from the analogous waveform predicted by GR [10,25]; however, since we expect most viable modifications to GR to be intrinsically small (e.g., [26]), it is standard to approximate the emitted waveform by the prediction from GR (hence the notation “ h^{GR} ” above).

Although the intrinsic modification is small, the effect targeted by Eq. (2) accumulates as the GW propagates.

During propagation, the effect of birefringence will be built up with the number of cycles, which is itself a function of the distance traveled and the frequency of the GWs. According to Eq. (2), a positive κ means the left-handed polarization is enhanced over the right-handed polarization, while a negative κ means the opposite; when $\kappa = 0$, the observed waveform is the same as GR predicts, meaning there is no birefringence.

Equation (2) can be derived as the first-order effect in an expansion away from GR under multiple frameworks. In general, κ will be a function of the theory parameters and the cosmological history, e.g., the value of the pseudoscalar field and its derivative in Chern-Simons gravity [10]. Since it originates from a truncated series expansion,¹ Eq. (2) is a good approximation only for small exponents, namely,

$$\left(\frac{|\kappa|}{\text{Gpc}^{-1}}\right) \left(\frac{d_C}{\text{Gpc}}\right) \left(\frac{f}{100 \text{ Hz}}\right) < 1, \quad (3)$$

recalling that κ has dimensions of inverse length. Otherwise, more frequency-dependent terms could enter the exponent of Eq. (2) in a theory-dependent way.

B. Inclination and other degeneracies

Under certain conditions, the effect of birefringence can be degenerate with a change in the orientation of the source with respect to the line of sight [10]. Concretely, for a nonprecessing compact binary inspiral in GR, the observed amplitude ratio of the left-handed and right-handed polarizations is only a function of the inclination ι , the angle between the orbital angular momentum of the source, and the line of sight [27]. For the dominant $\ell = |m| = 2$ angular mode of the radiation, the relation between the amplitude ratio and the inclination is

$$\frac{h_L^{\text{GR}}}{h_R^{\text{GR}}} = \left(\frac{1 - \cos \iota}{1 + \cos \iota}\right)^2 \quad (4)$$

for all frequencies (see, e.g., Sec. III C in [21]).

Since birefringence impacts the observed amplitude ratio of left- and right-handed modes, it could also affect inferences about the source inclination [10]. However, the two effects are degenerate only if the frequency dependence of Eq. (2) is neglected. This is easy to see from Eq. (2) since the implied polarization ratio for the $\ell = |m| = 2$ mode of a nonprecessing source is

¹Often the birefringent effect is written as an expansion in redshift z [e.g., [22]], rather than distance. Here we use d_C to emphasize that the effect accumulates *per cycle*, or with propagation distance. Thus, current ground-based experiments, which can detect gravitational waves to $z \simeq 1$, $d_C \simeq d_H$ (Hubble distance), are already observing propagation over a large fraction of the Universe and constrain this effect meaningfully.

$$\frac{h_L^{\text{br}}}{h_R^{\text{br}}} = \left(\frac{1 - \cos \iota}{1 + \cos \iota} \right)^2 \exp \left(2\kappa d_c \frac{f}{100 \text{ Hz}} \right). \quad (5)$$

For an isolated Fourier mode of definite frequency f , the effect of birefringence will be degenerate with a change in inclination; however, if multiple modes come into play, then no change in inclination alone can mask the effect of birefringence, which will affect the time domain waveform nontrivially (Fig. 1).

Okounkova *et al.* [16] took the effect of birefringence to be independent of the frequency, which is a simplified approximation of the birefringence model in Chern-Simons gravity. This assumption results in a full degeneracy between κ and ι : to reconstruct the amplitude ratio from the interferometer data, a value of ι representing a more face-off inspiral can pair with a positive value of κ , or a value of ι representing a more face-on inspiral with a negative value of κ . That fact can be used to constrain frequency-independent birefringence by searching for features in the distribution of inferred inclinations [16].

By implementing Eq. (2), we generally break the degeneracy between birefringence and source orientation; this was also the case in the frequency-dependent relations studied in [13,14]. Nevertheless, there exist systems for




FIG. 1. Illustration of amplitude birefringence. The GR waveform for the $\ell = |m| = 2$ mode of a nonprecessing BBH seen edge-on ($\cos \iota = 0$) is linearly polarized and thus contains equal amounts of left- and right-handed modes for all frequencies (dotted, top). However, if spacetime were birefringent (BR) following Eq. (2), then the waveform observed on Earth would contain different fractions of the two circular modes, with higher frequencies affected more strongly (solid, top). In the time domain, this manifests as a time-dependent amplification of the waveform, with a stronger effect at later times when the chirp reaches a higher instantaneous frequency (bottom). For this example, the black holes do not spin and have equal masses $m_1 = m_2 = 10M_\odot$, and we have chosen a luminosity distance $d_L = 400 \text{ Mpc}$ and $\kappa = 0.6$.

which the degeneracy cannot be broken, in practice, because not enough frequencies are available in the data. This may also be the case for quasimonochromatic sources, like nonaxisymmetric pulsars or very light binaries, which are well approximated by a single Fourier mode.

As we will find in Sec. IV, the effect of birefringence can be (partially) degenerate with other parameters besides source inclination. Indeed, the frequency-dependent amplification or dampening of polarizations caused by κ can sometimes be absorbed by changes in intrinsic parameters, like the spins, with concurrent adjustments to the source inclination and distance. As a measure of the black hole (BH) spin magnitudes along the orbital angular momentum, we will use the effective spin parameter [28–30]

$$\chi_{\text{eff}} \equiv \frac{1}{1+q} (\chi_{1\parallel} + q\chi_{2\parallel}), \quad (6)$$

where $\chi_{i\parallel}$ are the norms of the projections of the dimensionless spin vectors along the orbital angular momentum, and the mass ratio is $q \equiv m_2/m_1 \leq 1$. We will study the interplay between birefringence and precession by focusing on the posterior of the precessing spin parameter χ_p , defined as [31]

$$\chi_p \equiv \max \{ \chi_{1\perp}, k\chi_{2\perp} \}, \quad (7)$$

where $\chi_{i\perp}$ are the norms of the projections of the dimensionless spin vectors onto the orbital plane at a reference time, and $k \equiv q(4q+3)/(4+3q)$. We chart this and other approximate degeneracies as part of the results presented in Sec. IV.

III. METHOD

A. Single-event parameter estimation

To constrain birefringence, we reanalyze events from GWTC-3 [12,32] implementing Eq. (2) to directly obtain a posterior on κ from the strain of each event. We analyze the 71 BBHs that were detected with $\text{FAR} < 1/\text{yr}$, less stringent than the typical LVK threshold of $1/1000 \text{ yr}$ [6,33]. The FAR values are typically determined by GR pipelines, which could down-rank signals beyond GR [33–35]; however, since detectability through matched-filtering is most sensitive to the phase, not the amplitude, we expect only a minor decrease in sensitivity to the kind of birefringent signals explored here. To avoid extended computations on longer signals and considering these are generally at closer distances, we do not analyze systems involving neutron stars in this work. We procure strain data from Gravitational Wave Open Science Center [36].

We estimate source parameters using a custom version of the BILBY software [37], modified from the baseline version to apply Eq. (2) for any GR baseline waveform. We take the PE configuration in [12,32,38,39] as a starting point, with

IMRPENOMXPHM [40] as the reference waveform. We apply a distance prior corresponding to a uniform distribution over comoving volume and source-frame time (see, e.g., Eq. 10 in [41]), and set the prior on κ to be uniform between -1 Gpc^{-1} and 1 Gpc^{-1} .

For GW190521, we increase the maximum distance allowed by the prior to $1.5\times$ the original value in [38], as the birefringence effect results in posterior support at larger distances. For GW190720, we decrease the analysis segment from 16 s to 8 s, in order to accommodate missing data near the edges of the 16 s segment in Virgo. Otherwise, the configuration is as in [12,32,38,39].

In order to validate our PE implementation, we reproduce the LVK PE results obtained assuming GR by enforcing $\kappa = 0$; this also has the advantage of producing GR runs that are directly comparable to our birefringent runs. All our PE results, including the GR validation runs, are published in Ng *et al.* [42].

B. Collective analysis

1. Shared birefringence parameter

In the most simplified scenarios, birefringence is a property of spacetime that is not intrinsic to any source or region in space. Consequently, we should expect κ to take the same value for all signals, whether it vanishes or not. Under this assumption, the posterior on κ inferred collectively from all events is simply obtained from the product of the individual likelihoods, $p(d_i|\kappa) \propto p(\kappa|d_i)/p(\kappa)$, such that

$$p(\kappa|\{d_i\}) \propto p(\kappa) \prod_i \frac{p(\kappa|d_i)}{p(\kappa)}, \quad (8)$$

where d_i is the strain data for the i th event, and $p(\kappa)$ is the prior on κ ; since the prior is uniform, in our case Eq. (8) reduces to the product of the posteriors, namely $p(\kappa|\{d_i\}) \propto \prod_i p(\kappa|d_i)$. We use Eq. (8) to obtain the primary constraint presented in this work.

2. Nonshared birefringence parameters

Under many frameworks, birefringence is brokered by an extra parity-odd field that couples to gravity. In that case, the effective strength of birefringence may vary along the different lines of sight to each event, depending on cosmic history and the local evolution of the field. The *simplest* assumption is that the field manifests equally for all events, as assumed in the previous subsection, but this is not strictly required. These considerations motivate a collective analysis that does not assume κ is shared across events [43,44]. Such an analysis has the additional advantage of helping us further characterize our set of measurements and identify potential outliers.

To do this, we apply hierarchical Bayesian inference [45,46] to model the distribution of κ 's consistent with the observed data: we posit that, rather than a unique global value of κ , there is a specific value of the parameter κ_i , associated with each event, and that this is drawn from some unknown distribution of true underlying values; from the imperfect measurements of κ_i for each event, we may reconstruct the underlying distribution. If we are interested in constraining the first two moments of the distribution, then it is convenient to parametrize the κ_i 's as drawn from a Gaussian with unknown mean μ and variance σ^2 , i.e., $\kappa_i \sim \mathcal{N}(\mu, \sigma^2)$ [47], and measure those hyperparameters from the collection of observed data.

If GR is correct and there is no birefringence, then we should find the observed κ distribution to be consistent with a delta function at the origin ($\kappa_i = 0$ for all i , or $\mu = \sigma = 0$); on the other hand, if spacetime is globally birefringent, then we expect to find a delta function at some nonzero value ($\kappa_i = \kappa \neq 0$, or $\mu = \kappa$ and $\sigma = 0$). But this analysis also has the power to reveal unexpected physics or systematics in our measurements: if σ is confidently found to be nonzero, then this would imply that our set of measurements is statistically unlikely to originate from a unique κ value. This could signal richer physics than is implied by Eq. (2) or, more prosaically, that there are outliers in our measurements due to mismodeling, e.g., in the waveform approximant or the noise of the detector.

Starting from the posterior on κ from each i th event, $p(\kappa_i|d_i)$, the posterior on the hyperparameters μ and σ can be calculated by

$$p(\mu, \sigma|\{d\}) \propto p(\mu, \sigma) \prod_i \int \frac{p(\kappa_i|d_i)}{p(\kappa_i)} p(\kappa_i|\mu, \sigma) d\kappa_i, \quad (9)$$

where $p(\kappa)$ is the prior initially applied to κ during PE, which in our case is a uniform distribution, $\mathcal{U}[-1, 1]$. Further choosing the hyperpriors on μ and σ to be uniform, Eq. (9) simplifies to

$$p(\mu, \sigma|\{d\}) \propto \prod_i \int p(\kappa_i|d_i) p(\kappa_i|\mu, \sigma) d\kappa_i, \quad (10)$$

where $p(\kappa_i|\mu, \sigma) \propto \exp(-|\kappa_i - \mu|^2/2\sigma^2)$ is the usual Gaussian likelihood.

Beyond measuring the population mean and variance, we also calculate the expected population distribution of κ marginalized over μ and σ . This is defined by

$$p(\kappa|\{d\}) = \int p(\kappa_i|\mu, \sigma) p(\mu, \sigma|\{d\}) d\mu d\sigma \quad (11)$$

and represents our overall expectation for the true values of κ , given the observed set of individual measurements and, crucially, assuming the underlying distribution is Gaussian. Although the measurement of the population mean and

variance applies irrespective of whether the underlying κ_i distribution is truly Gaussian, the specific shape of the population-marginalized distribution of Eq. (11) is not; therefore, Eq. (11) should not be interpreted as giving a generic inference on the κ distribution—more expressive models than a Gaussian would be better suited to that purpose. To sample the posterior distribution of μ and σ , we use the sampling package FLOWMC [48].

IV. RESULTS

In this section, we present the results of our study. We first show the κ measurements from all events in our set, as well as the resulting global measurement of κ that represents our primary constraint on birefringence (Sec. IV A). We then assess the collection of measurements in more detail through a hierarchical analysis (Sec. IV B). Finally, we discuss some special events, individually, and outline the degeneracies that arise between birefringence and orbital precession (Sec. IV C).

A. GWTC-3 result

Figure 2 shows the main result of our GWTC-3 analysis as represented by the posterior distribution on κ (abscissa) obtained individually for each event (ordinate). Posteriors are colored by the respective posterior mean distance to the origin in units of standard deviation, i.e., $|\mu_i/\sigma_i|$ for each event i .² The collective measurement obtained by assuming a shared κ across events, Eq. (8), is represented by its 90%-credible symmetric interval (blue band) around the median (blue line); this joint measurement is fully consistent with $\kappa = 0$ (dashed line) with the credible level of 0.209, meaning we find no evidence for birefringence.

Figure 2 makes it clear that not all BBHs in GWTC-3 are equally informative about birefringence. When considered individually, the events that best constrain κ are listed in Table I, in order of increasing standard deviation σ_i . That table also shows the credible level (CL) at which the posterior supports $\kappa = 0$, whereby CL = 0 (CL = 1) means the posterior supports that value with high (low) probability. This quantity represents the relative height of the probability density function at $\kappa = 0$ and is not tied to a symmetric interval, making it particularly useful in assessing bimodal posteriors.

Judging by μ_i/σ_i , the two events that show the largest tension with $\kappa = 0$ are GW170818, for which $\mu_i/\sigma_i = 2.0$, and GW200129_065458 (henceforth GW200129), for which $\mu_i/\sigma_i = -1.5$. However, as we discuss in Sec. IV C 3, we

²We use the i notation for individual events here to disambiguate from the population mean μ and standard deviation σ of Eq. (9).

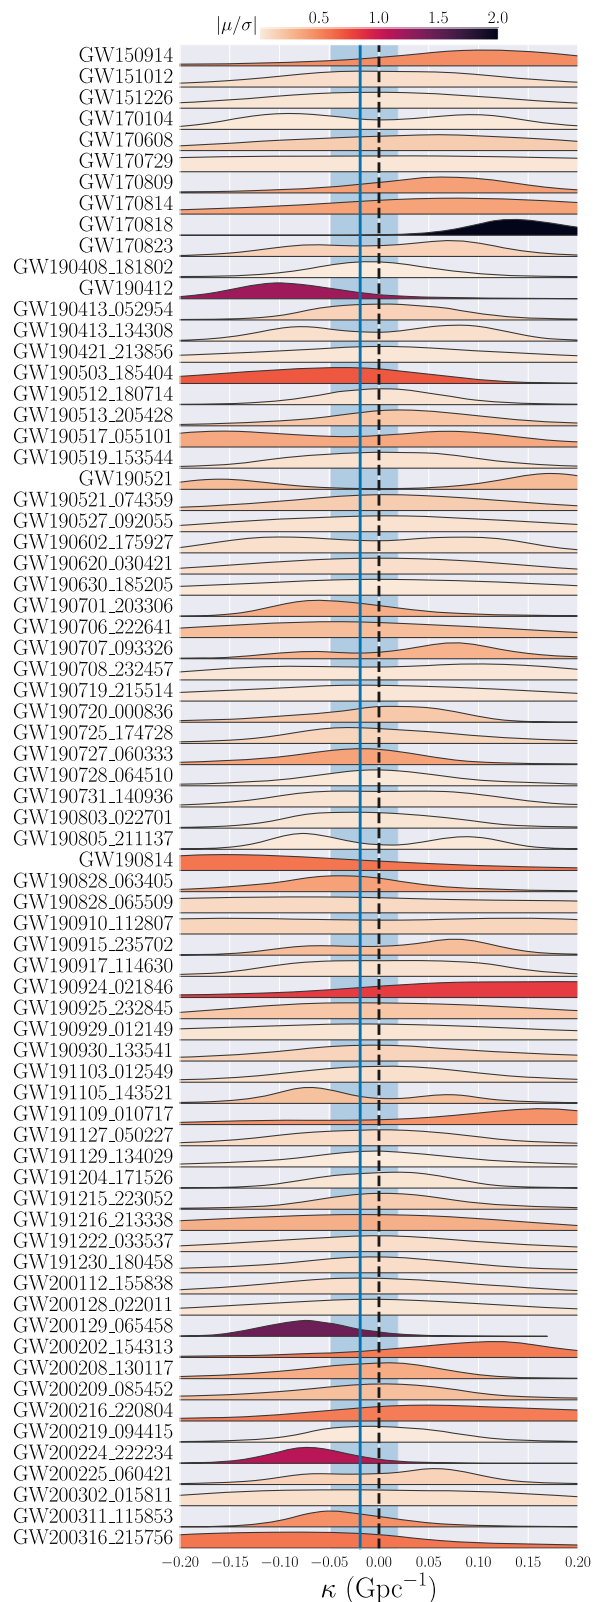


FIG. 2. Individual-event κ posteriors (distributions) and joint measurement (blue band, 90% CI; blue line, median).

TABLE I. Events that best constrain κ , sorted by posterior standard deviation σ_i . CL is the credible level of $\kappa = 0$.

Event	κ (Gpc ⁻¹)	σ_i (Gpc ⁻¹)	CL
GW200129_065458	$-0.072^{+0.085}_{-0.071}$	0.048	0.876
GW200224_222234	$-0.069^{+0.086}_{-0.071}$	0.058	0.860
GW200311_115853	$-0.036^{+0.115}_{-0.077}$	0.061	0.584
GW190412	$-0.091^{+0.129}_{-0.086}$	0.068	0.857
GW191204_171526	$+0.004^{+0.100}_{-0.109}$	0.069	0.205
GW190512_180714	$+0.002^{+0.113}_{-0.109}$	0.070	0.005
GW170818	$+0.146^{+0.136}_{-0.088}$	0.076	0.977
GW200219_094415	$-0.001^{+0.125}_{-0.118}$	0.079	0.109
GW190701_203306	$-0.041^{+0.158}_{-0.086}$	0.084	0.587
GW190413_052954	$+0.008^{+0.137}_{-0.113}$	0.087	0.081

have reason to think that the preference for $\kappa < 0$ in GW200129 might be driven by noise anomalies in the Virgo detector; with that in mind, in the next section we consider the effect of excluding this event from the joint result (we find its impact to be minimal).

Finally, a set of events stands out in Fig. 2 due to evident bimodality in the κ posterior. To a varying degree, that is the case for those events listed in Table II, which tend to have quite high total masses in the detector frame (Table II shows total mass as measured in the standard GR analysis). For these bimodal posteriors, μ_i/σ_i is not a good proxy for agreement with GR; instead, we can rely on $\text{CL}(\kappa = 0)$. By this measure, the bimodal events are some of the least consistent with $\kappa = 0$, GW190521, in particular.

As we anticipated in Sec. II B, we understand the bimodality in κ to be linked to spin effects and often to precessing morphologies, in particular. Other parameter degeneracies also come into play, especially for the lighter events GW191105_143521 (henceforth GW191105) and GW170104. We discuss this further in a dedicated section below (Sec. IV C).

B. Hierarchical modeling

Figure 2 shows a certain degree of variance in the distribution of κ posteriors for different events, including

TABLE II. Events with bimodality in the κ posterior, the GR measurement of their detector-frame total mass (M), precessing spin χ_p , and effective spin χ_{eff} , as well as the CL of $\kappa = 0$ from the birefringence analysis.

Event	M (M_\odot)	χ_p	χ_{eff}	CL
GW170104	$60.6^{+4.0}_{-3.9}$	$0.42^{+0.41}_{-0.32}$	$-0.04^{+0.15}_{-0.20}$	0.786
GW190413_134308	$133.3^{+20.6}_{-20.5}$	$0.56^{+0.35}_{-0.41}$	$-0.04^{+0.28}_{-0.36}$	0.837
GW190521	$250.1^{+40.3}_{-37.5}$	$0.46^{+0.36}_{-0.32}$	$+0.04^{+0.29}_{-0.39}$	0.996
GW190805_211137	$146.6^{+19.8}_{-20.1}$	$0.53^{+0.33}_{-0.33}$	$+0.33^{+0.29}_{-0.36}$	0.895
GW191105_143521	$22.4^{+2.6}_{-0.6}$	$0.31^{+0.44}_{-0.24}$	$-0.03^{+0.11}_{-0.08}$	0.802

some apparent outliers like GW170818 or GW190521. This is not unexpected: assuming independent Gaussian noise instantiations for each event, we might expect up to ~ 3 out of the 71 posteriors (i.e., $\sim 5\%$) to deviate away from the true κ value by over $\sim 2\sigma$ due to noise alone.

To further assess the statistical properties of the set of posteriors in Fig. 2, we apply the hierarchical analysis described in Sec. III B 2. By characterizing the population mean and standard deviation over events, this also allows us to obtain a collective measurement that does not assume all events share the same value of κ .

We summarize the result of this exercise in Fig. 3, which shows the posterior on the population mean μ and standard deviation σ inferred from the collection of observations in Fig. 2. The figure shows two distributions, which result from analyses with (blue) and without (orange) the potentially-contaminated event GW200129; the difference between the two is mainly restricted to a slight shift in μ , indicating that GW200129 has a small effect on our overall population inference.

Both versions of the posterior support the lack of birefringence ($\mu = \sigma = 0$) within 90% credibility; from the marginals of the result including all events, we constrain $\mu = -0.008^{+0.031}_{-0.031}$ Gpc⁻¹ for 90%-credible symmetric intervals around the median, and $\sigma < 0.048$ Gpc⁻¹ for the 90%-credible one-sided upper limit. However, even though $\sigma = 0$ is well supported, the σ posterior peaks visibly away from the origin, indicating some preference for a nonzero variance. This could be a sign of the presence of outliers in our sample.

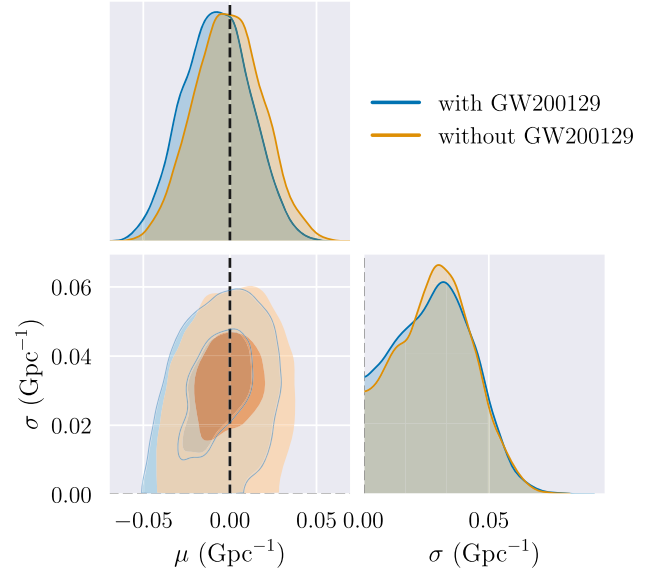


FIG. 3. The posterior of the κ population hyperparameters μ and σ , including (blue) and excluding (orange) GW200129 from the collection of events. The two-dimensional (2D) contours correspond to the 39.35% and 90% credible levels. The plot shows that the population constraint on κ is consistent with no birefringence ($\mu = \sigma = 0$) at the 90% credible level.




FIG. 4. Individual-event κ_i distributions of Fig. 2 reweighted in light of the population-level inference of Fig. 3, marginalizing over μ and σ . Curves are colored by the magnitude of the supported deviation, as represented by the respective posterior $|\mu_i/\sigma_i|$.

As a visual check for outliers, we reconsider the set of measurements in Fig. 2 in light of the hierarchical result for μ and σ in Fig. 3, including GW200129 (blue curve). This amounts to reweighting the κ posterior for each event under a population prior marginalized over μ and σ , conditional on the measurements from all other events [49,50]. The result in Fig. 4 does not show evidence for any of the events being in obvious tension with the population, even though the GW170818 curve stands out from the rest due to its higher support for $\kappa > 0$. This feature appears to offset a few other events which tend to favor $\kappa < 0$. The interaction between these distributions leads to a hyperposterior that is fully consistent with $\mu = 0$ while offering some support for $\sigma > 0$ (Fig. 3). Future observations will determine whether there is truly evidence for a nonvanishing variance in this population.

Finally, the hierarchical result in Fig. 3 can be translated into an overall expectation for κ_i under the assumption of a Gaussian distribution via Eq. (11). We show the result of doing this in Fig. 5, where we also compare to the posterior on κ obtained by assuming a shared value across events (same result shown as a blue band in Fig. 2). The hierarchical measurement leads to an expectation that $\kappa_i = -0.010^{+0.064}_{-0.052} \text{ Gpc}^{-1}$, whereas the shared- κ measurement implies $\kappa = -0.019^{+0.038}_{-0.029} \text{ Gpc}^{-1}$, both symmetric 90%-credible intervals around the median.

C. Notable events

Having established that the collection of detections is globally consistent with $\kappa = 0$, here we focus on three events whose κ posteriors stand out in Fig. 2: GW170818, GW190521, and GW200129. When considered in isolation, the first of these is the unimodal event with the most




FIG. 5. Restricted and generic posteriors on κ , which, respectively, do and do not assume that κ is shared by all events (color). Solid (dashed) traces indicate the analysis included (excluded) GW200129. The black dashed line marks the absence of birefringence ($\kappa = 0$). The shaded distribution is the primary result in this work (blue band in Fig. 2).

significant support for nonzero κ , the second is the most extreme representative of a class of events with bimodal κ posteriors, and the third shows signs of potential noise anomalies. Through these examples, we elucidate some of the interactions between κ and the source luminosity distance, inclination, and spin parameters.

1. GW170818

GW170818 produced the posterior most displaced from $\kappa = 0$, when judged by μ_i/σ_i in Fig. 2. Figure 6 shows that this happens because birefringence opens up a region of parameter space with $\kappa > 0$ for larger distances and smaller inclination angles than would be allowed in the GR case. We can make sense of this by noting that an edge-on nonprecessing source produces linearly-polarized waves, meaning that a smaller inclination leads the two circular polarizations to have similar amplitudes. On the other hand, having $\kappa > 0$ enhances the left-handed modes during propagation, per Eq. (2). The two effects can be balanced to match the polarization ratio observed at the detector (predominantly left handed, per the preference for $\cos \iota \approx -1$ in the GR analysis), as long as the distance is also enhanced to yield the right amount of birefringence and overall signal power. This is similar to the degeneracy mentioned in Sec. II B.

It is difficult to unequivocally identify a specific feature of the GW170818 data that leads to this posterior structure. However, it appears to be related to this event's support for precession, in conjunction with its uncommonly definite measurement of the polarization, phase, and spin angles [51] (see Appendix A 1). The relevance of precession is evident from the posterior on χ_p (Fig. 6, top right): allowing for $\kappa \neq 0$ leads to reduced support for precession. We can understand this in reference to Fig. 1: if only a short portion




FIG. 6. GW170818 posterior on κ , luminosity distance d_L , and inclination $\cos \iota$ from our birefringence analysis (blue), compared to the GR result (orange). The top right panel shows the marginalized posterior on χ_p : allowing for birefringence reduces the preference for precession. (See Fig. 11 for a full corner plot).

of the signal is observed, then the frequency-dependent signal enhancement or dampening due to birefringence can mimic the time-dependent amplitude modulation produced by a precession cycle. Therefore, under these circumstances, similar morphologies can be obtained by setting $\chi_p > 0$ or $\kappa > 0$, as long as the distance and inclination can also be adjusted accordingly.

The fact that the birefringent analysis favors a high κ rather than a high χ_p can be explained by a consequence of prior volume: many more configurations are available with long distances and high κ than with short distances and small $|\kappa|$. The preference for $\kappa > 0$ over $\kappa < 0$ (and, therefore, the lack of bimodality in the $\cos \iota$ and κ posteriors) is likely related to both the definite measurement of left-handed polarizations and the specific phasing of the precession cycle. The latter manifests as a precise constraint on the spin orientation and phase angles in the GR analysis [51] (see also Appendix A 1). The observed amplitude modulation (say, increasing vs decreasing towards the merger) likely determines the allowed sign of κ for this event. We provide the full corner plot for all relevant parameters in Fig. 11 in Appendix A 1.

2. GW190521

GW190521 is the most extreme representative of a class of events with bimodal κ posteriors (Table II). Figure 7 shows that the two peaks of the κ distribution arise from respective modes in the inclination (lower left corner): one

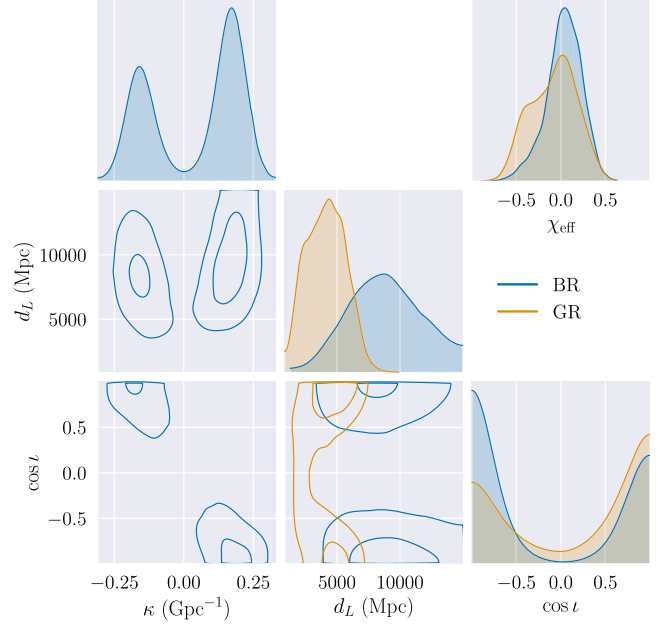


FIG. 7. GW190521 posterior on κ , luminosity distance d_L , and inclination $\cos \iota$ from our birefringence analysis (blue), compared to the GR result (orange). The marginalized posterior on χ_p is shown in the top right panel. (See Fig. 12 for a full corner plot.) Note that, although we show d_L in this plot, our definition of κ is associated with d_C , meaning that the condition in Eq. (3) is satisfied even for the largest values of d_L supported in this posterior.

solution corresponds to an intrinsically face-on source ($\cos \iota \approx 1$) with negative κ , and the other to a face-off source ($\cos \iota \approx -1$) with positive κ . In the former, an initially right-handed signal is birefringently enhanced by $\kappa < 0$ as it propagates towards the detector; in the latter, an initially left-handed signal is enhanced by $\kappa > 0$. The two outcomes (a right-handed or left-handed signal at the detector) match the two possibilities allowed also in the reference GR run, which itself yields a bimodal $\cos \iota$ posterior. In both cases, the overall amplification of the signal introduced by a nonzero κ is balanced by a large luminosity distance ($d_L \gtrsim 5$ Gpc), so as to match the observed signal-to-noise ratio.

Introducing birefringence affects our inference of the spins in the GW190521 system. We illustrate this in Fig. 7 through the posterior on the effective spin χ_{eff} , which is favored to take values closer to zero in the birefringent analysis; this can also be seen, perhaps more clearly, in the joint posterior for the individual spin magnitudes, provided in Appendix A 2. This suggests that at least part of the data features that lead the GR analysis to infer near extremal spins can be alternatively explained as a consequence of birefringence. The fact that this is possible is not fully surprising given that the signal from this heavy system is only observed for a very small number of cycles [52,53] (cf. Fig. 1).

All in all, the situation for GW190521 is not unlike that for GW170818 discussed above (Sec. IV C 1): structures in the data that would normally be interpreted as traces of high component spins are instead (at least partially) attributed to birefringent propagation, driven by a prior preference for high distances. The degeneracies in the case of GW190521 are likely exacerbated by the heavier mass of the system, which reduces the observed number of cycles. Crucially, GW190521 also lacks a definite measurement of the polarization state at the detector: without certainty about the ratio of right- versus left-handed components in the observed signal, it is not possible to disambiguate between $\kappa > 0$ versus $\kappa < 0$; hence the κ distribution remains bimodal for GW190521, unlike for GW170818. The poor determination of the polarization state of the detected signal is evidenced by the bimodal $\cos \iota$ posterior inferred assuming GR in Fig. 7 and is attributable to several factors including sky location.

3. GW200129

GW200129 is the event with the second largest $|\mu_i/\sigma_i|$ in Fig. 2. However, data for this event were affected by a non-Gaussian noise disturbance (glitch) in the Virgo instrument, which was subtracted from the publicly-available data used for parameter estimation [54]. Since previous work suggests the degree of glitch subtraction affects the inference for this event [55], we consider whether the apparent preference for $\kappa < 0$ could also be tied to the instrumental artifact.

To this end, we perform three additional PE runs for GW200129, considering only two detectors at a time: LIGO Hanford and Virgo (HV), LIGO Livingston and Virgo (LV), and LIGO Hanford and LIGO Livingston (HL). If the preference for $\kappa < 0$ is tied to the glitch in Virgo, we expect it to disappear in the HL run, which excludes Virgo data.

This is indeed the case, as we show in Fig. 8: all runs including Virgo lean towards $\kappa < 0$ (solid curves in color), whereas the LIGO-only run is fully consistent with $\kappa = 0$ (dashed black). While this is not conclusive proof that the glitch itself is driving the result, it does indicate that the Virgo data play a key role in the inference of κ . Since more work would be needed to understand the effect of the glitch, this motivated us to consider the effect of excluding this event from the collective analyses above (Sec. IV B). In any case, the outcome of that test showed the role of GW200129 to not be significant (e.g., Fig. 3).

V. DISCUSSION

A. Comparison with previous studies

Our primary constraint on birefringence is given by the combined constraint on the birefringent attenuation from GWTC-3, $\kappa = -0.019_{-0.029}^{+0.038}$ Gpc⁻¹ with 90% credibility, obtained in Sec. IV A (blue band in Fig. 2). We can compare




FIG. 8. GW200129 posterior on κ , luminosity distance d_L , and inclination $\cos \iota$, including different sets of detectors in the analysis per the legend. The marginalized posterior on χ_p is shown in the top right panel. The main run with all three detectors (HLV, filled blue), shows a preference for $\kappa < 0$, as in Fig. 2; this preference is more pronounced for two-detector runs that include Virgo (HV and LV, orange and green); however, it disappears if we remove Virgo (HL, dashed black). The 2D contours correspond to the 90% credible level.

this to previous measurements of amplitude birefringence in the literature. We summarize these constraints in Table III.

1. Okounkova *et al.*

Okounkova *et al.* [16] produced a constraint on frequency-independent GW amplitude birefringence from the distribution of measured inclinations of BBHs in GWTC-2. As opposed to our Eq. (2), this reference parametrized the effect of birefringence as

$$h_{L/R}^{\text{br}}(f) = h_{L/R}^{\text{GR}}(f) \times \exp(\pm \tilde{\kappa} d_C), \quad (12)$$

where $\tilde{\kappa}$ is a frequency-independent attenuation (simply denoted “ κ ” in Ref. [16]). Assuming a fiducial redshift of $z = 0.3$ for all events in GWTC-2, Okounkova *et al.* [16]

TABLE III. Comparison of our constraints with previous studies. CL is the credible level.

	M_{PV} (10^{-21} GeV)	$ \kappa $ (Gpc ⁻¹)	CL (%)
Wang <i>et al.</i> [14]	>0.10	<2.94	90
Okounkova <i>et al.</i> [16]	>0.40	<0.74	68
Zhu <i>et al.</i> [19]	>0.41	<0.72	90
Ng <i>et al.</i> (this work)	>6.75	<0.04	90

obtained a constraint of $|\kappa| \lesssim 0.74 \text{ Gpc}^{-1}$ at 68% credibility. On the other hand, our constraint from Sec. IV A is $|\kappa| < 0.03 \text{ Gpc}^{-1}$ at 100 Hz at 68% credibility. This is a factor of $\sim 25\times$ more stringent than Ref. [16], besides arising from a less simplified model.

The result of Okounkova *et al.* [16] is also phrased in terms of a canonical Chern-Simons length scale,³

$$l_0 = \frac{cd_H\kappa}{3\pi f} = 1400 \text{ km} \left(\frac{\kappa}{1 \text{ Gpc}^{-1}} \right) \left(\frac{100 \text{ Hz}}{f} \right); \quad (13)$$

our 68% credible bound on κ implies $l_0 \lesssim 40 \text{ km}$ at 100 Hz. This compares favorably to the constraint in Okounkova *et al.* [16], which had $l_0 \lesssim 1000 \text{ km}$.

2. Wang *et al.*

Wang *et al.* [14] performed PE on GWTC-1 events with a frequency-dependent model of birefringence resembling ours. Following [22], that reference parametrized the birefringent waveform in terms of some amplitude and phase modifications to the linear polarizations (δh and $\delta\Psi$, respectively), such that

$$h_{+/\times}^{\text{BR}}(f) = h_{+/\times}^{\text{GR}}(f) \mp h_{\times/+}^{\text{GR}}(f)(i\delta h - \delta\Psi), \quad (14)$$

with the plus (minus) sign for the cross (plus) polarization. In terms of circular polarizations, this is equivalent to

$$\begin{aligned} h_{L/R}^{\text{BR}}(f) &= h_{L/R}^{\text{GR}}(f) (1 \mp \delta h \mp i\delta\Psi) \\ &\approx h_{L/R}^{\text{GR}}(f) \exp(\mp \delta h \mp i\delta\Psi), \end{aligned} \quad (15)$$

assuming a small δh and $\delta\Psi$ in the last line. To consider only amplitude birefringence, as in this work, we must compare to the result in Ref. [14] that set $\delta\Psi = 0$ and parametrized $\delta h = \pi f z h_P / M_{\text{PV}}$, where h_P is Planck's constant, and M_{PV} is the energy scale of the birefringent (parity-violating) correction.⁴ Comparing to our Eq. (2) and approximating $z \approx d_C H_0 / c$ via the Hubble constant H_0 , this means that

$$|\kappa| \approx \frac{\pi H_0 h_P (100 \text{ Hz})}{c M_{\text{PV}}} \approx \frac{3.0 \times 10^{-22} \text{ GeV}}{M_{\text{PV}}} \text{ Gpc}^{-1}. \quad (16)$$

Wang *et al.* [14] quote a constraint of $M_{\text{PV}} > 10^{-22} \text{ GeV}$ at 90% credibility, which translates into $|\kappa| \lesssim 2.94 \text{ Gpc}^{-1}$.

³This results from setting the so-called ‘‘canonical’’ Chern-Simons embedding into a Friedmann-Robertson-Walker cosmology, where the Chern-Simons field evolves as $\theta \propto t$ [9,10,56].

⁴Concretely, Wang *et al.* [14] write $\delta h = -A_\nu \pi f$ with $A_\nu = M_{\text{PV}}^{-1} [\alpha_\nu(z=0) - \alpha_\nu(z)(1+z)]$, for α_ν some function of redshift encoding the evolution of a birefringence-mediating field; to produce their constraint, however, they further set $\alpha_\nu(z) = 1$, yielding $A_\nu = -z/M_{\text{PV}}$ and hence $\delta h = \pi f z h_P / M_{\text{PV}}$, multiplying by h_P to obtain the right dimensions.

We obtained a tighter 90% upper limit of $|\kappa| < 0.04 \text{ Gpc}^{-1}$, or $M_{\text{PV}} \gtrsim 6.8 \times 10^{-21} \text{ GeV}$.

3. Zhu *et al.*

A more recent work by Zhu *et al.* [19], which appeared while this manuscript was finalized, performed PE on GWTC-3 events with a frequency-dependent model of amplitude birefringence. Their model is the same as the one used in Wang *et al.* [14], which allows us to compare to our results as we did in Sec. VA 2. Zhu *et al.* [19] reported a constraint of $M_{\text{PV}} > 4.1 \times 10^{-22} \text{ GeV}$ at 90% credibility, which translates into $|\kappa| \lesssim 0.72 \text{ Gpc}^{-1}$ per Eq. (16) above.

This constraint is weaker than ours in spite of the similar number of events considered, a discrepancy that we trace back to inference at the individual-event level. Figure 9 shows the posteriors on M_{PV}^{-1} derived from our analysis via Eq. (16) for each event highlighted in Fig. 1 of Zhu *et al.* [19]. Our posteriors differ from those in Zhu *et al.* [19] in both overall scale and specific shape. For example, the posterior for GW190727_060333 peaks at zero in our analysis, but peaks at a nonzero value in Zhu *et al.* [19]. All distributions are narrower in our analysis by a factor of $\sim 10\times$. This might be explained by a number of analysis differences. The fact that posteriors disagree at the individual-event level suggests that the discrepancy originates in the different choice of parametrization and corresponding priors, in addition to potentially unstated differences in implementation.

B. Parameter degeneracies

In examining our results in Sec. IV C, we identified interactions between birefringence and the inference of source properties like the component spin parameters—in

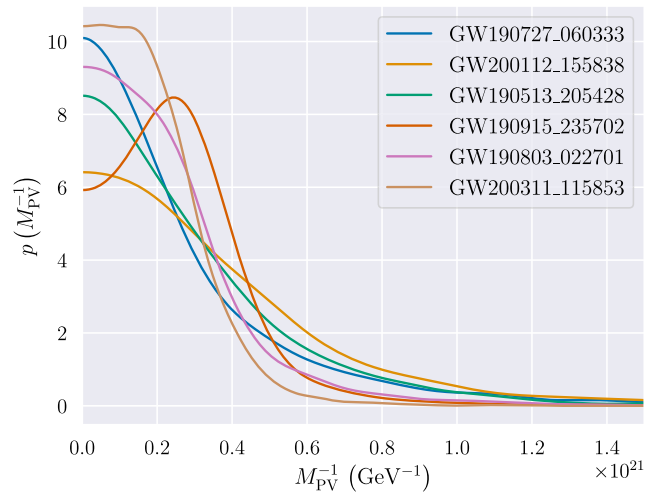


FIG. 9. Individual-event M_{PV}^{-1} posteriors from our analysis, converted from the κ posteriors in Fig. 2 via Eq. (16). Chosen events are those highlighted in Fig. 1 of Zhu *et al.* [19].

addition to expected partial degeneracies with inclination and distance (Sec. II B). In particular, we found that the frequency-dependent modification to the waveform induced by birefringence can effectively mimic spin effects, including precession, especially for heavier systems which are observed over short durations.

As a consequence, a system that would be inferred to be highly spinning in GR may instead be inferred to be low-spinning but highly birefringent if nonzero values of κ are allowed. These morphological degeneracies can interact with the distribution of prior probability mass to yield a preference for nonzero κ in some events: in absence of a strong constraint from the data, it can be more favorable to place a source at far distances with large $|\kappa|$ than at close distances with a fine tuned spin configuration and low $|\kappa|$. The resulting posterior can be unimodal or bimodal depending on the specific phase and polarization states observed at the detector. The events GW170818 (Sec. IV C 1) and GW190521 (Sec. IV C 2) appear to fall into these two categories, respectively.

Conversely, we expect that support for birefringence can also appear for signals that do not necessarily show significant spin effects in the GR analysis, as long as the spins or other source parameters can be tuned to *counteract* the frequency-dependent dephasing that would be otherwise induced by a nonzero κ . In those situations, κ may become at least partially degenerate with inclination, also leading to a bimodal posterior. This is because the intrinsic source parameters may, to an extent, “absorb” the frequency dependence of the birefringent effect, approximately reinstating the degeneracy expected for a single Fourier mode (Sec. II B).

Besides spins, other parameters can interact with κ either directly or through their own couplings with the spins. Since the mass ratio can affect the phasing of a GW inspiral at similar order as the spins [27], we might expect some degree of interaction between q and κ . This seems to indeed be the case for some events, like GW170104 or GW191105, as we demonstrate in Fig. 10. In the former, introducing birefringence increases the preference for asymmetric masses ($q < 1$) while decreasing the inferred level of precession; in the latter, birefringence increases the preference for symmetric masses ($q = 1$) but does not affect the inference on precession. The κ posteriors are bimodal for both events, but do not rule out $\kappa = 0$ (Fig. 2). In neither case is χ_{eff} significantly impacted by birefringence. The coupling between q and κ is likely limited by the fact that changes in the mass ratio alone cannot produce significant amplitude modulations.

Since our inference on κ can be tied to precession and the spin orientations, we might be susceptible to systematics in the modeling of spin angles in IMRPHENOMXPHM and might thus benefit from further analysis with other waveforms like NRSUR7DQ4 [57].

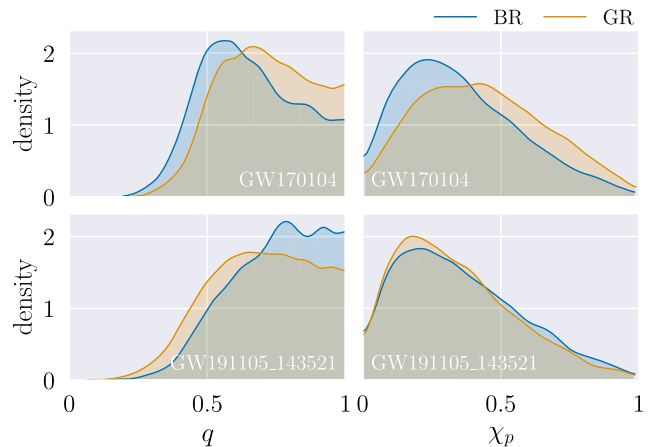


FIG. 10. Posterior on mass ratio $q \equiv m_2/m_1$ (left) and precessing spin χ_p (right) for GW170104 (top) and GW191105 (bottom), as obtained from both the birefringent analysis (orange) and the regular GR analysis (blue).

VI. CONCLUSION

We have reanalyzed all BBHs in GWTC-3 with FAR $< 1/\text{yr}$ to constrain amplitude birefringence in the propagation of GWs. To this end, we implemented a model in which right or left handed polarizations are amplified or suppressed over distance as a function of frequency, with an overall strength parametrized by a birefringent “attenuation” parameter κ , following Eq. (2). This parametrization is consistent with parity-odd theories like Chern-Simons gravity and can be used to constrain them where applicable.

We found no evidence of amplitude birefringence in the GWTC-3 data and constrained $\kappa = -0.019_{-0.029}^{+0.038} \text{ Gpc}^{-1}$ with 90% credibility when treating κ as a global quantity shared by all GWTC-3 events (Sec. IV A). This measurement is significantly more stringent than past constraints (Table III). As an additional check, we implemented a hierarchical analysis that allowed for each event to probe different effective values of κ , as might be the case if birefringence is mediated by a field that is nonuniform over the angular separations between detected sources (Sec. IV B). The result of that analysis was consistent with a vanishing κ for all events in our catalog within 90% credibility, but hinted at some possible variance in the population. Future catalogs with more events will shed light on this feature and enable richer models with explicit correlations between birefringent attenuation and source parameters like sky location [23,58].

From the set of results from individual events, we highlighted two, GW170818 and GW190521, whose posterior manifests interactions between birefringence and our inference of source parameters (Sec. IV C). In particular, we identified the relevance of spins and their (partial) degeneracy with κ , in addition to expected correlations with source inclination and distance. The mass ratio

can also play a role, by coupling with the spins or κ directly (Sec. V B).

Motivated by Chern-Simons gravity, this study only focused on amplitude birefringence. Other parity-violating gravity theories also predict velocity birefringence, which can dominate over amplitude birefringence when both are present [22]. Velocity birefringence was tested with GWTC-3 in Wang *et al.* [59] and Haegel *et al.* [60].

This study was restricted to BBHs because of the expectation that they should dominate the birefringence constraint thanks to their larger redshifts. However, lower mass systems involving neutron stars may also be informative thanks to the wide band of frequencies spanned by their signals, in spite of their closer distances. Analysis of future catalogs will benefit from the inclusion of those events, as well as a much larger number of BBH sources at increasingly greater distances. Future measurements will also be enriched by the addition of KAGRA [3] and LIGO India to the GW detector network, which will allow us to better disambiguate between polarization states and hence between models of birefringent propagation. All our data are available in Ng *et al.* [42].

ACKNOWLEDGMENTS

We thank Macarena Lagos and Nicolás Yunes for helpful discussions. M. I., K. W. K. W., and W. M. F. are funded by the Center for Computational Astrophysics at the Flatiron Institute. The Flatiron Institute provided the computational resources used in this work. This research has made use of data or software obtained from the Gravitational Wave Open Science Center [61], a service of LIGO Laboratory, the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN), and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. K. A. G. R. A. is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan. This paper was compiled using SHOWYOURWORK [62] to facilitate reproducibility.

APPENDIX: EXTENDED RESULTS FOR NOTABLE EVENTS

1. GW170818

Here we provide further details on the GW170818 measurement. Figure 11 displays the posterior on all parameters we consider relevant for this event, of which Fig. 6 in the main text represents a subset. As in that figure, Fig. 11 shows the regular IMRPHENOMXPHM GR analysis (GR; orange) and the birefringent analysis (BR; blue), both of which show notable features.

The GR analysis stands out for its relatively confident identification of the spin angles, $\theta_{1/2}$, $\Delta\phi$, and ϕ_{JL} , as well as the phase and polarization angles, ϕ_{ref} and ψ . Of the former, the two $\theta_{1/2}$ angles encode the tilts of the component BHs with respect to the orbital angular momentum \vec{L} , and $\Delta\phi$ is the angle between the spin projections onto the orbital plane, while ϕ_{JL} is a similar angle separating the projections of \vec{L} and the total angular momentum \vec{J} ; of the latter, ϕ_{ref} is an overall reference phase, and ψ encodes the orientation of the binary within the plane of the sky [21]. All these parameters are anchored to a reference point in the inspiral, which in this case corresponds to the time at which the dominant multipole of the observed GW signal reaches 20 Hz at the detector (spin angles may be better identified by using a more physical reference point [51]). It is unusual for all these angles to be well constrained, which suggests that data for this event display a particular phase and polarization signature.

The joint posterior on the spin magnitudes, χ_1 and χ_2 , disfavors the origin $\chi_1 = \chi_2 = 0$ with high credibility, indicating that at least one of the component BHs must have been highly spinning (this is not apparent from the marginals because the BHs were equal in mass [63]). Furthermore, the fact that the tilts are favored to be close to $\theta_{1/2} \approx \pi/2$ implies that the spins must lie along the orbital plane; this is a restatement of the support for high χ_p in Fig. 6.

On the other hand, the birefringent analysis does not strongly favor high spins; this is evidenced by the joint posterior on χ_1 and χ_2 , which now favors $\chi_1 = \chi_2 = 0$ and assigns relatively low probability densities to higher values. Although the spin tilts are similarly constrained to $\theta_{1/2} \approx \pi/2$, the low spin magnitudes would imply weak if any precession, which is reflected in the support for lower values of χ_p in Fig. 6 of the main text. As discussed in Sec. IV C 1, this can be explained by the introduction of κ , which can result in morphologies that are approximately degenerate with precession over a small number of wave-form cycles (cf. Fig. 1). See the main text for further discussion.

2. GW190521

As in the previous section, here we display additional parameters from the GW190521 measurement in Fig. 12,




FIG. 11. Extended corner plot for GW170818: a supplement to Fig. 6 discussed in Appendix A 1. The shaded regions contain 90% and 39.35% (1σ) of the probability mass. The prior is uniform in all shown quantities except d_L , whose prior corresponds to a distribution uniform in comoving volume.

expanding upon Fig. 7 in the main text. The reference GR analysis, obtained with the IMRPHENOMXPHM waveform, favors near extremal spins, in particular, for the heavier component, and rules out nonspinning objects (i.e., $\chi_1 = \chi_2 = 0$) with $\geq 90\%$ credibility. This can be seen most clearly in the joint posterior for χ_1 and χ_2 , rather than in the respective 1D marginals because the BHs in this system were inferred to have equal masses [63].

Similar to GW170818, the GR analysis for GW190521 also provides an informative measurement of the polarization angle ψ and reference phase ϕ_{ref} . However, the fact that ψ is constrained does not alone imply that the degeneracy between right- and left-handed states is broken (for a nonprecessing source, ψ is related to the difference in phase between the circular polarization modes [21]). Additionally, the posterior for tilt angle of the primary BH θ_1 shows support for an




FIG. 12. Extended corner plot for GW190521: a supplement to Fig. 7 discussed in Appendix A 2. The shaded regions contain 90% and 39.35% (1σ) of the probability mass. The prior is uniform in all shown quantities except d_L , whose prior corresponds to a distribution uniform in comoving volume.

antialigned spin ($\cos \theta_1 \approx -1$). For these events, spin and phase angles are referred to 10 Hz at the detector.

Unlike the GR analysis, the birefringent analysis does not favor extremal spins and is, in fact, consistent with nonspinning objects within 90% credibility (lighter blue region encloses $\chi_1 = \chi_2 = 0$ in the respective panel of Fig. 12). In Fig. 7, this same fact manifested as a χ_{eff}

posterior with higher support for the origin. As we discussed in Sec. IV C 2, this suggests that some of the spin effects in these data are being absorbed by a nonzero κ . Other than the spins, the birefringent analysis differs most notably from the GR one in the inference of the luminosity distance, which is favored to be much greater. See Sec. IV C 2 in the main text for related exposition.

- [1] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
- [2] F. Acernese *et al.* (VIRGO Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
- [3] T. Akutsu *et al.* (KAGRA Collaboration), Overview of KAGRA: Detector design and construction history, *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2021).
- [4] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Tests of general relativity with GW150914, *Phys. Rev. Lett.* **116**, 221101 (2016); *Phys. Rev. Lett.* **121**, 129902(E) (2018).
- [5] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Tests of general relativity with GW170817, *Phys. Rev. Lett.* **123**, 011102 (2019).
- [6] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), Tests of general relativity with GWTC-3, [arXiv:2112.06861](https://arxiv.org/abs/2112.06861) [Phys. Rev. D (to be published)].
- [7] C. M. Will, *Theory and Experiment in Gravitational Physics* (Cambridge University Press, Cambridge, England, 2018).
- [8] A. Lue, L.-M. Wang, and M. Kamionkowski, Cosmological signature of new parity violating interactions, *Phys. Rev. Lett.* **83**, 1506 (1999).
- [9] R. Jackiw and S. Y. Pi, Chern-Simons modification of general relativity, *Phys. Rev. D* **68**, 104012 (2003).
- [10] S. Alexander and N. Yunes, Chern-Simons modified general relativity, *Phys. Rep.* **480**, 1 (2009).
- [11] C. F. Sopuerta and N. Yunes, Extreme and intermediate-mass ratio inspirals in dynamical Chern-Simons modified gravity, *Phys. Rev. D* **80**, 064006 (2009).
- [12] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, [arXiv:2111.03606](https://arxiv.org/abs/2111.03606) [Phys. Rev. X (to be published)].
- [13] K. Yamada and T. Tanaka, Parametrized test of parity-violating gravity using GWTC-1 events, *Prog. Theor. Exp. Phys.* **2020**, 093E01 (2020).
- [14] Y.-F. Wang, R. Niu, T. Zhu, and W. Zhao, Gravitational wave implications for the parity symmetry of gravity in the high energy region, *Astrophys. J.* **908**, 58 (2021).
- [15] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs, *Phys. Rev. X* **9**, 031040 (2019).
- [16] M. Okounkova, W. M. Farr, M. Isi, and L. C. Stein, Constraining gravitational wave amplitude birefringence and Chern-Simons gravity with GWTC-2, *Phys. Rev. D* **106**, 044067 (2022).
- [17] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, *Phys. Rev. X* **11**, 021053 (2021).
- [18] S. Vitale, S. Biscoveanu, and C. Talbot, The orientations of the binary black holes in GWTC-3, [arXiv:2204.00968](https://arxiv.org/abs/2204.00968).
- [19] T. Zhu, W. Zhao, J.-M. Yan, C. Gong, and A. Wang, Tests of modified gravitational wave propagations with gravitational waves, [arXiv:2304.09025](https://arxiv.org/abs/2304.09025).
- [20] R. Abbott *et al.* (KAGRA, VIRGO, and LIGO Scientific Collaborations), Population of merging compact binaries inferred using gravitational waves through GWTC-3, *Phys. Rev. X* **13**, 011048 (2023).
- [21] M. Isi, Parametrizing gravitational-wave polarizations, *Classical Quantum Gravity* **40**, 203001 (2023).
- [22] W. Zhao, T. Zhu, J. Qiao, and A. Wang, Waveform of gravitational waves in the general parity-violating gravities, *Phys. Rev. D* **101**, 024002 (2020).
- [23] J. M. Ezquiaga, W. Hu, M. Lagos, and M.-X. Lin, Gravitational wave propagation beyond general relativity: Waveform distortions and echoes, *J. Cosmol. Astropart. Phys.* **11** (2021) 048.
- [24] L. Jenks, L. Choi, M. Lagos, and N. Yunes, Parameterized parity violation in gravitational wave propagation, *Phys. Rev. D* **108**, 044023 (2023).
- [25] M. Okounkova, L. C. Stein, J. Moxon, M. A. Scheel, and S. A. Teukolsky, Numerical relativity simulation of GW150914 beyond general relativity, *Phys. Rev. D* **101**, 104016 (2020).
- [26] M. Okounkova, M. Isi, K. Chatziioannou, and W. M. Farr, Gravitational wave inference on a numerical-relativity simulation of a black hole merger beyond general relativity, *Phys. Rev. D* **107**, 024046 (2023).
- [27] L. Blanchet, Gravitational radiation from post-Newtonian sources and inspiralling compact binaries, *Living Rev. Relativity* **17**, 2 (2014).
- [28] T. Damour, Coalescence of two spinning black holes: An effective one-body approach, *Phys. Rev. D* **64**, 124013 (2001).
- [29] P. Ajith *et al.*, Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [30] L. Santamaria *et al.*, Matching post-Newtonian and numerical relativity waveforms: Systematic errors and a new phenomenological model for non-precessing black hole binaries, *Phys. Rev. D* **82**, 064016 (2010).
- [31] P. Schmidt, F. Ohme, and M. Hannam, Towards models of gravitational waveforms from generic binaries II: Modelling precession effects with a single effective precession parameter, *Phys. Rev. D* **91**, 024043 (2015).
- [32] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, [arXiv:2108.01045](https://arxiv.org/abs/2108.01045) [Phys. Rev. D (to be published)].
- [33] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog, *Phys. Rev. D* **103**, 122002 (2021).
- [34] H. S. Chia and T. D. P. Edwards, Searching for general binary inspirals with gravitational waves, *J. Cosmol. Astropart. Phys.* **11** (2020) 033.
- [35] R. Magee, M. Isi, E. Payne, K. Chatziioannou, W. Farr, S. Vitale, and G. Pratten, Selection biases in tests of general relativity with gravitational waves (to be published).
- [36] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, *SoftwareX* **13**, 100658 (2021).

- [37] G. Ashton *et al.*, BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [38] LIGO Scientific Collaboration and Virgo Collaboration, GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run—parameter estimation data release, [10.5281/zenodo.6513631](https://arxiv.org/abs/10.5281/zenodo.6513631) (2022).
- [39] LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run—parameter estimation data release, [10.5281/zenodo.5546663](https://arxiv.org/abs/10.5281/zenodo.5546663) (2021).
- [40] G. Pratten *et al.*, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, *Phys. Rev. D* **103**, 104056 (2021).
- [41] B. P. Abbott *et al.* (LIGO Scientific, Virgo, and VIRGO Collaborations), A gravitational-wave measurement of the Hubble constant following the second observing run of advanced LIGO and Virgo, *Astrophys. J.* **909**, 218 (2021).
- [42] T. C. K. Ng, M. Isi, K. W. K. Wong, and W. M. Farr, Constraining gravitational wave amplitude birefringence with GWTC-3: Data release, [10.5281/zenodo.7935107](https://arxiv.org/abs/10.5281/zenodo.7935107) (2023).
- [43] A. Zimmerman, C.-J. Haster, and K. Chatziioannou, On combining information from multiple gravitational wave sources, *Phys. Rev. D* **99**, 124044 (2019).
- [44] M. Isi, W. M. Farr, and K. Chatziioannou, Comparing Bayes factors and hierarchical inference for testing general relativity with gravitational waves, *Phys. Rev. D* **106**, 024048 (2022).
- [45] I. Mandel, Parameter estimation on gravitational waves from multiple coalescing binaries, *Phys. Rev. D* **81**, 084029 (2010).
- [46] D. W. Hogg, A. D. Myers, and J. Bovy, Inferring the eccentricity distribution, *Astrophys. J.* **725**, 2166 (2010).
- [47] M. Isi, K. Chatziioannou, and W. M. Farr, Hierarchical test of general relativity with gravitational waves, *Phys. Rev. Lett.* **123**, 121101 (2019).
- [48] K. W. k. Wong, M. Gabri e, and D. Foreman-Mackey, flowMC: Normalizing flow enhanced sampling package for probabilistic inference in JAX, *J. Open Source Software* **8**, 5021 (2023).
- [49] S. Miller, T. A. Callister, and W. M. Farr, The low effective spin of binary black holes and implications for individual gravitational-wave events, *Astrophys. J.* **895**, 128 (2020).
- [50] T. Callister, Reweighting single event posteriors with hyperparameter marginalization, Technical Report No. LIGO-T2100301, LIGO Laboratory, 2021, <https://dcc.ligo.org/LIGO-T2100301/public>.
- [51] V. Varma, M. Isi, S. Biscoveanu, W. M. Farr, and S. Vitale, Measuring binary black hole orbital-plane spin orientations, *Phys. Rev. D* **105**, 024045 (2022).
- [52] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW190521: A binary black hole merger with a total mass of $150M_{\odot}$, *Phys. Rev. Lett.* **125**, 101102 (2020).
- [53] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Properties and astrophysical implications of the $150M_{\odot}$ binary black hole merger GW190521, *Astrophys. J. Lett.* **900**, L13 (2020).
- [54] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. Di Renzo, and G. Ashton, Subtracting glitches from gravitational-wave detector data during the third LIGO-Virgo observing run, *Classical Quantum Gravity* **39**, 245013 (2022).
- [55] E. Payne, S. Hourihane, J. Golomb, R. Udall, R. Udall, D. Davis, and K. Chatziioannou, Curious case of GW200129: Interplay between spin-precession inference and data-quality issues, *Phys. Rev. D* **106**, 104017 (2022).
- [56] N. Yunes and D. N. Spergel, Double-binary-pulsar test of Chern-Simons modified gravity, *Phys. Rev. D* **80**, 042004 (2009).
- [57] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, Surrogate model of hybridized numerical relativity binary black hole waveforms, *Phys. Rev. D* **99**, 064045 (2019).
- [58] S. Goyal, A. Vijaykumar, J. M. Ezquiaga, and M. Zumalacarregui, Probing lens-induced gravitational-wave birefringence as a test of general relativity, *Phys. Rev. D* **108**, 024052 (2023).
- [59] Y.-F. Wang, S. M. Brown, L. Shao, and W. Zhao, Tests of gravitational-wave birefringence with the open gravitational-wave catalog, *Phys. Rev. D* **106**, 084005 (2022).
- [60] L. Haegel, K. O’Neal-Ault, Q. G. Bailey, J. D. Tasson, M. Bloom, and L. Shao, Search for anisotropic, birefringent spacetime-symmetry breaking in gravitational wave propagation from GWTC-3, *Phys. Rev. D* **107**, 064031 (2023).
- [61] gwosc.org.
- [62] R. Luger, M. Bedell, D. Foreman-Mackey, I. J. M. Crossfield, L. L. Zhao, and D. W. Hogg, Mapping stellar surfaces III: An efficient, scalable, and open-source doppler imaging model, [arXiv:2110.06271](https://arxiv.org/abs/2110.06271).
- [63] S. Biscoveanu, M. Isi, S. Vitale, and V. Varma, New spin on LIGO-Virgo binary black holes, *Phys. Rev. Lett.* **126**, 171103 (2021).