Identifying when precession can be measured in gravitational waveforms

Rhys Green,[1] Charlie Hoy ,[1] Stephen Fairhurst,[1] Mark Hannam,[1,2] Francesco Pannarale,[1,2,3] and Cory Thomas [1]

[1]*School of Physics and Astronomy, Cardiff University, Cardiff, CF24 3AA United Kingdom*
[2]*Dipartimento di Fisica, Università di Roma "Sapienza", Piazzale A. Moro 5, I-00185 Roma, Italy*
[3]*INFN Sezione di Roma, Piazzale A. Moro 5, I-00185 Roma, Italy*

In binary-black-hole systems where the black-hole spins are misaligned with the orbital angular momentum, precession effects leave characteristic modulations in the emitted gravitational waveform. Here, we investigate where in the parameter space we will be able to accurately identify precession, for likely observations over coming LIGO-Virgo-KAGRA observing runs. Despite the large number of parameters that characterize a precessing binary, we perform a large scale systematic study to identify the impact of each source parameter on the measurement of precession. We simulate a fiducial binary at moderate mass-ratio, signal-to-noise ratio (SNR), and spins, such that precession will be clearly identifiable, then successively vary each parameter while holding the remaining parameters fixed. As expected, evidence for precession increases with signal-to noise-ratio (SNR), higher in-plane spins, more unequal component masses, and higher inclination, but our study provides a *quantitative* illustration of each of these effects, and informs our intuition on parameter dependencies that have not yet been studied in detail, for example, the effect of varying the relative strength of the two polarizations, the total mass, and the aligned-spin components. We also measure the "precession SNR" ρ_p, to quantify the signal power associated with precession. By comparing ρ_p with both Bayes factors and the recovered posterior distributions, we find it is a reliable metric for *measurability* that accurately predicts when the detected signal contains evidence for precession.

I. INTRODUCTION

In September 2015, the first direct detection of gravitational-wave (GWs) marked the beginning of GW astronomy [1]. Another 14 detections have been announced by the LIGO Scientific and Virgo collaborations (LVC), the vast majority of which were due to black-hole (BH) mergers [2–7]. Additional events have also been reported by independent groups [8–11]. These GW observations have already provided significant insights into gravitational physics, cosmology, astronomy, nuclear physics and fundamental physics (see e.g., Refs. [12–19]). With an order of magnitude more observations expected over the next 5–10 years, as the sensitivities of the LIGO [20,21], Virgo [22] and KAGRA [23] detectors improve and additional detectors come online, GW astronomy from compact-binary mergers has the potential to transform our understanding of gravitational and fundamental physics [24–26].

Everything we learn from GW binary-black-hole (BBH) observations is a consequence of a detailed parameter estimation analysis that extracts the source parameters of the binary. While some parameters are extracted with good precision, inspiral dominanted signals show strong correlations between certain parameters which means that they cannot be measured so accurately, for example correlations between the binary's distance and inclination [27–29], the two masses [27,30], and the mass-ratio and spin components aligned to the binary's orbital angular momentum [30–33]. As well as studies of the inspiral, work has been done to extract the source properties for high mass signals dominated by the merger ringdown, see e.g., [34–37].

Spin components misaligned with the binary's orbital angular momentum, leading to a precession of the binary's orbital plane and hence modulations of the amplitude and phase, have not yet been unambiguously measured in GW observations [2], see Fig. 1. Precession effects and correlations with other parameters are understood in principle [38,39] but since theoretical signal models of precessing binaries that include the merger and ringdown date from only shortly before the first detections [40,41], we have less experience of when precession will be measurable, and what the impact will be on other parameter measurements.

The purpose of this paper is to explore when precession will be measurable, and its impact on other parameter measurements, in the kind of configurations that are representative of expectations from binary populations based on LIGO-Virgo-KAGRA observations to date [2].
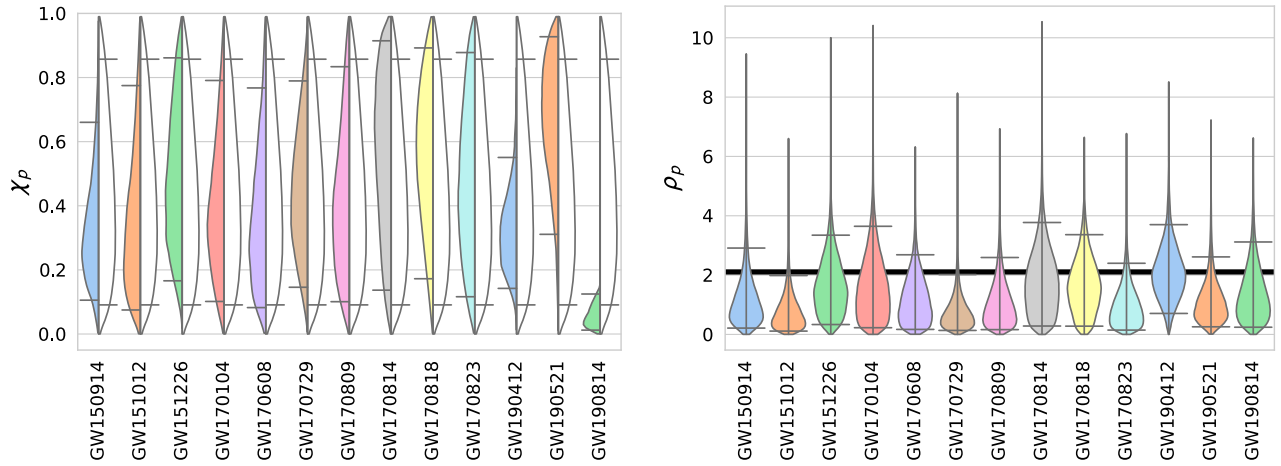
FIG. 1. Plot showing the posterior distributions for χ_p and ρ_p for all LIGO/Virgo BBH observations. The χ_p posterior distribution (left hand side, colored) is compared to its prior (right hand side, white) in the form of a split violin plot. The ρ_p posterior distribution is shown as a single violin. Horizontal grey lines show the 90% symmetric credible interval. The solid black line shows the $\rho_p = 2.1$ threshold. Bounded KDEsare used for estimating the probability density. Data obtained from the Gravitational Wave Open Science Center [42].

By utilizing the precession signal-to-noise ratio (SNR) ρ_p [43,44] as a quantifier for the measurability of precession, we also verify that ρ_p is indeed a good metric for the measurability of precession across the vast majority of the parameter space, and relate it to the standard means to identify the presence of precession, the Bayes factor. In doing so, we show that computationally expensive parameter estimation runs can be avoided by simply calculating the precession SNR.

Previous work has explored the general phenomenology of precession effects: its increased measurability with large in-plane spins [45–47], large mass ratios [45,46], high inclination [36,38,44,48–50], and of course high SNR [45,46,51]. Beyond these general expectations, the *quantitative* behavior of parameter measurements in the presence of precession has not been studied in great detail for typical LIGO-Virgo-KAGRA observations. The measurability of precession for high mass ratio LIGO-Virgo-KAGRA observations like GW190814 has been investigated in recent work [52].

In this paper, we focus on the region of parameter space most likely to yield binaries with observable precession: binaries of comparable mass, with moderate in-plane and aligned-spin components [44]. We perform a series of one-dimensional investigations of the parameter space, in which we vary one parameter at a time: total mass, mass ratio, spins (both in-plane as characterized by χ_p, and the aligned spin combination $\chi_{\rm eff}$), the binary orientation (both the inclination of the orbit and also binary polarization), and the sky location and show the impact of varying each of the binary parameters individually. These investigations serve to confirm that much of the known phenomenology is apparent even at relatively low SNR, while also

demonstrating that the precession SNR can be effectively used across a significant fraction of the parameter space to *predict* the observable consequences of precession *without* the need for computationally costly parameter estimation analyses.

This paper is structured as follows: Sec. II provides an introduction to precession, a brief recap of the two-harmonic approximation that allows us to define ρ_p, and a summary of the importance of precession across the parameter space. Section III provides an introduction to the parameter estimation techniques used here, and parameter estimation results and interpretation for our fiducial system. In Sec. IV we perform a series of one-dimensional explorations of the parameter space. In Sec. VI we compare the predicted precession SNR with observations and in Sec. V we compare precession SNR with the Bayes factors between precessing and nonprecessing runs. We conclude with a summary and discussion of future directions.

II. BLACK HOLE SPIN INDUCED PRECESSION

A binary consisting of two compact objects will slowly inspiral due to the emission of GWs. Assuming that the binary is on a quasi-spherical orbit, it may be described by the individual component masses, m_1 and m_2 (where we define $m_1 > m_2$ and we denote the mass ratio to be $q = m_1/m_2 > 1$), and their spin angular momenta $\mathbf{S}_1$ and $\mathbf{S}_2$.

For the case where the total spin is misaligned with the total orbital angular momentum, $\mathbf{S}_i \times \mathbf{L} \neq 0$, the system undergoes spin-induced precession. In most cases $L \ll J$, and the system undergoes "simple precession," where the orbital angular momentum precesses around the
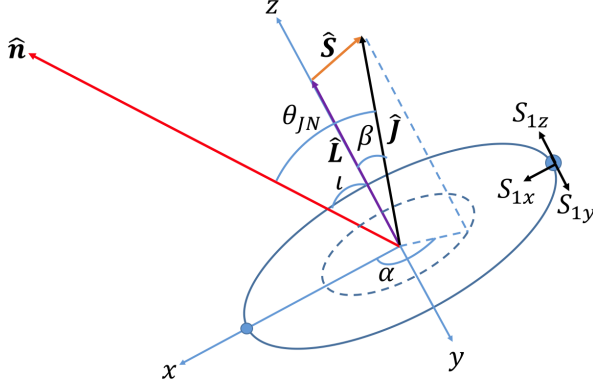
FIG. 2.   Plot showing how the precession angles used in this study are defined in the coprecessing frame. The normal vector $\hat{\mathbf{n}}$ here indicates the line of sight of the observer, $\hat{\mathbf{L}}$ and $\hat{\mathbf{J}}$ are the orbital angular momentum and total angular momentum vectors respectively, S_{1x}, S_{1y} and S_{1z} are the x, y and z components of the spin on the larger black hole. Note that $\mathbf{J}$, $\mathbf{L}$ and $\hat{\mathbf{n}}$ are shown to be co-planar only for ease of illustration. It is not true in general.

(approximately constant) total angular momentum, $\mathbf{J} = \mathbf{S_1} + \mathbf{S_2} + \mathbf{L}$, and the spins precess such that $\dot{\mathbf{S}} = -\dot{\mathbf{L}}$, where $\mathbf{S} = \mathbf{S_1} + \mathbf{S_2}$ [38,39].

The strength of precession is characterised by the tilt angle of the binary's orbit, β, defined as the polar angle between $\mathbf{L}$ and $\mathbf{J}$; see Fig. 2. β is determined primarily by the total spin in the plane, and the binary's mass ratio and separation. At leading order the magnitude of the binary's orbital angular momentum is given by $L = \mu\sqrt{Mr}$, and we can write,

$$\cos\beta = \frac{\mu\sqrt{Mr} + S_{\parallel}}{[(\mu\sqrt{Mr} + S_{\parallel})^2 + S_{\perp}^2]^{1/2}}, \qquad (1)$$

where $M = m_1 + m_2$ is the total mass of the source, $\mu = m_1 m_2/M$ is the reduced mass, r is the separation and $S_{\parallel}$ and $S_{\perp}$ are the total spin parallel and perpendicular to L respectively. In general, the larger the "opening angle," β, the more prominent the precession effects.

In simple precession cases, β slowly increases during inspiral. This temporal variation is mostly dependent on the variation of r as, throughout the portion of the binary's inspiral that is visible in the GW detectors, μ, M, $S_{\parallel}$ and $S_{\perp}$ are all approximately constant. Therefore, β also typically varies very little over the duration that is visible in the GW detectors, and it is possible to make the simplifying assumption that β remains constant. This assumption has been used to good effect in previous work, e.g., Ref. [48], although in this work we make no assumptions about β.

It is often convenient to quantify the binary's in-plane spin by the scalar quantity χ_p [53] (see Refs. [54,55] for alternative measures). χ_p estimates a time-average of the in-plane spin components that drive precession, motivated by the leading-order PN precession equations [38,39], and is defined as,

$$\chi_p = \frac{1}{A_1}\max\left(A_1 S_{1\perp}, A_2 S_{2\perp}\right), \qquad (2)$$

where $A_1 = 2 + 3/(2q)$ and $A_2 = 2 + 3q/2$ and $\mathbf{S}_{i\perp}$ is the component of the spin perpendicular to $\mathbf{L}$ on the ith black hole. χ_p as defined above, takes values between 0 (non-precessing system) and $+1$ (maximally precessing system).

A. Two harmonic approximation

Non-zero values of β and χ_p indicate that a binary is precessing, and, if all other parameters are kept constant, they parameterise the amount of precession in a system. However, they are not sufficient to tell us whether precession will be *observable*. Precession appears in the signal as modulations of the amplitude and phase, but these also depend on the binary orientation and signal polarization. References [43,44] introduce a method for decomposing a precessing waveform into a series of five *nonprecessing harmonics*, where the characteristic modulations of a precessing signal are caused by the beating of these *harmonics*. The harmonics form a power series in the parameter,

$$b = \tan(\beta/2). \qquad (3)$$

In most regions of parameter space, the two leading harmonics (a leading-order term independent of b, and a first-order term proportional to b) are sufficient to capture the significant precession features in the waveform, and the other harmonics can be neglected. As discussed in detail in Sec. V of Ref. [43]), these other precession harmonics can be ignored for binaries where $\beta \lesssim 45°$. In general this is true for all binaries, apart from relatively extreme systems, for example, those that have θ_{JN} close to edge-on *and* are either highly precessing or have very large negative spins.

Thus, for almost all signals we expect to observe, the waveform can faithfully be expressed as,

$$h(f) \approx \mathcal{A}_0 h^0(f) + \mathcal{A}_1 h^1(f), \qquad (4)$$

where $\mathcal{A}_0$ and $\mathcal{A}_1$ are complex, orientation dependent amplitudes, and $h^0(f)$ and $h^1(f)$ are the waveforms of the two leading harmonics. For a detector with a one-sided noise spectral density of $S_n(f)$, the relative amplitude of the harmonics is given by

$$\bar{b} := \frac{|h^1|}{|h^0|} = \sqrt{\frac{\int df \frac{|h_1(f)|^2}{S_n(f)}}{\int df \frac{|h_0(f)|^2}{S_n(f)}}}. \qquad (5)$$

which is the average value of b over the frequency range defined by our starting frequency (20 Hz) to our Nyquist

Frequency (1024 Hz), weighted by the signal strength in the detector. The complex amplitudes $\mathcal{A}_{0,1}$ depend upon the extrinsic parameters of the binary: the distance d_L, angle between $\mathbf{J}$ and the line of sight $\hat{\mathbf{n}}$, polarization angle ψ, reference phase ϕ_0, and the reference precession phase ϕ_{JL}.[1] It is convenient to introduce a reference distance d_0, which is incorporated into the definitions of the harmonics h_0 and h_1. The amplitudes are then defined as,[2]

$$\mathcal{A}_0 = \frac{d_0}{d_L}\left[F_+\left(\frac{1+\cos^2\theta_{JN}}{2}\right) + iF_\times\cos\theta_{JN}\right]e^{-i(2\phi_o+2\phi_{JL})},$$

$$\mathcal{A}_1 = \frac{d_0}{d_L}[F_+\sin2\theta_{JN} + 2iF_\times\sin\theta_{JN}]e^{-i(2\phi_o+\phi_{JL})}. \quad (6)$$

The relative significance of the two harmonics is encoded by

$$\zeta := \frac{\bar{b}\mathcal{A}_1}{\mathcal{A}_0}$$

$$= \bar{b}e^{i\phi_{JL}}\left(\frac{F_+\sin2\theta_{JN} + 2iF_\times\sin\theta_{JN}}{\frac{1}{2}F_+(1+\cos^2\theta_{JN}) + iF_\times\cos\theta_{JN}}\right), \quad (7)$$

where the detector response $F_{+,\times}$ is calculated using the polarization angle appropriate for a coordinate system defined with the z-axis along the direction of total angular momentum $\mathbf{J}$.

The observability of precession can then be characterized by the precession SNR ρ_p, defined as the SNR in the weaker of the two harmonics, and also expressible as a fraction of $\rho = |h|$, the total signal SNR,

$$\rho_p = \mathrm{Min}(|\mathcal{A}_0 h^0|, |\mathcal{A}_1 h^1|)$$

$$= \rho\left(\frac{\min(1,|\zeta|)}{\sqrt{1+|\zeta|^2}}\right). \quad (8)$$

Equation (8) assumes that the two harmonics are orthogonal. If the two harmonics are not orthogonal, such that the overlap between the two harmonics is nonzero, it is necessary to consider only the orthogonal part of the less significant harmonic when calculating the precession SNR. This introduces additional overlap terms between the two harmonics in Eq. (8), see Ref. [43] for details. For simplicity of presentation, in this work all equations are presented under the assumption that the two harmonics are orthogonal, unless otherwise stated. The majority of cases considered in this paper have a small overlap between the two harmonics (less than 10%). However, in the analysis we always orthogonalize the harmonics for all calculations.

---

[1]This is equivalent to α_0 in Ref. [43].
[2]These expressions are equivalent to Eq. (33) in [43], which are restricted to the special case where $F_+ = \cos2\psi$ and $F_\times = -\sin2\psi$. More generally, the detector response also depends upon the sky location in addition to the polarization.

At higher masses, the two harmonics have significant overlap and we discuss the impact of this when we vary the total mass of the binary below.

The quantity ρ_p parametrizes the *observable* precession, it is therefore the *measured* quantity in the data. By considering what we actually measure in the data we are able to understand many of the correlations and degeneracies in the physical parameters that are *caused* by the presence of (or lack of) measurable precession.

In the absence of precession, ρ_p^2 will be χ^2 distributed with two degrees of freedom. Consequently, if there is no observable precession in the system, $\rho_p \geq 2.1$ in only 10% of cases. Thus far we have used $\rho_p = 2.1$ as a simple threshold to determine if there is any observable precession in the system. We revisit this in more detail in Sec. VI B.

B. Observability of precession

The strength of the modulations in the GW signal depend primarily on the opening angle, β, and this is reflected in the expansion parameter b in the two-harmonic approximation; the precession frequency $\dot{\alpha}$ also plays a role. The strength of the modulations in the *observed* signal also depend on the binary's inclination to the observer, θ_{JN}, and the detector polarization ψ, and these are all incorporated into the precession SNR ρ_p, through Eqs. (7) and (8). From these we can draw immediate conclusions about the scenarios in which precession will be most easily measured. These observations are in general not new (see, as always, the pioneering discussions in Refs. [38,39]), but we summarize them here and, where salient, present them in terms of the two harmonic formalism, which highlights the insights and intuition that are simplified in this formulation. We then compare these expectations with the quantitative results that we find in our full parameter estimation study.

Our first basic picture of the strength of precession effects comes from Eq. (1), which gives the dominant effect on β during the inspiral. If we first consider cases where the spin is entirely in the orbital plane, i.e., $S_\| = 0$, we see that the opening angle β will be zero if $S_\perp = 0$ (as we would expect), and increases linearly for small $S_\perp$. The opening angle also increases as μ decreases, i.e., as the mass ratio is increased. Equation (1) is no longer accurate near merger, and for equal-mass systems β does not become large, but for large mass ratios the opening angle can approach 90°.

If we now consider nonzero $S_\|$, we see that the level of precession will be reduced for systems with a positive aligned-spin component, and will be increased for systems with a negative aligned-spin component. The importance of this effect will depend on the other terms, but we can see that for a high-mass-ratio system where μ is very small, and close to merger, so rM is also small, the aligned-spin component will have a strong effect on β, and therefore the measurability of precession. A negative $S_\|$ is necessary to achieve $\beta > 90°$, and for large mass-ratio systems near

TABLE I.   Table showing the simulated and inferred parameters for the "standard" injection when recovered by a nonprecessing (IMRPhenomD) and a precessing (IMRPhenomPv2) waveform model. We report the median values along with the 90% symmetric credible intervals and the maximum likelihood (maxL) value.

	Simulated	Median		maxL	
	Precessing	Nonprecessing	Precessing	Nonprecessing	Precessing
Total mass $M/M_\odot$	40.0	40^{+3}_{-2}	40^{+4}_{-2}	40.161	40.507
Chirp mass $\mathcal{M}/M_\odot$	16.22	$16.5^{+0.3}_{-0.2}$	$16.3^{+0.3}_{-0.3}$	16.459	16.113
Mass ratio q	2.0	$1.8^{+0.8}_{-0.7}$	$1.9^{+0.9}_{-0.7}$	1.895	2.191
Inclination angle $\theta_{JN}/°$	60.0	110^{+50}_{-100}	120^{+40}_{-90}	30.0	40.0
Precession phase $\phi_{JL}/°$	45.0	…	200^{+100}_{-200}	…	80.0
Effective aligned spin, $\chi_{\rm eff}$	0.0	$0.044^{+0.099}_{-0.084}$	$-0.005^{+0.098}_{-0.092}$	0.06	-0.011
Effective precessing spin, χ_p	0.4	…	$0.5^{+0.4}_{-0.3}$	…	0.554
Right ascension RA/rad	1.88	3^{+3}_{-3}	3^{+3}_{-3}	1.418	1.325
Declination DEC/rad	1.19	$0.2^{+1.0}_{-1.2}$	$0.2^{+1.0}_{-1.2}$	1.229	1.221
Luminosity distance $d_L/$Mpc	223	500^{+200}_{-200}	400^{+200}_{-200}	451.834	372.706
Network SNR ρ	20.0	$19.3^{+0.1}_{-0.2}$	$19.7^{+0.2}_{-0.2}$	19.52	19.936
Precessing SNR ρ_p	5.05	…	4^{+2}_{-2}	…	4.649

merger (small μ and rM) and large negative $S_\parallel$, β can approach 180°, but such systems will be rare.

The measurability of precession also depends on the orientation of the binary with respect to the detector, θ_{JN}. As we see in Eq. (7), precession effects will be minimal if $\theta_{JN} \sim 0°$ or 180°, i.e., the observer views the system from the direction of $\hat{\mathbf{J}}$. We expect precession to be strongest in the observed waveform for orientations close to $\theta_{JN} \sim 90°$. Additionally, when the detector, or network is primarily sensitive to the $\times$ polarization, precession effects will be more significant. The amplitude of the $k = 1$ harmonic vanishes in the $+$ polarization for both face on $\theta_{JN} = 0°$ and 180° and edge-on $\theta_{JN} = 90°$ systems, while the $\times$ polarization is maximal for edge-on systems. Additionally, the $\times$ polarization for the $k = 0$ harmonic vanishes for edge on systems, while the $+$ polarization is only reduced by a factor of two. Thus, even when b is small, there can be observable precession when the system is close to edge on and the network is preferentially sensitive to the $\times$ polarization. For a given choice of masses and spins, the maximum precession SNR is $\rho_p = \rho/\sqrt{2}$.

III. PARAMETER ESTIMATION RESULTS

A. Standard configuration

We begin by describing the results of the parameter recovery routine for a specific simulated signal. The details of the signal are given in Table I. These parameters were chosen so that precession effects would be significant in the observed waveform while still being consistent with the observed population of BBHs. In the following sections, we vary over the parameters of the signal one-by-one to investigate the impact of each parameter on the observability of precession and the accuracy of parameter recovery. For each parameter, we are able to both increase and decrease the significance of precession.

By taking the inferred properties of the BBHs observed in the first, second and third observing runs [56], it is predicted that 90% of detected binaries will have mass ratios $q < 4$ and $\sim$97% of BHs in these binaries will have masses less than $45\ M_\odot$ [57]. Our "standard" simulated signal was chosen to have total mass $M = 40\ M_\odot$ and mass ratio $q = 2$ inclined at an angle of $\theta_{JN} = 60°$. This corresponds to component masses of $26.7\ M_\odot$ and $13.3\ M_\odot$. This mass ratio and inclination was chosen to increase the observability of precession.

Of the 50 events reported by the LIGO/Virgo, 13 exclude the aligned-spin measure $\chi_{\rm eff} = 0$ at 90% confidence [2,56,58]. The other 37 observations peak at $\chi_{\rm eff} = 0$ [2,56]. Based on this, studies have shown that it is likely BHs in binaries have low spin magnitudes [19,44,59,60]. For this reason, in our standard configuration the BH spins were chosen such that there is zero spin aligned with the binary's orbital angular momentum, $\chi_{\rm eff} = 0$. We introduce precession by giving the more massive BH a spin of 0.4 in-plane and leaving the second BH with zero spin; two-spin effects are generally far weaker than the dominant precession effect, which exhibits the same phenomenology as a single-spin system [53,61]. From Eq. (2) we see that this gives us a system with $\chi_p = 0.4$. The opening angle for the binary when the signal enters the detector's sensitivity band is 10° and the average value of the parameter $b = \tan(\beta/2)$ is $\bar{b} = 0.11$, from Eq. (5). The signal is simulated using the IMRPhenomPv2 waveform model that incorporates precession effects, but not higher harmonics ($\ell > 2$) in the signal [40,62].

Our "standard" simulated signal was chosen to be more favorable to precession measurements than typical LIGO-Virgo observations. Assuming systems are distributed

uniformly in binary orientation, masses drawn from a power law distribution and spins drawn from a low isotropic distribution (see Ref. [44] for details), we expect that 4 in every 100 binaries detected by LIGO-Virgo will be inclined at angles greater than $60°$ and have $\bar{b} > 0.11$.

The sky location of the binary was chosen to have RA = 1.88 rad, DEC = 1.19 rad. The coalescence time is $t = 1186741861$ GPS (corresponding to the merger time of GW170814 [63]). The polarization angle, defined by the orientation of the orbital plane when entering the sensitive band at 20 Hz, is $\psi = 40°$. The two harmonic approximation is calculated in the J-aligned frame ($\hat{z} = \hat{\mathbf{J}}$). In this frame, the polarization angle is $\psi_J = 120°$, which gives antenna factors for H1 of $F_+ = 0.34$ and $F_\times = 0.53$ and for L1 of $F_+ = -0.45$ and $F_\times = -0.30$, thus both detectors are roughly equally sensitive to the two GW polarizations.

We injected the signals into zero noise. The zero-noise analysis results will be similar to those obtained from the average results of multiple identical injections in different Gaussian noise realizations. The simulated signal is recovered using the LIGO Livingston and Hanford detectors with sensitivities matching those achieved in the second observing run (O2) [2]. A low frequency cut-off of 20 Hz was used for likelihood evaluations, this frequency is also used as the reference frequency when defining all frequency dependent parameters such as θ_{JN}. Both the LIGO Livingston and Hanford sensitivities improved prior to the third observing run (O3) [64] and are expected to improve further prior to the fourth observing run (O4) [65]. The results presented in this work are unlikely to be affected significantly by these changes and therefore we expect the main conclusions to be valid for O4 and beyond.

The SNR of the signal is fixed to be 20, corresponding to a moderately loud signal for aLIGO and AdV observations [65]. This sets the distance to $d_L = 223$ Mpc. The simulated SNR in the two detectors is 16.2 in L1 and 11.7 in H1. The simulated precession SNR in each of the detectors is 3.7 and 3.4 respectively, giving a network precession SNR of 5.0. Thus, we expect that precession will be clearly observable in this signal.

B. Parameter estimation techniques

We will adopt a parameter estimation methodology that uses matched filtering with phenomenological gravitational waveforms and Markov Chain Monte Carlo (MCMC) techniques to sample the posterior.

We begin by introducing the matched filtering formalization for parameter estimation. We assume that the time series received from the GW detectors can be decomposed as a sum of the GW signal, $h(t)$, plus noise, $n(t)$, which is assumed stationary and Gaussian with zero mean,

$$d(t) = h(t) + n(t). \tag{9}$$

Under the assumption of Gaussian noise, the probability of observing data d given a signal $h(\lambda)$ parametrized by $\lambda = \{\lambda_1, \lambda_2, ..., \lambda_N\}$, otherwise known as the likelihood, is [66],

$$p(d|\lambda) \propto \exp\left(-\frac{1}{2}\langle d - h(\lambda)|d - h(\lambda)\rangle\right), \tag{10}$$

where $\langle a|b\rangle$ denotes the inner product between two waveforms a and b and is defined as,

$$\langle a|b\rangle = 4\mathrm{Re}\int_0^\infty \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)} df, \tag{11}$$

where $S_n(f)$ is the one-sided power-spectral density (PSD) and $\tilde{a}$ denotes the Fourier transform of the gravitational waveform a.

The posterior probability density function (PDF) can then be computed through a simple application of Bayes' theorem,

$$
\begin{aligned}
p(\lambda|d) &= \frac{p(\lambda)p(d|\lambda)}{p(d)}, \\
&\propto p(\lambda)\exp\left(-\frac{1}{2}\langle d - h(\lambda)|d - h(\lambda)\rangle\right),
\end{aligned}
\tag{12}
$$

where $p(\lambda|d)$ is the posterior distribution for the parameters λ, $p(\lambda)$ is the prior probability distribution where $\int p(\lambda)d\lambda = 1$, and $p(d)$ is the marginalized likelihood where $p(d) = \int p(\lambda_i)p(d|\lambda_i)d\lambda_i$. Posterior distributions for specific parameters can then be found by marginalizing over all other parameters,

$$p(\lambda_i|d) = \int p(\lambda|d)d\lambda_1...d\lambda_{i-1}d\lambda_{i+1}...d\lambda_N. \tag{13}$$

In the idealized situation of zero noise, Eq. (10) has a maximum at $h(\lambda) = h(\lambda_0)$. However, as can be seen in Eq. (12) the posterior also includes priors, this means that, as well as effects due to noise, certain priors may cause the maxima to be deflected away from $h(\lambda) = h(\lambda_0)$. This would then lead to Eq. (13) recovering a biased posterior. In this work, we consider the effect of three closely related priors,

(i) *Global*: the prior used during the parameter estimation analysis. This reflects our prior belief before observing any data,

(ii) *Conditioned*: the global prior conditioned upon the posterior distributions of other parameters from the same analysis. For example since χ_{eff} and χ_{p} are correlated, any informative measurement of χ_{eff} modifies our prior beliefs about χ_p. This prior has been used in previous LVC publications, see e.g., [2],

(iii) *Informed*: the global prior conditioned upon the posterior distributions from a different analysis. Here, we use this to inform our expectations of the degree of precession given the results from a nonprecessing analysis. See Sec. VI for details.

C. Parameter recovery

We performed parameter estimation on the signal using the `LALInference` [67] and `LALSimulation` libraries within `LALSuite` [68]. Parameter recovery was performed with the `IMRPhenomPv2` model [40,62], which matches the simulated signal to remove any systematic error caused by waveform uncertainty, and the corresponding `IMRPhenomD` aligned-spin waveform model [69,70], which does not include any precession effects. Additionally, all analyses used exactly the same priors as those used in the LIGO-Virgo discovery papers, for details, see Appendix B.1 of [2]. All post-processing was handled by the `PESummary` PYTHON package [71].

Table I summarizes the key results for the standard configuration. All uncertainties are the 90% symmetric credible intervals.

We begin by comparing the overall differences between parameter recovery with the precessing, `IMRPhenomPv2`, and nonprecessing, `IMRPhenomD`, runs. From the table, we see that the maximum likelihood SNR for the nonprecessing model is, as expected, lower than for the precessing waveforms. This can be easily understood from the two-harmonic approximation. Since the precessing waveform is well approximated by the sum of two nonprecessing harmonics, we would expect the nonprecessing recovery to accurately recover the more significant of these two. If that were the case, the we would expect that,

$$\rho_D^2 \approx \rho^2 - \rho_p^2, \tag{14}$$

and this is indeed the case, as $\rho_D = 19.52$, $\rho = 19.94$ and the recovered power in the second harmonic is $\rho_p = 4.6$. Furthermore, we see that the recovered waveforms confirm this expectation: the recovered waveform when we include precession matches well with the simulated signal, while the nonprecessing run recovers a waveform that matches the dominant harmonic, as show in Fig. 3.

We first consider the accuracy with which the masses and (aligned) spins are recovered. As expected, the chirp mass of the system is well recovered, in that it matches the simulated value with only a 2% uncertainty, which remains constant for both precessing and nonprecessing runs. As is well known, there is a degeneracy between mass-ratio and spin, particularly during the inspiral part of the waveform [30–33], which leads to significant uncertainty in both parameters. In Fig. 4 we show the recovery of the mass ratio and spin, for both precessing and nonprecessing runs. When the model used to recover includes precession effects, the peak of the posteriors is located close to the
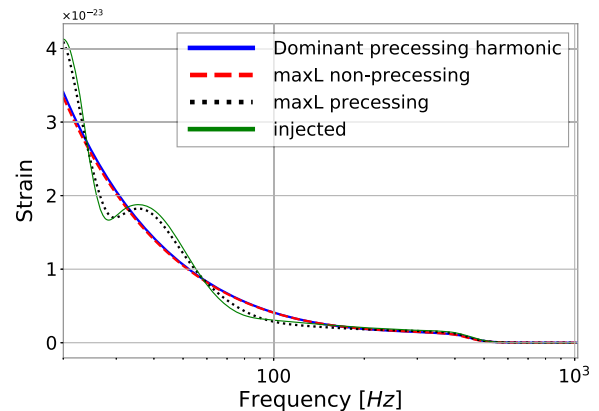

FIG. 3. Comparison of the simulated precessing (green), nonprecessing maximum likelihood (red), precessing maximum likelihood (black) and dominant precessing harmonic (blue) waveforms as a function of frequency. Waveforms are projected onto the LIGO Hanford detector.

simulated value ($\chi_{\text{eff}} = 0$ and $q = 2.0$) and, while the degeneracy leads to significant uncertainty in both parameters, the mass-ratio distribution is clearly peaked away from $q = 1$. Interestingly, when we recover with a nonprecessing waveform model, the inferred *aligned* spin component is systematically offset, with a peak at $\chi_{\text{eff}} \approx 0.05$. This can be understood by recalling that precession induces a secular drift in the phase evolution of the binary, and this can be mimicked by a change in the value of the aligned spin [38,43]. This discrepancy has not been seen in LIGO/Virgo observations [2] as we have not observed any systems with significant ρ_p (see Fig. 1). We investigate this further in Sec. IV D, where we study the effect of varying the mass ratio.

For nonprecessing binaries, it is generally not possible to accurately recover the distance and orientation of the source, due to a well known degeneracy (see e.g., Ref. [29] for details), although the observation of higher signal harmonics can break this degeneracy through an independent measurement of the source inclination [27,29,72]. Similarly, the observation of precession can break this degeneracy [73]. Precession causes an oscillation of the orbital plane leading to a time-dependence of the orientation of the orbital plane relative to the line of sight. Equivalently, in the two-harmonic picture, precession leads to the observation of a second harmonic and, consequently, additional constraints on the binary orientation as the amplitudes of the harmonics depend upon the viewing angle. In Fig. 4, we show the inferred two-dimensional distance and inclination posteriors for the precessing and nonprecessing runs. As expected, the precessing run constrains the source to be away from face-on, while the nonprecessing run simply returns the prior. However, even with observable precession, the simulated distance and orientation are not accurately recovered—a significant fraction of the posterior support is for a system at a greater distance and oriented closer to face-on. We will see
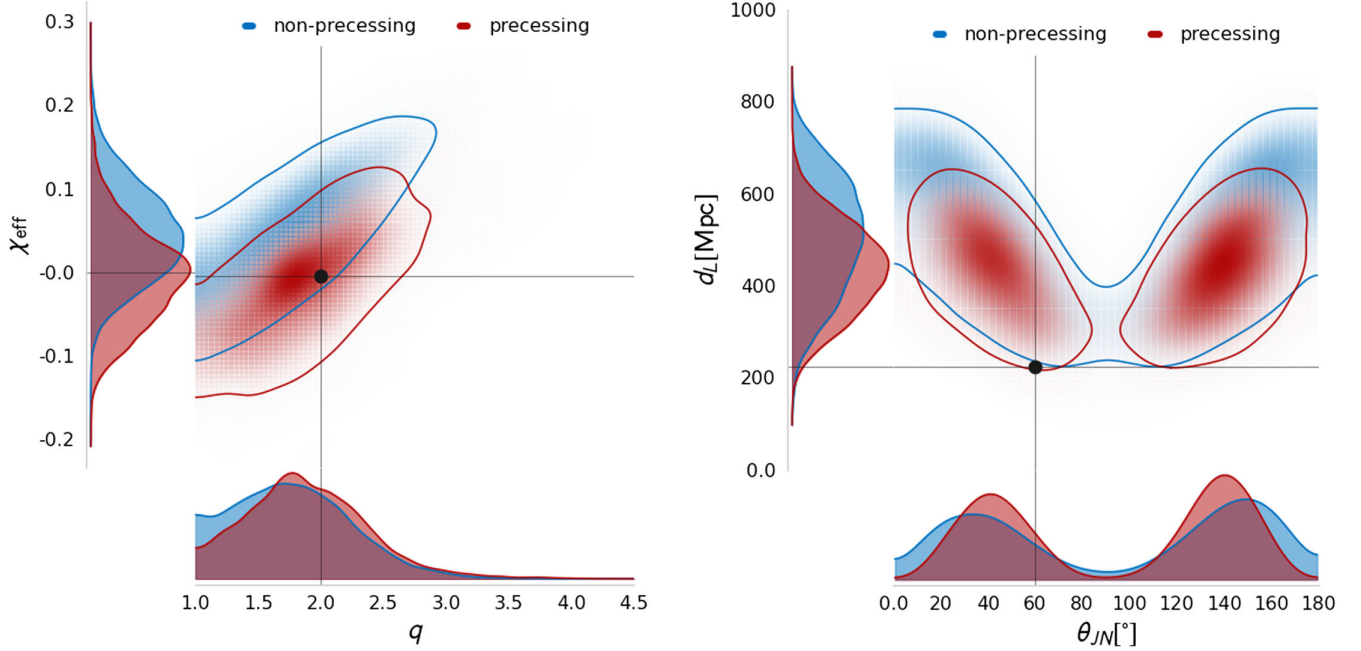
FIG. 4. 2d contour comparing q–χ_{eff} (left) and distance–inclination (right) degeneracies when precession effects are included. Contours show the 90% confidence interval. Bounded two-dimensional KDEs are used for estimating the joint probability density. The black circle indicates the simulated values.

how these measurements improve with stronger precession in Sec. IV B.

The sky location of the source is not well recovered. The analysis was performed with only the two LIGO detectors, and therefore we expect to recover the source restricted to a ring on the sky, which corresponds to a fixed time delay between the detectors [74,75]. The location along the ring cannot be well constrained and, as expected the inferred location is preferentially associated with sky positions where the detector network is more sensitive. Thus, while the simulated sky location is within the 90% region, it is not at or close to the peak. This impacts the recovery of the distance, with the signal being recovered at larger distances, although the simulated distance remains within the 90% range. In Sec. IV G, we show results from a set of runs with varying sky location, and verify that at sky locations where the network is more sensitive, the distance posterior is more consistent with the simulated value.

Lastly, we turn to measurement of precession. In Fig. 5 we show the recovered distributions for binary orientation, θ_{JN}, precessing spin χ_p, initial precession phase, ϕ_{JL}, and precession SNR, ρ_p. There is a clear correlation between the inferred orientation and χ_p, with binaries that are more inclined having lower values of χ_p. Neither of these quantities are directly observable, it is only the amount of observable precession in the system, encoded by ρ_p, that can be measured. Thus the orientation and spin must combine to give the right amount of power in precession, and we see that this is the case—there is little correlation between the recovered values of ρ_p and the precessing

spin χ_p. The inferred value of the precessing spin χ_p and precession SNR ρ_p are both consistent with the simulated values. Specifically, the signal has $\chi_p = 0.4$ and this is consistent with the recovered value, although the posterior distribution is broad, with support over essentially the entire range from 0 to 1. The precession SNR peaks well away from zero, giving clear indication of precession in the system. However, the peak of the distribution occurs at 3.5, while the simulated value is 5.0. We have deliberately chosen an event with significant observable precession. Only a small fraction of the parameter-space volume leads to such significant precession as shown by the *informed* prior on Fig. 5. This is calculated by estimating the allowed values of ρ_p conditioned on the measurements from a nonprecessing analysis. See Sec. VI for further details.

The precession phase, ϕ_{JL}, while not measured with great accuracy, does show two peaks, which are consistent with the simulated value of 45° (0.8 rad). The precession phase can be inferred from the relative phase of the two precessing harmonics using Eq. (7), provided the binary orientation is well measured. There is a clear dependence with the binary orientation: if $\theta_{\text{JN}} < 90°$ then the peak is in ϕ_{JL} at the simulated value and if it is greater then ϕ_{JL} is offset by 180°, to compensate for the change in sign of the $\cos\theta_{JN}$ terms in Eq. (7).

IV. IMPACT OF VARYING PARAMETERS

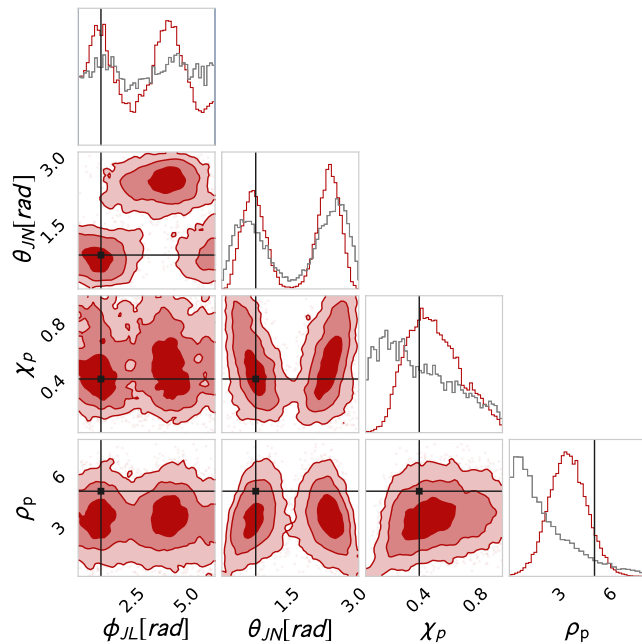We now look at the effect of varying individual parameters one at a time on the recovered posteriors, in particular

FIG. 5. A corner plot showing the recovered values of binary orientation θ_{JN}, precessing spin χ_p, precession phase ϕ_{JL} and precession SNR ρ_p. Shading shows the 1σ, 3σ and 5σ confidence intervals. Black dots show the simulated values. The grey histograms show the *informed* prior, see Sec. VI. There is a clear correlation between the binary orientation and inferred precession spin, with signals which are close to face on $(\cos\theta \approx \pm 1)$ having larger values of precessing spin, while those which are more inclined having less precessing spin. The precession SNR only weakly correlated with χ_p.

focusing on the measurement of precession as described by the posterior distributions of ρ_p and χ_p. All subsequent one-dimensional investigations of the parameter space maintain a constant SNR (except for Sec. IV A where the effect of the SNR is investigated). This is achieved by varying the distance to the source.

Primary results presented in this section will be displayed in the form of violin plots. We show the χ_p posterior distribution (left hand side, colored) compared to the global prior (right hand side, white) unless otherwise stated. We show the ρ_p posterior distribution as a single violin. Horizontal grey lines show the 90% symmetric credible interval. Horizontal red lines show the simulated value. A solid black line corresponds to the $\rho_p = 2.1$ threshold. Bounded KDEsare used for estimating the probability density. We use the same 2d contour plots and multi-dimensional corner plots as described in Sec. III C. Plots were generated with the PESummary [71] PYTHON package.

A. SNR

We start with the fiducial run configuration described above and vary the SNR of the simulated signal.

In the strong-signal limit, where the likelihood surface can be well approximated by a multivariate gaussian, it is well known that the accuracy with which parameters can be measured is generally inversely proportional to the SNR [27,30]. However, this is not always the case due to, for example, degeneracies between parameters (see Ref. [76] for a discussion of the limits of this approximation).

Fig. 6 shows that as the SNR of the simulated signal increases, the accuracy and precision of the inferred χ_p posterior distribution improves. As expected the width of the 90% credible interval decreases approximately linearly with increasing SNR. The improvement in the χ_p posterior distribution can be mapped to a linear increase in ρ_p.

When the simulated signal has low SNR ($\rho = 10$), the recovered χ_p posterior distribution resembles the prior, implying that there is no information about precession in the data. For this case, ρ_p matches the expected distribution in the absence of any measurable precession—a χ distribution with 2 degrees of freedom. As the SNR increases ($\rho = 20$–30), the 5th percentile of the ρ_p distribution is comparable or greater than the $\rho_p = 2.1$ threshold. This maps to the χ_p posterior distribution removing all support for near-zero χ_p ($\chi_p \lesssim 0.1$). For larger SNRs ($\rho > 40$), the entire ρ_p distribution is greater than the 2.1 threshold. This implies significant power from precession. For these cases, we remove support for maximal precession $\chi_p \sim 1$.

As expected we find good agreement between ρ_p and a noncentral χ distribution with 2 degrees of freedom and noncentrality equal to the inferred power in the second harmonic (median of the ρ_p distribution).

B. In-plane spin components

We now look at the effect of varying the amount of precession in the system, varying χ_p from 0 to 1 in steps of 0.25. At $\chi_p = 1$ we have maximal spin, all in the plane of the binary. The inferred values of precessing spin and precession SNR are shown in Fig. 7. We observe, as expected, that increasing the in-plane spin leads to an increase in the magnitude of precession effects observable in the system. With zero precessing spin, there is no evidence for precession in the system; the recovered χ_p is consistent with zero.[3] Similarly, there is no support for significant precession SNR, with ρ_p constrained near zero. As χ_p increases, the amount of precession in the system grows and the measurement of χ_p becomes both more accurate and more precise. Figure 7 shows the relationship between ρ_p and χ_p, and a larger value for ρ_p enables a better measurement for χ_p.

---

[3] We do not expect the χ_p posterior to contain $\chi_p = 0$ as there is no prior support there, however the posterior is relatively well constrained at low precession.
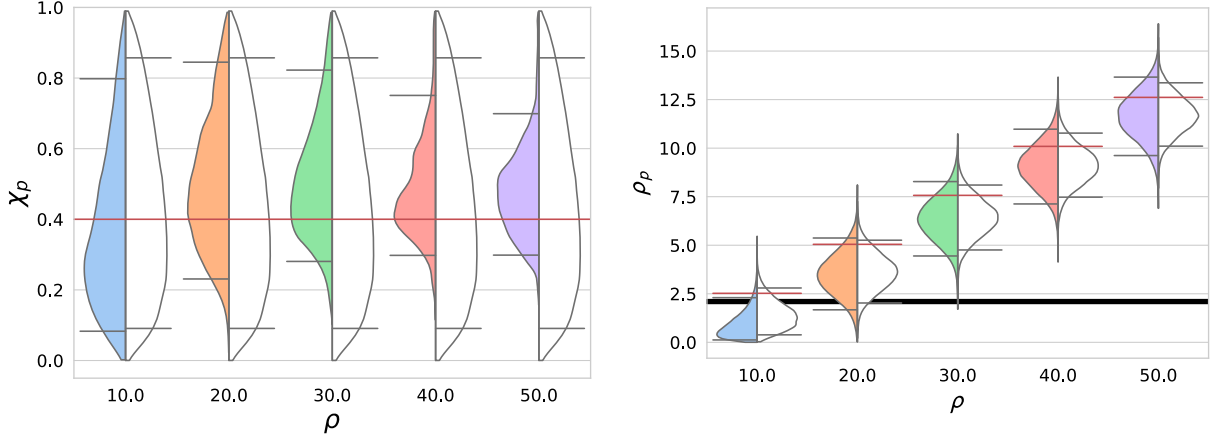
FIG. 6.   Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior (left) and ρ_p compared to a noncentral χ distribution with 2 degrees of freedom and noncentrality equal to the median of the ρ_p distribution (right). Distributions are plotted for varying SNR. Parameters other than the SNR of the signal match the "standard injection" (see Table I).
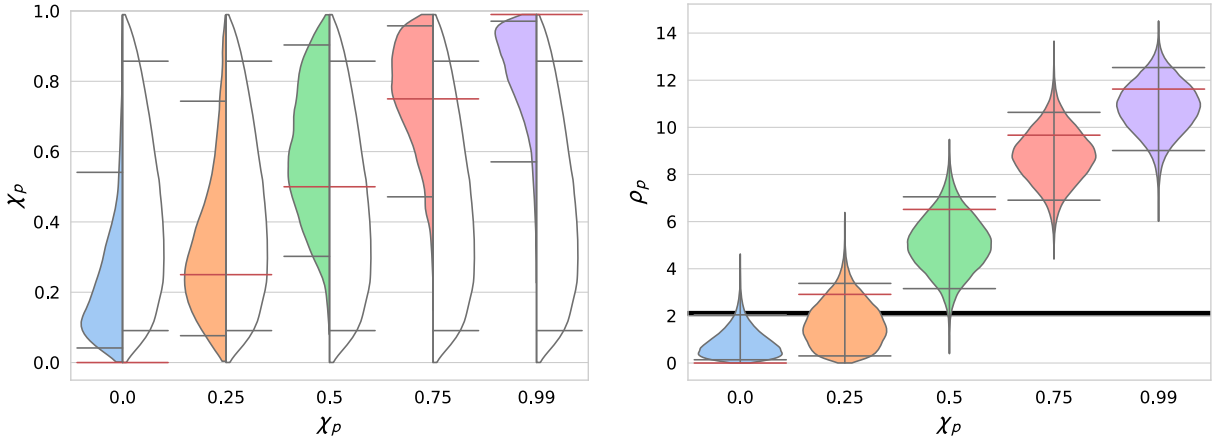

FIG. 7.   Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior (left) and ρ_p (right). Distributions are plotted for varying χ_p. Parameters other than χ_p match the "standard injection" (see Table I).

Fig. 8 shows how the inferred mass ratio–aligned spin and distance–orientation contours change as the magnitude of the in-plane spins change. When there is no observable precession in the system, there is a clear degeneracy in both cases. However, as precession effects become stronger the degeneracy between both pairs of parameters is broken. If ρ_p is small then this can be explained by both a small amount of precession observed at almost any inclination angle, or a large χ_p observed close to face on, as seen in Fig. 5. Since precession effects are not strong enough to provide an accurate measurement of the orientation, the degeneracy between distance and θ_{JN} persists. When ρ_p clearly excludes small values, there is *no support* for close to face-on signals, allowing a more precise measurement of the inclination angle θ_{JN}, breaking the degeneracy with distance.

Stronger precession also allows for improved measurement of the mass ratio. The opening angle β, and consequently the precession parameter $\bar{b}$, increases as the

mass-ratio is increased, as can be seen from Eq. (1). Thus, when strong precession effects are observed, the signal is inconsistent with an equal mass system. In addition, the difference in frequency between the two leading precession harmonics depends upon the mass-ratio [43], and this may also improve our measurement of q. This can also be seen from the precession dynamics, where the precession rate of L around J, $\dot{\alpha}$, depends the mass ratio, and the number of observable precession cycles corresponds to improved accuracy in the measurement of the mass ratio [77].

As χ_p is increased, the peak of the recovered ρ_p distribution is closer to the simulated value. This is likely due to a better measurement of the binary orientation as shown in Fig. 8.

C. Inclination

It is well known that the inclination angle will affect our ability to measure precession, as outlined in the discussion in Sec. II. In particular, from Eq. (7) we see that in the
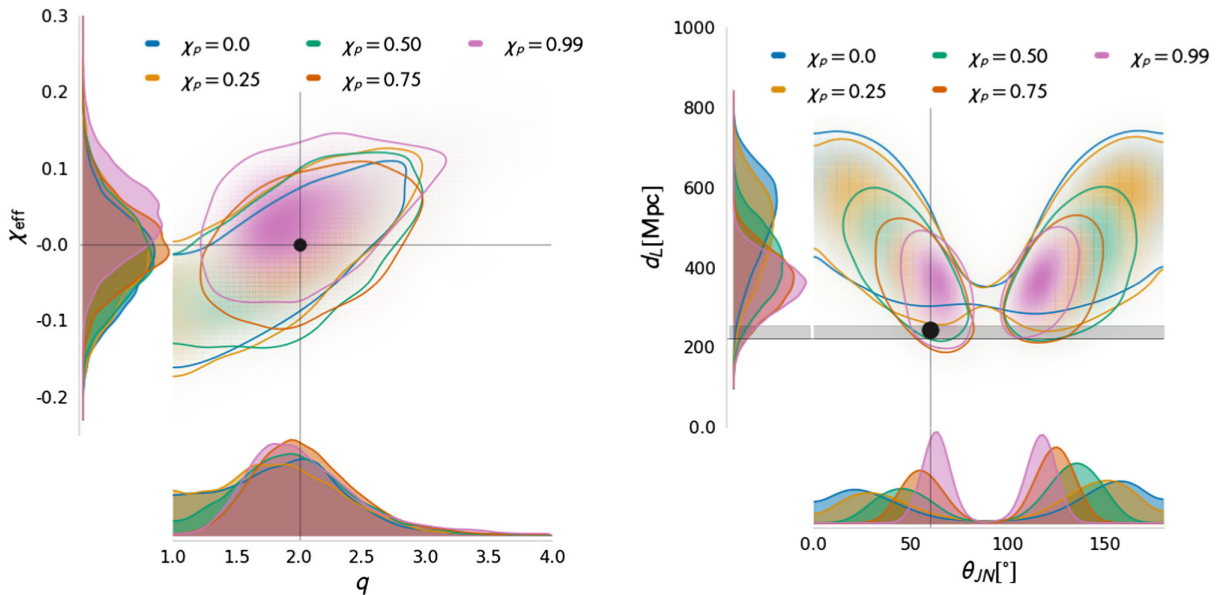
FIG. 8. Two dimensional posteriors for (left) mass ratio and aligned spin, χ_{eff}, (right) binary orientation and distance. Contours show the 90% confidence interval. Bounded two-dimensional KDEs are used for estimating the joint probability density. The black circle with corresponding horizontal and vertical lines indicates the simulated values. For the simulated distance, a solid horizontal band indicates the maximum and minimum simulated values.

two-harmonic approximation the second harmonic vanishes when $\theta_{\text{JN}} = 0°$ or $180°$. In this section we consider the effect of changing the orientation of our standard configuration, which allows us to quantify how it will manifest in realistic LIGO-Virgo signals. A related study has looked at the effect at higher mass ratios [52].

The effect of varying θ_{JN} is shown in Fig. 9. For binaries where the total angular momentum is nearly aligned with the line of sight, precession effects are not observable, as is clear from both the ρ_p and χ_p posteriors. It is not until $\theta_{\text{JN}} \geq 40°$ that we begin to be able to measure precession. Although the accuracy of the measurement clearly improves as we increase θ_{JN}, the uncertainty in the measurement of χ_p

remains large and even at $\theta_{\text{JN}} = 90°$ the posterior is very broad. This can be understood by considering the degeneracies shown in Fig. 5 for the standard signal and in Fig. 10 for the $\theta_{\text{JN}} = 90°$ signal. In both cases, the measured quantity, ρ_p, is relatively well constrained but neither the binary orientation nor χ_p are accurately measured. The observed precession is consistent with both a highly inclined system with lower precessing spin (i.e., low χ_p and large θ_{JN}) or by a less inclined system with higher precessing spin (i.e., high χ_p and small θ_{JN}). Both of these will produce similar observable effects in the waveform.

This allows us to explain the measured posterior for χ_p. At low inclination the posterior is consistent with small
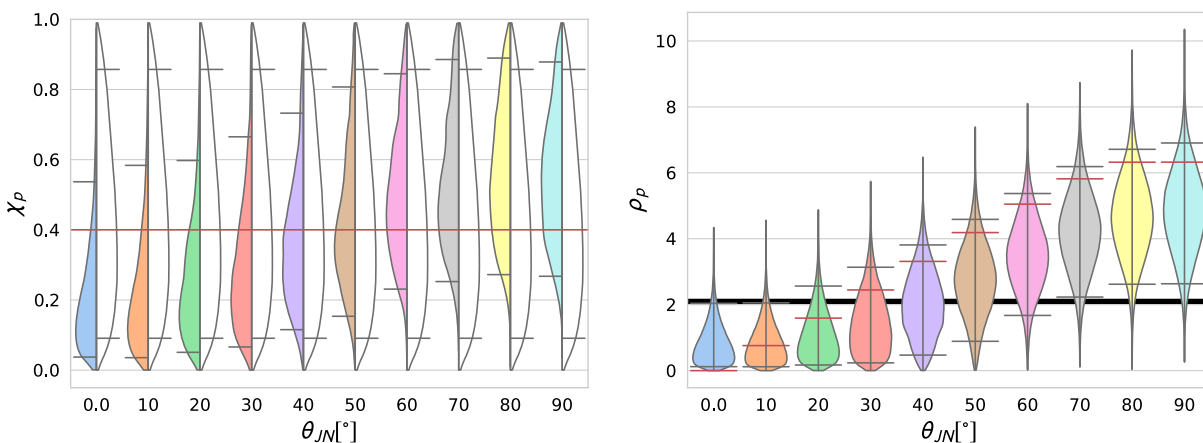

FIG. 9. Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior (left) and ρ_p (right). Distributions are plotted for varying θ_{JN}. Parameters other than θ_{JN} match the "standard injection" (see Table I).
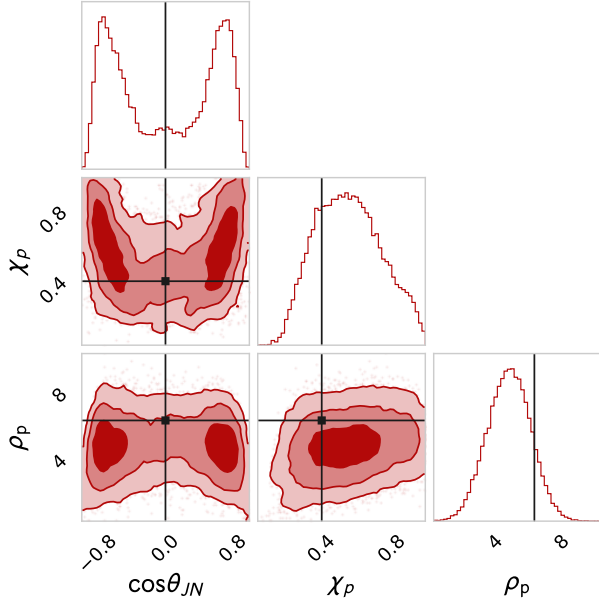
FIG. 10.   A corner plot showing the recovered values of binary orientation θ_{JN}, precessing spin χ_p and precession SNR ρ_p for a system simulated at edge on. Shading shows the 1σ, 3σ and 5σ confidence intervals. Black dots show the simulated values, We see the strong correlation between θ_{JN} and χ_p reflecting the measurement of a certain ρ_p.

values of χ_p. While we are unable to rule out large χ_p, there is limited support as it would require the system to be observed very close to face-on, otherwise precession effects become significant. At large values of θ_{JN}, when precession is clearly observable in the signal, $\chi_p = 0$ is excluded but the distribution remains broad and extends to $\chi_p = 1$.

D. Mass ratio and aligned spin

Figure 11 shows how the inferred precessing spin and precession SNR varies with the mass ratio of the system.
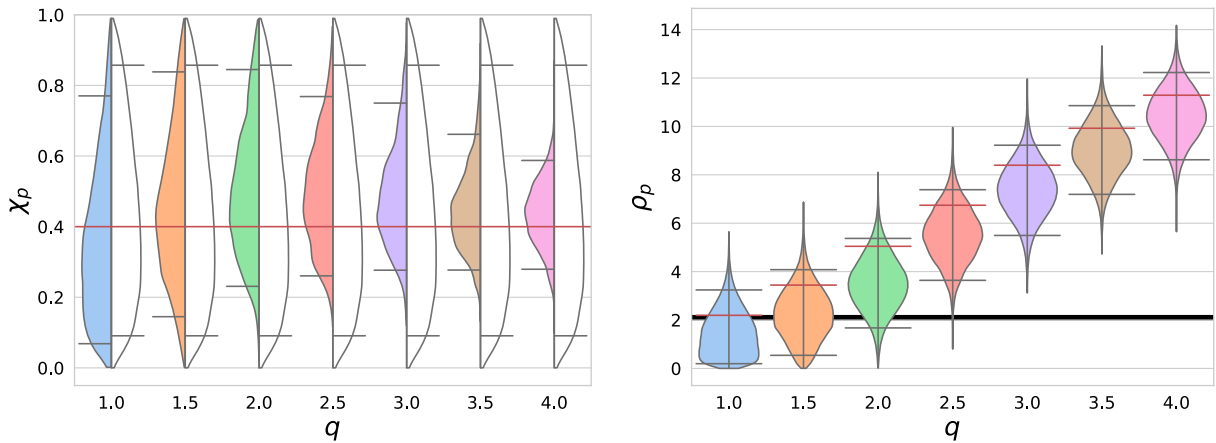
As expected from the general considerations presented in Sec. II, as the mass ratio increases, an in-plane spin on the larger BH leads to a larger opening angle and more significant precession effects. For near equal-mass systems ($q \lesssim 1.5$), the inferred χ_p posterior distribution resembles its prior, and there is not significant power in precession, as shown by the value of ρ_p. As the mass ratio increases, the inferred power in precession also increases and for $q \gtrsim 2.5$, the 90% credible interval of the inferred ρ_p distribution is entirely above $\rho_p = 2.1$. At this stage, precession is clearly identified and $\chi_p \approx 0$ is clearly excluded. In addition, the maximum value of χ_p is also bounded away from maximal.

Fig. 12 shows how varying $\chi_{\rm eff}$ affects our ability to measure precession. A system with a large negative $\chi_{\rm eff}$ results in a larger opening angle compared to an equivalent system with a large positive $\chi_{\rm eff}$. Thus, based upon Eq. (1), we expect the observable impact of precession to be greater for negative values of $\chi_{\rm eff}$ and smaller for positive values. The results are consistent with this expectation, in that the precession SNR decreases with increasing $\chi_{\rm eff}$ and the width of the recovered χ_p distribution increases. However, for the $\chi_{\rm eff} = 0.4$ analysis, we find that the range of χ_p is restricted, with both $\chi_p = 0$ and $\chi_p = 1$ excluded. This is *not* due to the measurement of precession, but is actually due to the measured nonzero aligned-spin component.

A nonzero measurement of $\chi_{\rm eff}$ forces $\chi_p < 1$ as the primary and secondary spin magnitudes must be less than unity. For example, in the $\chi_{\rm eff} = 0.4$ analysis, we measure $\chi_{\rm eff} = 0.38^{+0.07}_{-0.07}$. Under the single spin assumption, this limits $\chi_p < 0.95$. Similarly, since we are using prior distributions that are uniform in spin magnitude and orientation, the observation of a large aligned spin component leads to greater support for a large in-plane spin component. This is shown in Fig. 13, where we plot both the uninformed prior on the primary spin as well as the
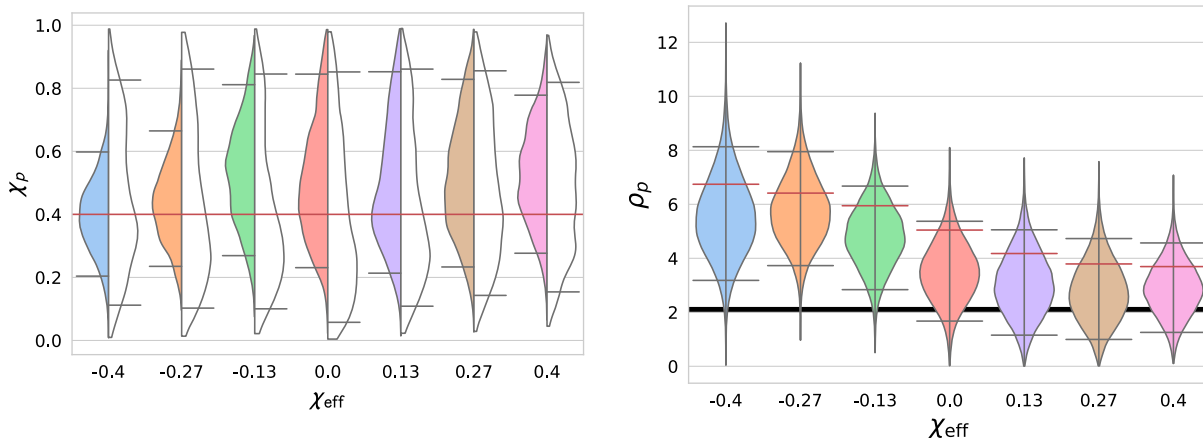

FIG. 11.   Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior (left) and ρ_p (right). Distributions are plotted for varying mass ratio. Parameters other than the mass ratio of the signal match the "standard injection" (see Table I).

FIG. 12.   Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior conditioned on the $\chi_{\rm eff}$ and mass ratio posterior distributions (left) and ρ_p (right). Distributions are plotted for varying $\chi_{\rm eff}$. Parameters other than the $\chi_{\rm eff}$ of the signal match the "standard injection" (see Table I).

prior conditioned on $\chi_{\rm eff} = 0.4$, which removes all support for $\chi_p \approx 0$.

The χ_p measurement for the $\chi_{\rm eff} = 0.27$ and 0.4 analyses are similar to the conditional prior but do restrict the lower χ_p bound beyond prior effects. Although the distribution for ρ_p does extend to zero, it still peaks above $\rho_p = 2.1$ indicating some evidence, although not particularly strong, for precession.

As we vary the mass ratio and aligned spin, the length of the waveform will change. In particular, the aligned spin and high mass ratio configurations produce longer
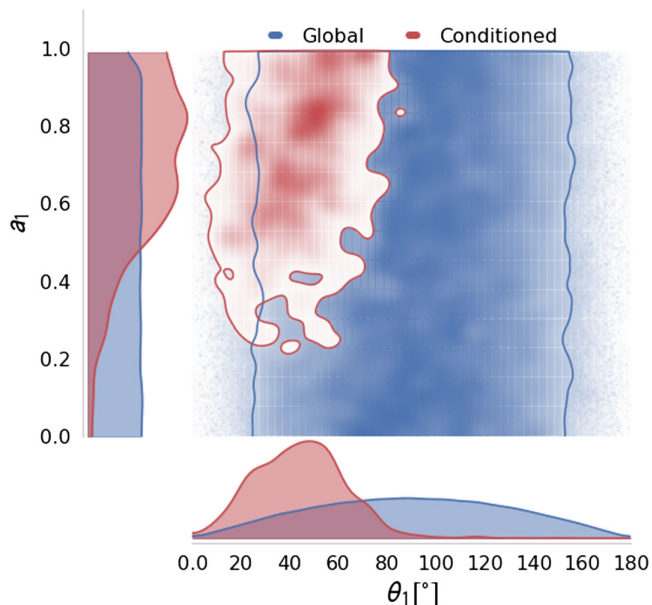

FIG. 13.   2d contours showing the prior 90% credible interval over the primary spin magnitude and spin direction parameter space. Blue shows the global prior and red shows the global prior conditioned on the $\chi_{\rm eff} = 0.4$ mass ratio and $\chi_{\rm eff}$ posterior distributions.

waveforms than those with anti-aligned spins and equal masses [78]. In principle, this will impact the measurability of precession, as longer waveforms allow for a greater number of precession cycles in the detectors' sensitive band. For very short signals, with less than one precession cycle in band, the two leading harmonics are no longer orthogonal (or even approximately so), which make it more challenging to unambiguously identify the second harmonic. This is not an issue for the signals considered here, but does become important when we vary the mass of the binary in Sec. IV E. With a greater number of precession cycles, we will also be able to more accurately measure the precession frequency (the frequency difference between the harmonics), which may improve the measurement of mass ratio [77]. However, it is still the precession SNR that determines the observability of precession. Finally, we note that changing the mass ratio and aligned spin will change the overall amplitude of the waveform. Since our study is performed at a *fixed SNR*, this simply leads to the signals being placed at a larger or smaller distance and therefore does not impact the results presented here.

E. Total mass

We now vary the total mass of the system, keeping all other parameters including mass ratio fixed, in steps of $20\ M_\odot$. As before, we keep the SNR of the system constant at 20, so the higher mass systems are generated at a greater distance. The inferred distributions for χ_p and ρ_p are shown in Fig. 14.

As the total mass of the source increases, the length of the waveform decreases, as does the number of precession cycles, with the number scaling approximately inversely to the total mass (see Eq. (45) of [38]). From the two-harmonic perspective, a small number of precession cycles leads to a large overlap between the harmonics. Specifically, for the $M = 100\ M_\odot$ system the overlap
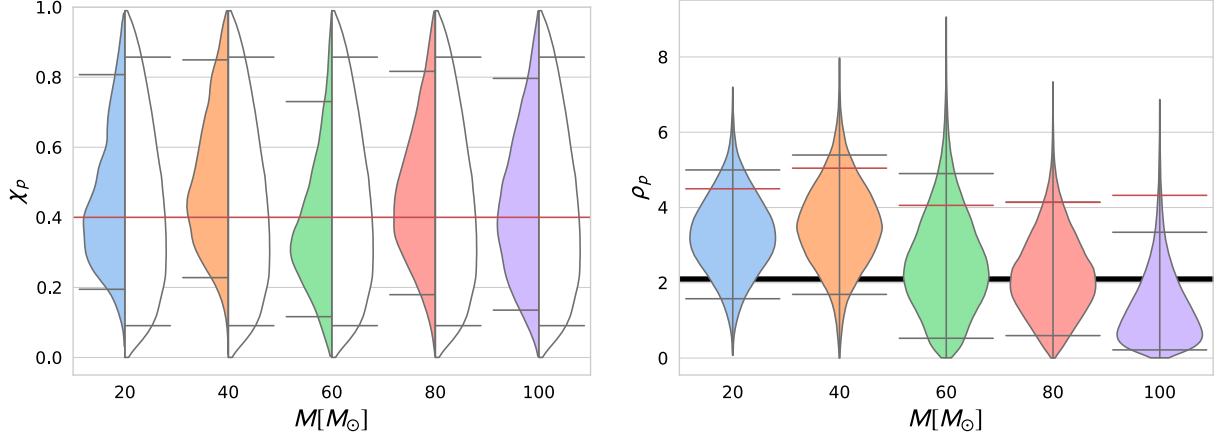
FIG. 14.   Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior (left) and ρ_p (right). Distributions are plotted for varying total mass. Parameters other than the total mass of the signal match the "standard injection" (see Table I).

between the normalized harmonics is $\langle \hat{h}_0 | \hat{h}_1 \rangle = 0.77$, where $\hat{h} = h/|h|$ and the inner product is defined in Eq. (11). At $M = 20\,M_\odot$, the harmonics are close to orthogonal with $\langle \hat{h}_0 | \hat{h}_1 \rangle = 0.15$. The opening angle does not change significantly, with $\bar{b} = 0.14$ at $M = 20\,M_\odot$ and $\bar{b} = 0.21$ at $M = 100\,M_\odot$.

At lower masses, $M \leq 40\,M_\odot$, while the precessing spin is not tightly constrained, it is clearly restricted to be nonzero and the precession SNR has essentially no support for $\rho_p = 0$. For the $60\,M_\odot$ and $80\,M_\odot$ mergers, the precessing spin is still peaked close to the simulated value while ρ_p peaks above 2.1 showing evidence for observable precession, although both ρ_p and χ_p distributions do extend to zero.

For the high-mass system, $M = 100\,M_\odot$, the χ_p posterior more closely matches the prior and we are unable to exclude $\chi_p = 0$. The inferred ρ_p distribution peaks close to zero, and is consistent with no precession, even though the precession SNR in the simulated signal is similar to the lower mass signals. This is likely due to the breakdown of the two-harmonic approximation for this short signal. In particular, for a high-mass system, the power orthogonal to the leading harmonic will depend sensitively upon the initial precession phase ϕ_{JL}. The fact that the recovered value of ρ_p is inconsistent with the simulated value may be due to this fact: the value of $\phi_{JL} = 45°$ used in the simulation leads to maximal observable precession. Across the full parameter space there are very few configurations with significant precession, so this observation is disfavored by our priors. We explore the prior effects such as this in detail in Sec. VI B.

F. Polarization

The effect of changing the relative sensitivity to the two GW polarizations is clear from Eq. (7). Recalling that

$\bar{b} = 0.11$ and $\theta_{\mathrm{JN}} = 60°$, we can express ζ (the ratio of the amplitudes of the two harmonics) as

$$|\zeta| = 0.15 \left| \frac{F_+ + 2iF_\times}{1F_+ + 0.8iF_\times} \right|.$$

Thus, ζ, and consequently the imprint of precession on the waveform, will be maximized when the detector network is primarily sensitive to the $\times$ polarization and minimized when the network is sensitive to the $+$ polarization. We can investigate this by varying the polarization angle of the simulated signal, in steps of $10°$ from the "standard" value of $40°$. At $\psi = 40°$, the sensitivity to the two polarizations is approximately equal, $|F_\times|/|F_+| = 0.9$. It is largest for $\psi = 20°$ where $|F_\times|/|F_+| = 25$ and smallest for $\psi = 60°$ where $|F_\times|/|F_+| = 0.04$. This leads to a variation in the precession SNR from $\rho_p \approx 3$ to $\rho_p \approx 7$.

In Fig. 15 we show the recovered posteriors for χ_p and ρ_p for a set of runs where the precession is varied. The precession SNR varies in accordance with expectation—it is largest at $\psi = 20°$, where the median of the posterior is at $\rho_p = 6$ and there is no support for nonprecessing systems, and smallest at $60°$ where the posterior extends down to $\rho_p = 0$. The amount of observable precession directly impacts the inferred distribution for ρ_p. For the $\psi = 60°$ signal, the posterior for χ_p is consistent with zero, or small in-plane spins, and large values are excluded. Meanwhile for $\psi = 20°$, $\chi_p < 0.1$ is excluded while extremal in-plane spins are consistent with the observation.

It is well known that precession leaves a stronger imprint upon the $\times$ polarization. However, we are not aware of previous results showing how simply changing the polarization of the system can so dramatically change the observable consequences of precession—from being barely observable when the observed signal is primarily the $+$ polarization to being strongly observed in $\times$. Using the
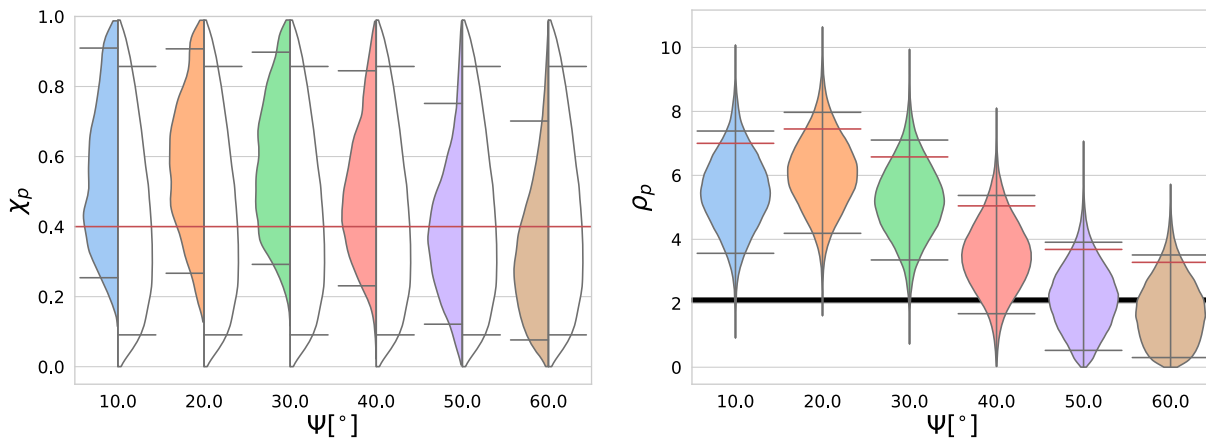
FIG. 15.   Violin plots showing the recovered posterior distributions distributions for χ_p compared to its prior (left) and ρ_p (right). Distributions are plotted for varying ψ_J. Parameters other than ψ_J match the "standard injection" (see Table I).

two-harmonic approximation, we are able to straightfor-wardly predict this effect and then verify it with detailed parameter estimation studies.

G. Sky location

We performed a series of runs where we altered the sky location of the signal, keeping the masses and spins of the components fixed. We also maintained the binary orienta-tion $\theta_{\mathrm{JN}} = 60°$, but varied the distance and polarization of the source to ensure that the SNR remained constant and that the relative contribution of the $+$ and $\times$ polarizations was consistent with the standard run. Furthermore, sky locations were restricted to those for which the relative time of arrival between the Hanford and Livingston detectors remains the same (i.e., we were sampling from the nearly degenerate ring on the sky of constant time delays). Details of the runs are given in Table II.

Table II shows that the inferred luminosity distance remains approximately constant despite the simulated luminosity distance varying by almost a factor of two. In addition, the recovered ρ_p distribution remains consistent with the "standard" injection. Figure 16 shows that the inferred sky position of the source remains essentially unchanged, and consistent with locations of the detectors'

TABLE II.   Table showing the simulated parameters for the sky location set (see Sec. IV G). All other parameters match the "standard injection" (see Table I). The recovered luminosity distance (far right column) is also shown.

Label	RA/rad	DEC/rad	$\psi/°$	d_L/MPc	ρ_p	d_L/MPc
A	0.31	0.92	320	370	5.02	480^{+130}_{-180}
B	0.80	1.15	345	320	5.09	470^{+140}_{-160}
C	1.31	1.22	10	280	5.11	450^{+150}_{-160}
D	1.88	1.19	40	220	5.05	430^{+160}_{-160}
E	6.11	0.21	40	310	5.09	440^{+150}_{-170}

greatest sensitivity. We note here that for this study we only considered the two detector LIGO network. Including VIRGO would likely have considerably improved the precision of the inferred sky location. We do not expect that this would affect any of the inferred physical param-eters or any of the main conclusions in this work.

V. RELATING ρ_p POSTERIORS TO BAYES FACTORS

An alternative method for identifying evidence for precession can be calculated within the Bayesian frame-work. We can calculate the Bayes factor, $\mathcal{B}$, by comparing the marginalized likelihoods [see Eq. (12)] from two competing hypotheses (A, B) [79],

$$\ln \mathcal{B} = \ln p(d_A) - \ln p(d_B). \tag{15}$$

Bayes factors have thus far been the gold standard for identifying evidence for precession within the GW com-munity and have been used extensively in previous works, see e.g., Ref. [52].

In the same way that Bayes factors can be used to quantify evidence for precession, it is also possible to quantify the significance of a GW signal by calculating the Bayes factor for signal verses noise [80]. It has been shown that the log Bayes factor for signal versus noise scales approximately with ρ^2 [81]. Here, we investigate the relationship between the Bayes factor in favor of precession and the precession SNR ρ_p. Both of these quantities have been used together in recent works when assessing the evidence for observable precession [4,5,52]

For a subset of the runs described in Sec. IV C, we reran the analysis using the aligned-spin waveform model IMRPhenomD. Bayes factors in favor of precession could then be calculated and compared to the derived ρ_p posterior distributions.
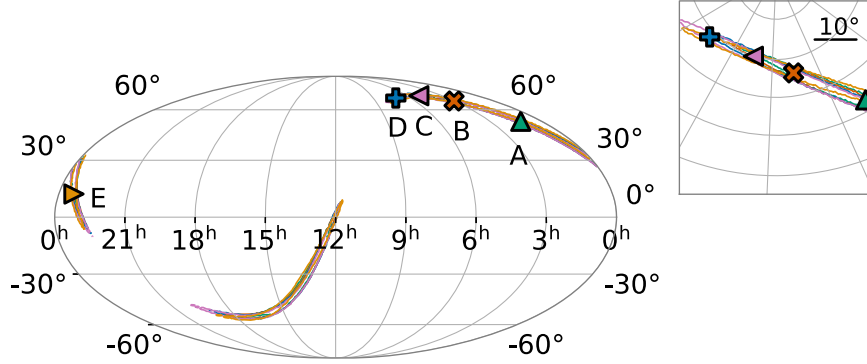
FIG. 16.   Skymap showing the different simulated sky positions, see Table II. The solid lines show the 90% credible intervals and the markers show the simulated sky position. Their respective colors matches their corresponding credible intervals. We vary the distance and polarization of the source to ensure that the SNR remains consistent with the standard injection in Table I.

Fig. 17 shows an approximately linear relationship between the log Bayes factor (ln BF) and the square of the precession SNR (ρ_p^2). This is expected given that the likelihoods recovered from the precessing waveform model will be larger than the likelihoods recovered from the aligned-spin waveform model by a factor of $\exp(\rho_p^2/2)$.

The commonly used heuristic when assessing the strength of evidence using Bayes factors is that $1 \leq \ln BF \geq 3$ is marginal evidence and $\ln BF > 3$ is strong evidence in favor of a hypothesis. From the plots above we conclude that if 90% (50%) of the ρ_p posterior distribution is above the $\rho_p = 2.1$ threshold, this corresponds to a $\ln BF \approx 3.5$ ($\ln BF \approx 0.8$) and is therefore very strong (marginal) evidence for precession. The posterior distribution on ρ_p can therefore be approximately mapped to the

commonly used ln BF. Assessing the strength of evidence for precession using ρ_p would also reduce the need for additional parameter estimation runs using nonprecessing models, which are *necessary* to compute the Bayes factor. This reduction in computational cost will not be significant for a single event, but for population analyses and large scale PE studies this alternative metric could be extremely useful.

VI. PREDICTING THE PRECESSION SNR POSTERIOR

For the majority of simulations presented in this paper, the distribution for the precession SNR, ρ_p, has been peaked significantly below the simulated value, although in nearly every case the simulated value does lie within the 90% confidence region. While the naive expectation is that the recovered posterior will peak at the simulated value, for complex parameter recovery where there are dependencies and degeneracies between the different parameters, this is often not the case. We have already seen that the distance is typically overestimated in the simulations we have performed—this is a well-known effect and arises for two reasons, first that the network is less sensitive to sources from the chosen sky location than from other locations consistent with the observed signal (as discussed in Sec. IV G), and second that the signal was simulated significantly inclined from face-on, yet preferentially recovered close to face-on (as discussed in Sec. IV C). Similarly, it seems likely that the signals we have simulated have more significant precession effects (deliberately, as we wish to understand the observability of precession) than the vast majority of possible sources. Thus, our conjecture is that the likelihood peaks at the simulated value of ρ_p but the posterior distribution will be biased to recover a smaller value owing to the much larger volume of parameter space consistent with low ρ_p. To demonstrate this, we calculate a prior distribution for ρ_p which uses the information gleaned from a nonprecessing analysis to take into consideration
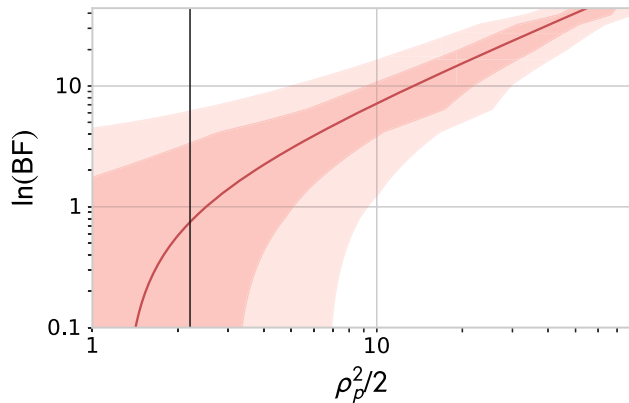

FIG. 17.   Plot comparing the Bayes factor in favor of precession to the inferred ρ_p distribution. Bayes factors were calculated by comparing the evidences for a precessing analysis and a nonprecessing analysis. The uncertainties on the Bayes factors are calculated by taking the 90% confidence interval across multiple LALinferencenest chains. The solid line uses the median of the ρ_p distribution. The shading gives the 1σ and 2σ uncertainties on the ρ_p measurement. The solid black lines shows the $\rho_p = 2.1$ threshold.

the much larger volume of parameter space consistent with low ρ_p. We then show that when multiplying the likelihood by the prior, the *predicted* posterior for ρ_p agrees well with the *inferred* posterior from a fully precessing parameter estimation analysis.

Let us first show that the likelihood peaks at the simulated value of ρ_p. The two-harmonic approximation allows us to factorize the likelihood in Eq. (16) into two terms: a nonprecessing component (dependent on h_0) $\Lambda_{\rm np}(\lambda)$ and precessing component (dependent on h_1) $\Lambda_p(\lambda)$,

$$
\begin{aligned}
p(d|\lambda) &\propto \exp\left(-\frac{1}{2}\langle d - (\mathcal{A}_0(\lambda)h^0(\lambda) + \mathcal{A}_1(\lambda)h^1(\lambda))|d - (\mathcal{A}_0(\lambda)h^0(\lambda) + \mathcal{A}_1(\lambda)h^1(\lambda))\rangle\right) \\
&\propto \exp\left(\langle d|\mathcal{A}_0(\lambda)h^0(\lambda)\rangle - \frac{|\mathcal{A}_0(\lambda)|^2}{2}\langle h^0(\lambda)|h^0(\lambda)\rangle\right) \times \exp\left(\langle d|\mathcal{A}_1(\lambda)h^1(\lambda)\rangle - \frac{|\mathcal{A}_1(\lambda)|^2}{2}\langle h^1(\lambda)|h^1(\lambda)\rangle\right) \\
&\propto \Lambda_{\rm np}(\lambda) \times \Lambda_p(\lambda),
\end{aligned}
\tag{16}
$$

For simplicity we use the approximations that $\langle h^0|h^1\rangle = 0$ and that h^0 is the dominant harmonic, i.e., that the SNR in the h^0 harmonic is larger than in h^1. The calculation proceeds analogously when h^1 is dominant, and can be extended to the general case by replacing h^1 by its projection onto the space orthogonal to h^0.

We can reexpress the precessing contribution to the likelihood Λ_p in terms of the precession SNR using Eq. (8). To do so, we introduce $\hat{\rho}_p$ which is the simulated value of ρ_p, and $\rho_p(\lambda)$ which is the precession SNR for the set of parameters λ. Furthermore, we define the simulated phase [as given in Eq. (6)] of the precession harmonic as $\hat{\phi}_1$ and the phase associated with the parameters λ as $\phi_1(\lambda)$. Following the procedure described in, e.g., Ref. [82], we can rewrite the precession likelihood as

$$
\begin{aligned}
\Lambda_p(\rho_p, \phi_1) &\propto \exp\left(-\frac{1}{2}(\rho_p^2(\lambda) - 2\hat{\rho}_p\rho_p(\lambda)\cos(\hat{\phi}_1 - \phi_1)\right. \\
&\left. + \hat{\rho}_p^2)\right).
\end{aligned}
\tag{17}
$$

In general, we have no prior knowledge of the precession phase, so it is natural to assume a uniform prior on ϕ_1. We may then analytically marginalize $\Lambda_p(\rho_p, \phi_1)$ over ϕ_1 to obtain,

$$
\begin{aligned}
\Lambda_p(\rho_p) &\propto \int_0^{2\pi} \Lambda_p(\rho_p, \phi_1)p(\phi_1)d\phi_1 \\
&\propto I_0(\hat{\rho}_p\rho_p)\exp\left(-\frac{\hat{\rho}_p^2 + \rho_p^2}{2}\right).
\end{aligned}
\tag{18}
$$

We therefore see that the precession likelihood peaks at $\hat{\rho}_p$. We may then calculate the posterior distribution for ρ_p using Bayes' theorem,

$$
p(\rho_p|d) \propto p(\rho_p)\Lambda_p(\rho_p),
\tag{19}
$$

where $p(\rho_p)$ is the prior for the precession SNR.

Previously, in Ref. [43], we obtained a distribution for $p(\rho_p|d)$ by maximizing the likelihood over $\mathcal{A}_1$. This is equivalent to assuming uniform priors for the real and imaginary components of $\mathcal{A}_1$, and leads to a prior $p(\rho_p) \propto \rho_p$. It follows from Eq. (19) that this results in a χ^2 distribution with 2 degrees of freedom. Here, we instead use a prior for ρ_p which is informed by the information obtained from a nonprecessing analysis, we refer to this as the *informed* prior. This *informed* prior better represents our prior knowledge about ρ_p before explicitly accounting for precession in our analysis.

The majority of parameters required to calculate the *informed* prior are already given in the nonprecessing results. The two exceptions are the amplitude of the precessing spin χ_p and the initial precession phase ϕ_{JL}. As discussed in Sec. IV D, we can obtain a prior for χ_p *conditioned* upon the other parameters, specifically the mass ratio and aligned spin $\chi_{\rm eff}$, and this can be used to generate the informed prior on ρ_p. The initial precession phase is unconstrained by the non-precessing parameter recovery, this then allows us to assume it to be uniformly distributed. By calculating the *predicted* posterior distribution for ρ_p based upon a set of nonprecessing samples, we may examine the effect of other measured parameters on the final ρ_p distribution. For example, if the aligned-spin run favors a binary that is close to equal mass and an orientation consistent with a face-on system, then our prior belief will be that the precessing SNR will be low—it is only with unequal masses and systems misaligned with the line of sight that there are significant precession effects in the observed waveform. A prior belief of ρ_p peaking at low values will cause the predicted ρ_p to peak at values lower than the simulated one and consequently so too will the inferred posterior distribution for ρ_p inferred from a full 15-dimensional parameter estimation analysis.

A. Precessing signal

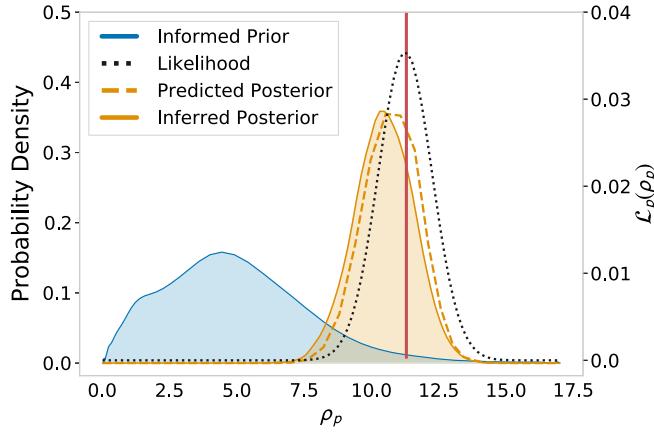We now apply this conjecture to a precessing signal by attempting to *predict* the posterior distributions for ρ_p.

FIG. 18. The predicted distribution for the precession SNR ρ_p (dashed orange) calculated as the product of the precessing contribution to the likelihood (black dotted line) and the informed prior of ρ_p (blue) for the $q = 4$ simulation presented in Sec. IV D. For comparison, we show the inferred ρ_p posterior distribution from the full 15 dimensional parameter estimation analysis (solid orange) and ρ_p for the injection (red line). The informed prior is peaked at low values of ρ_p causing the peak of the posterior to be smaller than the maximum likelihood value.
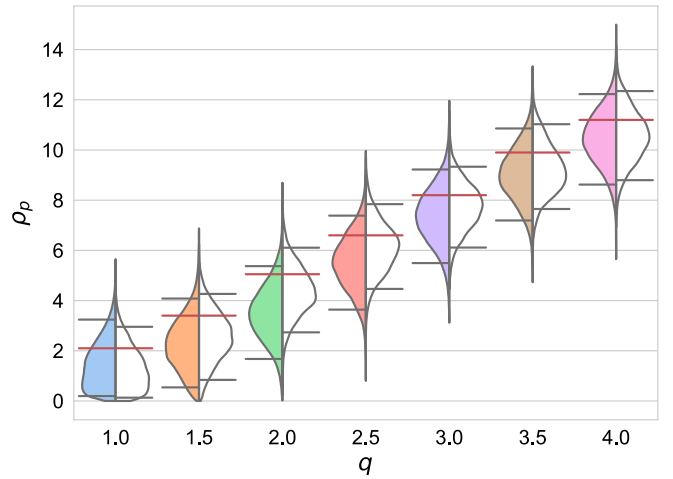

FIG. 19. Violin plot comparing the observed ρ_p distribution (colored) from a precessing analysis, and the predicted distribution (white) based on the aligned-spin results and simulated value of ρ_p for the set of varying mass ratio simulations presented in Sec. IV D. The predicted and observed distributions for precession SNR are in good agreement, even though the ρ_p in the simulated signal (red lines) lies above the peak of either distribution.

This allows us to investigate how much our recovered posterior distributions may differ from the idealized case of a precession likelihood function distributed about the simulated (true) value. In Fig. 18 we show the results of this for the $q = 4$ simulation presented in Sec. IV D. This specific simulation was chosen since this case has the largest ρ_p and corresponds to a simulation where a non-precessing analysis is less justified. It is therefore a good case to show how the combination of the informed prior and the additional likelihood from precession Λ_p correctly estimates the large ρ_p. In Fig. 19, we show how the predicted posterior distribution compares to the inferred distribution over the full range of mass ratio simulations presented in Sec. IV D.

In Fig. 18 we show this predicted distribution, the informed prior, the χ^2 likelihood function and the posterior distribution obtained from a full parameter estimation analysis. By explicitly calculating the informed prior and likelihood terms separately for ρ_p, we can see the effect of the prior on the ρ_p posterior. The prior strongly disfavors large observable precession and therefore pulls the posterior toward *smaller* values than the simulated value i.e., where the likelihood function peaks.

In Fig. 19, we show a comparison between the predicted and measured ρ_p distributions for the set of runs with varying mass ratio presented in Sec. IV D. When we calculate the posterior, explicitly accounting for the parameter space weighting encoded in the informed prior on ρ_p, we find good agreement between the predicted and the inferred ρ_p distributions and note that neither predicted nor inferred are centered around the *true* value for the set of

signals that we have simulated. Of course, if we were to draw signals uniformly from the prior distribution, we would expect to observe the inferred distributions of ρ_p matching with the simulated values.

B. Nonprecessing signal

We now look at the expected posterior distribution for ρ_p when there is no precession in the signal. As explained in Sec. VI, previously a χ^2 distribution with two degrees of freedom was used to model the ρ_p distribution in the absence of any precession (see Ref. [43]). This then led to the natural heuristic that $\rho_p = 2.1$ should be the threshold for observable precession. Using Eq. (19) we can now use a more informative prior on $\tilde{\rho}_p$ and obtain a more accurate estimate of the expected posterior distribution in the absence of precession. We do this by using parameter estimation samples from an aligned-spin model and setting the simulated precession SNR to be 0, this then allows us to account for the effects of priors and different noise realizations.

In Fig. 20 we show the predicted and observed distributions for the precession SNR for a nonprecessing signal. We use a nonprecessing equivalent of the "standard" injection as our simulated signal (i.e., we set $\chi_p = 0$ while ensuring all other parameters match those in Table I). We inject with zero noise and use the `IMRPhenomPv2` model for parameter recovery.

The inferred ρ_p distribution is peaked at lower values that the χ^2 distribution as shown in Fig. 20. However using the prediction from the likelihood [Eq. (16)] and the
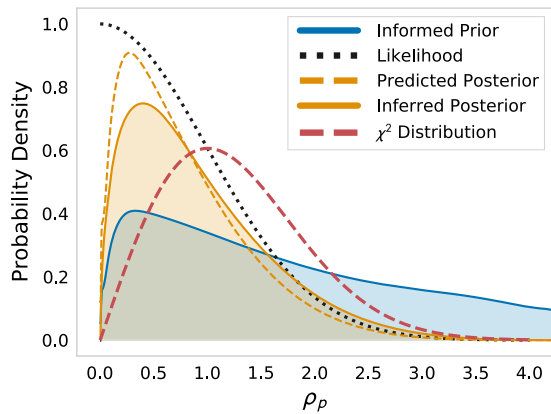
FIG. 20. Distribution of ρ_p in the absence of precession for the "standard injection". The inferred ρ_p distribution using the `IMRPhenomPv2` approximant for recovery is shown by the solid orange line. The dashed orange line shows the predicted distribution using samples collected from an aligned-spin analysis and setting the simulated precession SNR to be 0. We also shows the χ^2 distribution used previously ([43]) as a red dashed line

informed prior we are able to obtain a better estimate of the posterior in the absence of precession. This estimate can be obtained without performing parameter estimation incorporating precession, this therefore allows for a better metric for determining whether or not there is measurable precession in the system.

The distribution for the informed prior on precession SNR will depend upon the details of the signal. In particular, it will be strongly peaked near zero for events that are likely to have small opening angle (equivalently $\bar{b}$), i.e., events that are close to equal mass and have significant spin aligned with the orbital angular momentum, while high mass-ratio events and those with large antialigned spins will lead to greater support for large values of ρ_p. Furthermore, for binaries where the orientation can be well measured, *without* precession information, for example where higher modes are important, those that are close to face-on will lead to predictions of smaller ρ_p while those that are edge-on will give larger values. Given that the majority of signals observed to date are consistent with equal mass binaries, in most cases the prior on ρ_p will tend to be peaked at low values. Consequently, the simple threshold of $\rho_p \gtrsim 2.1$ as evidence for precession, remains appropriate and is likely more stringent than suggested by the simple likelihood calculation.

VII. DISCUSSION

In most candidate astrophysical binary distributions, precession is likely to be first measured in a comparable-mass binary [44]. We have considered a fiducial example of such a possible signal (mass-ratio $q = 2$, SNR $\rho = 20$, and in-plane spin $\chi_p = 0.4$, such that the

precession contribution to the total SNR is $\rho_p = 5$), and performed an extensive parameter-estimation study that has systematically explored the impact on parameter measurements of changes in each of the key source parameters: the SNR, the in-plane spin magnitude, binary inclination, the binary mass ratio and aligned-spin contribution, the binary's total mass, the polarization, and sky location. These examples illustrate well-known features of precession signals [38,44–46,48–51], and quantify their effect on both the measurement of precession, and their impact on the measurement accuracy and precision of other parameters.

We have also verified that ρ_p provides a suitable and intuitive metric for determining whether or not we have measured precession, and shown that there is an approximate mapping between ρ_p and the use of the Bayes factor to assess the evidence of precession. We suggest that given these results, future large scale studies of precession can be made considerably computationally cheaper by computing ρ_p, rather than a full Bayesian analysis.

We note that as ρ_p captures precession by identifying additional power beyond a simple nonprecessing waveform model, it could therefore be effected by phenomena such as eccentricity and higher order multipoles. As BFs simply compare the evidence for two models, one precessing and one nonprecessing, using BFs as the sole metric would also be biased by properties like eccentricity and higher order multipoles.

However, a similar approach to the 2-harmonic decomposition for precessing signals has recently been applied to GWs including the effects of higher harmonics [82]. In future work, we will combine these approaches and explore the measurability of precession in systems with significant evidence for higher harmonics, and the impact of the combination of higher modes and precession upon parameter accuracy. It may also be possible to account for eccentricity through a similar decomposition.

As highlighted in Sec. VI these decompositions provide powerful insights into how the addition of physical phenomena introduce information into the analysis. Here we show that the likelihood can be simply factored into precessing and nonprecessing contributions. This then allows us quantify the extra information that can be gained from a precessing analysis and even predict the recovered ρ_p distribution with or without these effects taken into consideration in the analysis.

The current study does not include higher harmonics, and uses a signal model (`IMRPhenomPv2`) that neglects two-spin precession effects, mode asymmetries that lead to out-of-plane recoil [83], and detailed modelling of precession effects through merger and ringdown. Although these effects are typically small, so is the imprint of precession on the signal, and it would be interesting in future to investigate the impact of these additional features on our results. We also emphasize that, although we consider it to be extremely useful to provide quantitative

examples of the effects of each of the binary parameters, these will necessarily depend on the location in parameter space of our fiducial example. However, having chosen a configuration from among what we expect to be the most likely signals, we hope that these examples will act as a useful guide in interpreting precession measurements when they arise in future gravitational-wave observations.

[1] B. P. Abbott *et al.*, Observation of Gravitational Waves from a Binary Black Hole Merger, Phys. Rev. Lett. **116**, 061102 (2016).
[2] B. P. Abbott *et al.*, 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).
[3] B. P. Abbott *et al.*, GW190425: Observation of a compact binary coalescence with total mass ∼3.4 $M_\odot$, Astrophys. J. Lett. **892**, L3 (2020).
[4] R. Abbott *et al.*, GW190412: Observation of a binary-black-hole coalescence with asymmetric masses, Phys. Rev. D **102**, 043015 (2020).
[5] R. Abbott *et al.*, GW190814: Gravitational waves from the coalescence of a 23 solar mass black hole with a 2.6 solar mass compact object, Astrophys. J. **896**, L44 (2020).
[6] R. Abbott *et al.*, Properties and astrophysical implications of the 150 Msun binary black hole merger GW190521, Astrophys. J. Lett. **900**, L13 (2020).
[7] R. Abbott *et al.*, GW190521: A Binary Black Hole Merger with a Total Mass of 150 $M_\odot$, Phys. Rev. Lett. **125**, 101102 (2020).
[8] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, 2-OGC: Open gravitational-wave catalog of binary mergers from analysis of public advanced LIGO and Virgo data, Astrophys. J. **891**, 123 (2020).
[9] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New binary black hole mergers in the second observing run of Advanced LIGO and Advanced Virgo, Phys. Rev. D **101**, 083030 (2020).
[10] B. Zackay, T. Venumadhav, L. Dai, J. Roulet, and M. Zaldarriaga, Highly spinning and aligned binary black hole merger in the advanced ligo first observing run, Phys. Rev. D **100**, 023007 (2019).
[11] B. Zackay, L. Dai, T. Venumadhav, J. Roulet, and M. Zaldarriaga, Detecting gravitational waves with disparate detector responses: Two new binary black hole mergers, arXiv:1910.09528.
[12] B. F. Schutz, Determining the Hubble constant from gravitational wave observations, Nature (London) **323**, 310 (1986).
[13] M. Soares-Santos *et al.*, First measurement of the Hubble constant from a dark standard siren using the Dark Energy Survey galaxies and the LIGO/Virgo binary-black-hole merger GW170814, Astrophys. J. Lett. **876**, L7 (2019).
[14] B. P. Abbott *et al.*, Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817A, Astrophys. J. **848**, L13 (2017).
[15] B. P. Abbott *et al.*, Gw170817: Measurements of Neutron Star Radii and Equation of State, Phys. Rev. Lett. **121**, 161101 (2018).
[16] B. P. Abbott *et al.*, Gw170817: Implications for the Stochastic Gravitational-Wave Background from Compact Binary Coalescences, Phys. Rev. Lett. **120**, 091101 (2018).
[17] LIGO Scientific Collaboration, A gravitational-wave standard siren measurement of the Hubble constant, Nature (London) **551**, 85 (2017).
[18] B. P. Abbott *et al.*, Tests of general relativity with the binary black hole signals from the LIGO-Virgo catalog gwtc-1, Phys. Rev. D **100**, 104036 (2019).
[19] B. P. Abbott *et al.*, Binary black hole population properties inferred from the first and second observing runs of advanced LIGO and advanced Virgo, Astrophys. J. **882**, L24 (2019).
[20] B. P. Abbott *et al.*, Prospects for observing and localizing gravitational-wave transients with advanced LIGO, advanced Virgo and KAGRA, Living Rev. Relativity **21**, 3 (2018).
[21] J. Aasi *et al.*, Advanced LIGO, Classical Quantum Gravity **32**, 115012 (2015).

[22] F. Acernese *et al.*, Advanced Virgo: A second-generation interferometric gravitational wave detector, Classical Quantum Gravity **32,** 024001 (2015).

[23] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto, Interferometer design of the Kagra gravitational wave detector, Phys. Rev. D **88,** 043007 (2013).

[24] B. S. Sathyaprakash *et al.*, Cosmology and the early universe, arXiv:1903.09260.

[25] E. Bianchi, A. Gupta, H. M. Haggard, and B. S. Sathyaprakash, Quantum gravity and black hole spin in gravitational wave observations: A test of the Bekenstein-Hawking entropy, arXiv:1812.05127.

[26] K. E. S. Ford *et al.*, Multi-messenger astrophysics opportunities with stellar-mass binary black hole mergers, 2019, https://baas.aas.org/pub/2020n3i531/release/1.

[27] C. Cutler and E. E. Flanagan, Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral wave form?, Phys. Rev. D **49,** 2658 (1994).

[28] LIGO Scientific and Virgo Collaborations, Properties of the Binary Black Hole Merger GW150914, Phys. Rev. Lett. **116,** 241102 (2016).

[29] S. A. Usman, J. C. Mills, and S. Fairhurst, Constraining the inclinations of binary mergers from gravitational-wave observations, Astrophys. J. **877,** 82 (2019).

[30] E. Poisson and C. M. Will, Gravitational waves from inspiraling compact binaries: Parameter estimation using second postNewtonian wave forms, Phys. Rev. D **52,** 848 (1995).

[31] E. Baird, S. Fairhurst, M. Hannam, and P. Murphy, Degeneracy between mass and spin in black-hole-binary waveforms, Phys. Rev. D **87,** 024035 (2013).

[32] B. Farr *et al.*, Parameter estimation on gravitational waves from neutron-star binaries with spinning components, Astrophys. J. **825,** 116 (2016).

[33] K. K. Y. Ng, S. Vitale, A. Zimmerman, K. Chatziioannou, D. Gerosa, and C.-J. Haster, Gravitational-wave astrophysics with effective-spin measurements: Asymmetries and selection biases, Phys. Rev. D **98,** 083007 (2018).

[34] P. B. Graff, A. Buonanno, and B. S. Sathyaprakash, Missing Link: Bayesian detection and measurement of intermediate-mass black-hole binaries, Phys. Rev. D **92,** 022002 (2015).

[35] C.-J. Haster, Z. Wang, C. P. L. Berry, S. Stevenson, J. Veitch, and I. Mandel, Inference on gravitational waves from coalescences of stellar-mass compact objects and intermediate-mass black holes, Mon. Not. R. Astron. Soc. **457,** 4499 (2016).

[36] S. Vitale, R. Lynch, V. Raymond, R. Sturani, J. Veitch, and P. Graff, Parameter estimation for heavy binary-black holes with networks of second-generation gravitational-wave detectors, Phys. Rev. D **95,** 064053 (2017).

[37] H. Yu *et al.*, Prospects for Detecting Gravitational Waves at 5 Hz with Ground-Based Detectors, Phys. Rev. Lett. **120,** 141102 (2018).

[38] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries, Phys. Rev. D **49,** 6274 (1994).

[39] L. E. Kidder, Coalescing binary systems of compact objects to postNewtonian 5/2 order. 5. Spin effects, Phys. Rev. D **52,** 821 (1995).

[40] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms, Phys. Rev. Lett. **113,** 151101 (2014).

[41] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism, Phys. Rev. D **89,** 084006 (2014).

[42] R. Abbott *et al.*, Open data from the first and second observing runs of advanced ligo and advanced virgo, SoftwareX **13,** 100658 (2021).

[43] S. Fairhurst, R. Green, C. Hoy, M. Hannam, and A. Muir, Two-harmonic approximation for gravitational waveforms from precessing binaries, Phys. Rev. D **102,** 024055 (2020).

[44] S. Fairhurst, R. Green, M. Hannam, and C. Hoy, When will we observe binary black holes precessing?, Phys. Rev. D **102,** 041302 (2020).

[45] A. Vecchio, Lisa observations of rapidly spinning massive black hole binary systems, Phys. Rev. D **70,** 042001 (2004).

[46] R. N. Lang and S. A. Hughes, Measuring coalescing massive binary black holes with gravitational waves: The impact of spin-induced precession, Phys. Rev. D **74,** 122001 (2006); , Erratum, Phys. Rev. D **77,** 109901 (2008).

[47] K. Chatziioannou, G. Lovelace, M. Boyle, M. Giesler, D. A. Hemberger, R. Katebi, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, Measuring the properties of nearly extremal black holes with gravitational waves, Phys. Rev. D **98,** 044028 (2018).

[48] D. A. Brown, A. Lundgren, and R. O'Shaughnessy, Nonspinning searches for spinning binaries in ground-based detector data: Amplitude and mismatch predictions in the constant precession cone approximation, Phys. Rev. D **86,** 064020 (2012).

[49] S. Vitale, R. Lynch, J. Veitch, V. Raymond, and R. Sturani, Measuring the Spin of Black Holes in Binary Systems Using Gravitational Waves, Phys. Rev. Lett. **112,** 251101 (2014).

[50] B. P. Abbott *et al.*, Effects of waveform model systematics on the interpretation of gw150914, Classical Quantum Gravity **34,** 104002 (2017).

[51] E. Berti, A. Buonanno, and C. M. Will, Estimating spinning binary parameters and testing alternative theories of gravity with LISA, Phys. Rev. D **71,** 084025 (2005).

[52] G. Pratten, P. Schmidt, R. Buscicchio, and L. M. Thomas, Measuring precession in asymmetric compact binaries, Phys. Rev. Research **2,** 043096 (2020).

[53] 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).

[54] D. Gerosa *et al.*, A generalized precession parameter χ_p to interpret gravitational-wave data, Phys. Rev. D **103,** 064067 (2021).

[55] L. M. Thomas, P. Schmidt, and G. Pratten, New effective precession spin for modeling multimodal gravitational waveforms in the strong-field regime, Phys. Rev. D **103,** 083022 (2020).

[56] R. Abbott *et al.*, GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, arXiv:2010.14527.

[57] R. Abbott *et al.*, Population properties of compact objects from the second LIGO-Virgo gravitational-wave transient catalog, arXiv:2010.14533.

[58] LIGO Scientific and Virgo Collaborations, GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence, Phys. Rev. Lett. **116,** 241103 (2016).

[59] W. M. Farr, S. Stevenson, M. C. Miller, I. Mandel, B. Farr, and A. Vecchio, Distinguishing spin-aligned and isotropic black hole populations with gravitational waves, Nature (London) **548,** 426 (2017).

[60] V. Tiwari, S. Fairhurst, and M. Hannam, Constraining black-hole spins with gravitational wave observations, Astrophys. J. **868,** 140 (2018).

[61] A. Buonanno, Y.-B. Chen, Y. Pan, and M. Vallisneri, A quasi-physical family of gravity-wave templates for precessing binaries of spinning compact objects. 2. Application to double-spin precessing binaries, Phys. Rev. D **70,** 104003 (2004); , Erratum, Phys. Rev. D **74,** 029902 (2006).

[62] B. Alejandro, H. Mark, H. Sascha, O. Frank, P. Michael, and S. Patricia, Phenompv2—technical notes for lal implementation, Technical Report, No. LIGO-T1500602, LIGO Project, 2016.

[63] B. P Abbott *et al.*, Gw170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence, Phys. Rev. Lett. **119,** 141101 (2017).

[64] A. Buikema *et al.*, Sensitivity and performance of the advanced LIGO detectors in the third observing run, Phys. Rev. D **102,** 062003 (2020).

[65] B. P. Abbott *et al.*, Prospects for observing and localizing gravitational-wave transients with advanced LIGO, advanced Virgo and KAGRA, Living Rev. Relativity **21,** 3 (2018).

[66] L. S. Finn, Detection, measurement, and gravitational radiation, Phys. Rev. D **46,** 5236 (1992).

[67] J. Veitch *et al.*, Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library, Phys. Rev. D **91,** 042003 (2015).

[68] LIGO Scientific Collaboration, LIGO algorithm library, 2018, https://ascl.net/2012.021.

[69] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, Phys. Rev. D **93,** 044007 (2016).

[70] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal, Phys. Rev. D **93,** 044006 (2016).

[71] C. Hoy and V. Raymond, PESummary: The code agnostic parameter estimation summary page builder, arXiv:2006.06639.

[72] C. Kalaghatgi, M. Hannam, and V. Raymond, Parameter estimation with a spinning multi-mode waveform model: Imrphenomhm, Phys. Rev. D **101,** 103004 (2020).

[73] S. Vitale and H.-Yu. Chen, Measuring the Hubble Constant with Neutron Star Black Hole Mergers, Phys. Rev. Lett. **121,** 021303 (2018).

[74] S. Fairhurst, Triangulation of gravitational wave sources with a network of detectors, New J. Phys. **11,** 123006 (2009); , Erratum, New J. Phys. **13,** 069602 (2011).

[75] L. P. Singer and L. R. Price, Rapid Bayesian position reconstruction for gravitational-wave transients, Phys. Rev. D **93,** 024013 (2016).

[76] M. Vallisneri, Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects, Phys. Rev. D **77,** 042001 (2008).

[77] R. O'Shaughnessy, B. Farr, E. Ochsner, H.-S. Cho, V. Raymond, C. Kim, and C.-H. Lee, Parameter estimation of gravitational waves from precessing black hole-neutron star inspirals with higher harmonics, Phys. Rev. D **89,** 102005 (2014).

[78] M. Campanelli, C. O. Lousto, and Y. Zlochower, Spinning-black-hole binaries: The orbital hang up, Phys. Rev. D **74,** 041501 (2006).

[79] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, England, 2003).

[80] J. Veitch and A. Vecchio, Bayesian coherent analysis of inspiral gravitational wave signals with a detector network, Phys. Rev. D **81,** 062003 (2010).

[81] N. Cornish, L. Sampson, N. Yunes, and F. Pretorius, Gravitational wave tests of general relativity with the parameterized post-Einsteinian framework, Phys. Rev. D **84,** 062003 (2011).

[82] C. Mills and S. Fairhurst, Measuring gravitational-wave higher-order modes, Phys. Rev. D **103,** 024042 (2021).

[83] C. Kalaghatgi and M. Hannam, Investigating the effect of in-plane spin directions for precessing binary black hole systems, Phys. Rev. D **103,** 024024 (2021).

[84] https://www.gw-openscience.org

[85] J. D. Hunter, MATPLOTLIB: A 2D graphics environment, Comput. Sci. Eng. **9,** 90 (2007).

[86] http://www.astropy.org

[87] Astropy Collaboration, Astropy: A community PYTHON package for astronomy, Astron. Astrophys. **558,** A33 (2013).

[88] A. M. Price-Whelan *et al.*, The Astropy project: Building an open-science project and status of the v2.0 core package, Astron. J. **156,** 18 (2018).

[89] https://lscsoft.docs.ligo.org/ligo.skymap

[90] https://corner.readthedocs.io

[91] D. Foreman-Mackey, corner.py: Scatterplot matrices in PYTHON, J. Open Source Softw. **1,** 24 (2016).