

Simple parameter estimation using observable features of gravitational-wave signals

Stephen Fairhurst^{1,*}, Charlie Hoy^{1,2}, Rhys Green¹, Cameron Mills^{1,3,4} and Samantha A. Usman¹

¹*Gravity Exploration Institute, School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, United Kingdom*

²*University of Portsmouth, Portsmouth PO1 3FX, United Kingdom*

³*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany*

⁴*Leibniz Universität Hannover, D-30167 Hannover, Germany*



(Received 5 May 2023; accepted 16 August 2023; published 23 October 2023)

Using simple, intuitive arguments, we discuss the expected accuracy with which astrophysical parameters can be extracted from an observed gravitational wave signal. The observation of a chirplike signal in the data allows for measurement of the component masses and aligned spins, while measurement in three or more detectors enables good localization. The ability to measure additional features in the observed signal—the existence or absence of power in (i) the second gravitational wave polarization, (ii) higher gravitational wave multipoles or (iii) spin-induced orbital precession—provide new information which can be used to significantly improve the accuracy of parameter measurement. We introduce the `simple-pe` algorithm which uses these methods to generate rapid parameter estimation results for binary mergers. We present results from a set of simulations, to illustrate the method, and compare results from `simple-pe` with measurements from full parameter estimation routines. The `simple-pe` routine is able to provide initial parameter estimates in a matter of CPU minutes, which could be used in real-time alerts and also as input to significantly accelerate detailed parameter estimation routines.

DOI: [10.1103/PhysRevD.108.082006](https://doi.org/10.1103/PhysRevD.108.082006)

I. INTRODUCTION

Gravitational wave astronomy has quickly evolved from the first observation in 2015 [1] to now regular observations of black hole binary mergers [2–7]. Future improvements to the existing ground-based gravitational wave detectors [8,9] are expected to increase the frequency of observations further, with mergers being observed daily, or even more frequently, in the coming years [10]. With a large number of observed signals, we can expect that many of them are not unique, and mostly serve to improve the sampling of the underlying astrophysical populations. However, there will be a small number of signals which probe new areas of the parameter space, for example due to having particularly large or small masses and spins (see e.g. [11]); significant mass ratios (see e.g. [11,12]); clear evidence of neutron star structure; eccentricity (see e.g. [13]) and, most tantalizingly, evidence for physics beyond Einstein's relativity. There is, then, a desire to be able to, quickly and easily,

determine which signals are likely to provide interesting results so that energy can be focused on them.

Detailed parameter estimation routines have already been developed [14–30], and are routinely used to recover the parameters of observed signals. Furthermore, increasingly accurate gravitational waveforms have been developed, which incorporate ever more physical effects—higher multipoles, accurate treatment of black hole spins, inclusion of accurate neutron star equation of state, use of numerical relativity results, eccentricity, beyond GR, etc. [31–51]. Thus, we can infer the parameters of the system using ever more sophisticated methods. However, there are two issues. First, as more physical effects are added to the waveforms, the time taken to generate these waveforms, and to sample the expanding parameter space, increases; although there has been recent effort to reduce this computational cost through bespoke optimizations [52–54], and by harnessing machine learning techniques [55,56]. Second, the parameter estimation routines provide estimates and uncertainties, but typically do not identify the features in the waveform that enable the measurement of given parameters with the stated accuracy.

There is a long history of work aimed at understanding, at a more basic level, how parameters can be extracted from the observed gravitational waveform and providing some idea of the expected accuracy of measurements. For

*fairhursts@cardiff.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

example, Refs. [57–61] give early examples of investigations into measurements of masses and the degeneracy of mass and spin. With increased interest in multimessenger astronomy, various methods to understand gravitational wave localization have also been developed [62–67]. The impact of higher gravitational wave multipoles were also examined in Refs. [68,69] and the impact of spin-induced orbital precession in detail in Refs. [70–74]. All along, warnings of using approximation techniques to investigate the full, high-dimensional parameter space in a single analysis have been given [64,66].

In this paper, we synthesize the physical insights mentioned above to provide a hierarchical understanding of parameter recovery from gravitational wave observations. To do so, we begin with the basic information—the observation of a gravitational wave chirp in one detector. In this case, the shape of the waveform can be used to infer some details of the masses and spins. For lower mass binaries, where the merger does not contribute too significantly to the signal power, the chirp mass is measured with good accuracy, while for higher masses the total mass determines the waveform during merger and ringdown. Additional information about the phasing of the system allows for inference of the mass ratio and the components of black hole spin aligned with the orbital angular momentum. A single detector provides essentially no information about sky location, other than what can be inferred probabilistically (the signal is more likely to come from a sky location where the detector is more sensitive). We can infer a maximum distance for the source. However, for most cases, an accurate distance measurement is not possible as it is degenerate with the orientation and sky location.

As more features are observed, it is possible to extract additional astrophysical parameters from the signal. Extra measurements arise from either the observation of the signal in additional detectors, or from the observation of additional waveform features in a detector or network of detectors. If the signal is observed in more than one detector, this enables localization of the source and measurement, in principle, of both gravitational wave polarizations. The relative amplitude and phase of the second polarization provide additional constraints on the distance to and orientation of the binary. Additional waveform features include higher gravitational wave multipoles and spin-induced orbital precession. In both cases, these features can be considered as adding additional components to the gravitational wave signal which are, to a good approximation, orthogonal to the dominant chirp waveform. The significance of higher multipoles is typically more pronounced for systems with more unequal masses. The relative amplitudes of the higher multipoles also depend upon the orientation, with many higher multipoles vanishing for face-on systems. Thus observation of higher multipoles can allow for improved measurement of mass ratio

and orientation of the binary. The observation of precession requires nonzero in-plane spin components and allows inferences about the in-plane spins as well as the orientation of the binary. In this paper, we consider the impact of adding each of the above features, and how it can improve the parameter recovery.

The paper is laid out as follows. In Sec. II we describe the observable features of the waveform, in Sec. III we describe how the observable features of the waveform can be used to infer the system parameters. In Sec. IV we provide details of an implementation with results in Sec. V, including a comparison between our results and those obtained with Bilby [16,18,20], in Sec. VA. In Sec. VI we provide a summary and discussion of future work. In Appendix A we provide additional details of the waveform decomposition and in B provide results for a low signal-to-noise ratio (SNR) signal.

II. THE OBSERVABLE FEATURES IN A GRAVITATIONAL WAVEFORM

The gravitational waveform observed at a detector is given by

$$h_X(t) = \text{Re}[F^X h], \quad (1)$$

where

$$F^X = F_+^X + iF_\times^X \quad (2)$$

and $F_{+,\times}^X$ are the detector response functions for the detector X which depend upon the location of the source relative to the detector, and

$$h := h_+ - ih_\times, \quad (3)$$

where $h_{+,\times}$ are the two polarizations of the gravitational wave, which depend upon the details of the source. In this paper, we restrict attention to black hole binary mergers for which the signal is relatively short-lived, so that we can treat $F_{+,\times}$ as constant over the duration of the signal.¹

The gravitational waveform emitted during a binary merger can naturally be decomposed into a set of spin-weighted spherical harmonics as [76,77]

$$h(t) = \sum_{\ell \geq 2} \sum_{m=-\ell}^{\ell} h_{\ell,m}(t, \vec{\lambda})^{-2} Y_{\ell,m}(\theta, \phi), \quad (4)$$

where $^{-2}Y_{\ell,m}$ is the spin-weighted spherical harmonic of weight -2 , θ and ϕ give the orientation of the observer relative to a coordinate system used to identify the spherical

¹This approximation is appropriate for binary mergers in the advanced detector network, but breaks down for low-mass mergers in next-generation detectors [75].

harmonics, $\vec{\lambda}$ encodes the physical parameters of the system (masses, spins, etc.) and t is the time. Here, and throughout this section, we follow the notation used in [42,78]. The frequency evolution of the harmonics depends upon the orbital frequency of the binary, with the dominant harmonic being the $(\ell, m) = (2, 2)$ having a frequency which is double the orbital frequency during the inspiral phase. Various models for the gravitational waveform emitted during the merger of quasicircular black hole binary mergers have been developed in recent years, see e.g. [32,39,43,47].

For binaries where the spins are misaligned with the orbital angular momentum, neither the orientation of the spins nor the magnitude and orientation of orbital angular momentum remain fixed and both precess around the direction of the total angular momentum, which does remain approximately constant [79]. This orbital precession leads to amplitude and phase modulations in the observed gravitational wave signal, on the precession timescale which is typically slower than the orbital period. These modulations can be interpreted as the beating of different harmonics whose frequencies differ by multiples of the precession frequency [70,80]. Thus, for a binary with misaligned spins, precession will cause a splitting of each (ℓ, m) multipole in Eq. (4) into multiple harmonics whose frequencies differ by the precession frequency.

In many cases, the multipoles for a precessing system are approximated by “twisting up” [40,81] the multipoles of the nonprecessing counterpart based upon the evolution of the orientation of the orbital angular momentum. The direction of the orbital angular momentum $\hat{\mathbf{L}}$, relative to the total angular momentum $\hat{\mathbf{J}}$ is given by two angles; the opening angle β ($\cos \beta = \hat{\mathbf{L}} \cdot \hat{\mathbf{J}}$) and the precession phase, α relative to a fixed orientation, α_o (also denoted ϕ_{JL}). Then, the precession frequency is given by

$$\Omega_p = \dot{\alpha}. \quad (5)$$

To fully describe a coprecessing coordinate system, we require a third Euler angle ϵ , defined via

$$\dot{\epsilon} = \dot{\alpha} \cos \beta, \quad (6)$$

which determines the rate of rotation of the coprecessing frame.

The multipoles for a precessing system are given by

$$h_{\ell,m}(t) = \sum_{n=-\ell}^{\ell} h_{\ell,n}^{\text{NP}}(t) D_{n,m}^{\ell}(\alpha(t), \beta(t), \epsilon(t)), \quad (7)$$

where $h_{\ell,n}^{\text{NP}}$ denotes the waveform for the equivalent nonprecessing system or, equivalently, the waveform observed in a frame that is coprecessing with the binary, $D_{n,m}^{\ell}$ denotes the Wigner D -matrix,

$$D_{n,m}^{\ell}(\alpha, \beta, \epsilon) = e^{i m \alpha} d_{n,m}^{\ell}(-\beta) e^{-i n \epsilon}, \quad (8)$$

and $d_{n,m}^{\ell}$ the Wigner d -matrix given, for example, in [78]. It is straightforward to insert the expression for the precessing multipoles, Eq. (7), into the multipole expansion of the waveform, Eq. (4), to obtain the waveform for a precessing binary. To do so, we first note that the coordinate system (θ, ϕ) is naturally aligned with the total orbital angular momentum $\hat{\mathbf{J}}$. Therefore, if the system is viewed in a direction $\hat{\mathbf{N}}$, the angle $\theta = \theta_{JN}$ is the angle between $\hat{\mathbf{J}}$ and $\hat{\mathbf{N}}$. In addition, the orientation of the x -axis is specified relative to the (initial) precession phase so that $\phi = -\alpha_o$. Therefore,

$$h(t) = \sum_{\ell,m,n} {}^{-2}Y_{\ell,m}(\theta_{JN}, -\alpha_o) D_{n,m}^{\ell}(\alpha, \beta, \epsilon) h_{\ell,n}^{\text{NP}}(t, \vec{\lambda}). \quad (9)$$

Waveform models have been developed to generate accurate representations of the leading multipoles in the gravitational waveform [32,39,43,47]. For example, numerous models provide the (2,2), (3,3), (4,4) multipoles and in addition the (2,1) and (3,2) multipoles.

A. Waveform components

The gravitational waveform, $h(t)$, as given in Eq. (9), is expressed as an infinite sum of waveform components. However, only a small number of these make a significant contribution to the waveform. By identifying the most significant waveform components, and restricting attention to them, we can simplify the waveform with little loss of accuracy. In previous works [69], we have shown that, for nonprecessing waveforms, it is the (2, 2) and (3, 3) multipoles which are most significant across the majority of the parameter space, with the (4, 4) also contributing significantly for systems with high, comparable masses. Similarly in [70,72], we have shown that the two leading-precession harmonics of the (2, 2) mode provide the dominant contribution, although see [82] for examples of highly precessing systems where the third precession harmonic also contributes significantly. Here, we present the expansion in terms of precession harmonics and higher multipoles simultaneously.

In [70], we demonstrated that the leading (2, 2) waveform could be decomposed into five precession harmonics and, furthermore, that these precession harmonics formed a natural hierarchy, with each subsequent mode suppressed by an additional power of

$$b = \tan(\beta/2). \quad (10)$$

For the majority of systems, the opening angle is significantly smaller than 45° , so that $b \lesssim 0.4$ (see Fig. 3 of [70] for details). The opening angle only approaches 45° for systems with unequal masses and a large spin on the more

massive black hole—in this case, the spin can be comparable to the orbital angular momentum. Consequently, for the majority of the binary black hole parameter space, b can be used as a small expansion parameter, and terms of higher order in b can be neglected. In Appendix A, we perform a decomposition for a generic waveform, comprising several multipoles and demonstrate that each multipole can be written in terms of precession harmonics with increasing powers of b .

Let us restrict attention to the most significant waveform component and the leading subdominant contributions. The leading-waveform component, which we denote $h_{22,0}$ arises as the dominant precession contribution to the (2, 2) multipole. The two most significant subleading contributions are $h_{22,1}$, the leading-precession correction to the (2, 2) multipole, and $h_{33,0}$ and the leading contribution the (3, 3) multipole. The waveform can be written as

$$h \approx \frac{d_o}{d_L} \frac{(e^{2i\phi_o} h_{22,0} + \tau^4 e^{-2i\phi_o} h_{22,0}^*)}{(1 + \tau^2)^2} + \frac{d_o}{d_L} \frac{4\tau(e^{2i\phi_o} h_{22,1} - \tau^2 e^{-2i\phi_o} h_{22,1}^*)}{(1 + \tau^2)^2} + \frac{d_o}{d_L} \frac{\tau(e^{3i\phi_o} h_{33,0} + \tau^4 e^{-3i\phi_o} h_{33,0}^*)}{(1 + \tau^2)^3}, \quad (11)$$

where d_L is the luminosity distance, d_o is a fiducial distance, ϕ_o is the reference phase and

$$\tau = \tan(\theta_{JN}/2) \quad (12)$$

describes the orientation of the binary.²

The detailed expression for the waveform in terms of the coprecessing harmonics are given in Eqs. (A11) and (A16). Most notably, the subdominant precession harmonic is reduced in amplitude by a factor b and offset in frequency from the dominant harmonic by the precession frequency Ω_p , so that

$$h_{22,1} = b e^{-i(\alpha - \alpha_o)} h_{22,0}. \quad (13)$$

In addition, the (3, 3) multipole has a frequency which is 1.5 times the (2, 2) multipole.

We are interested in the relative amplitudes of the different waveform components. This is most usefully given in terms of the expected SNR for each component in a gravitational wave detector. Given a detector with a sensitivity given by the power spectral density (PSD) $S(f)$, the expected SNR of a waveform h is

²In Appendix A, we have not explicitly extracted the d_L or ϕ_o factors from the waveform. It is straightforward to do this by redefining the waveform components to be those associated with a binary at distance d_o and coalescence phase $\phi_o = 0$.

$$\rho_h := |h| = \sqrt{\langle h|h \rangle}, \quad (14)$$

where the inner product, $\langle a|b \rangle$, between two time series $a(t)$ and $b(t)$ is defined as

$$\langle a|b \rangle = 4 \operatorname{Re} \int df \frac{\tilde{a}(f) \tilde{b}^*(f)}{S(f)} \quad (15)$$

and $\tilde{a}(f)$, $\tilde{b}(f)$ are the Fourier transforms of $a(t)$ and $b(t)$, respectively.

We denote the amplitude of the $h_{22,0}$ component as σ (consistent with e.g. [83]), and define the amplitudes of the other components *relative* to this. Therefore, we have

$$|h_k| := \sigma \alpha_k, \quad (16)$$

where, by definition $\alpha_{22,0} = 1$,

$$\alpha_{22,1} = \bar{b} \quad \text{and} \quad \alpha_{33,0} = \alpha_{33}. \quad (17)$$

The quantity \bar{b} is the average value of the $b = \tan(\beta/2)$ over the observed waveform, and this gives the relative amplitude of the precession harmonic. The amplitude of the higher multipole is used to define α_{33} . As discussed in detail in [69,70], both of these quantities are generally significantly less than unity. In obtaining the waveform as given in Eq. (11), we have neglected terms which are of order b^2 or $b\alpha_{33}$. In addition, we have neglected terms of order $\alpha_{\ell m}$ for $(\ell, m) \neq (3, 3)$. In some regions of parameter space, the (4, 4) multipole can be more significant than the (3, 3). While we do not consider that case in this paper, the results in the following sections could easily be rederived for an alternative higher multipole.

In Fig. 1 we show the contribution of the different waveform components to the full waveform. It is clear that the (2, 2, 0) waveform is dominant, while the (2, 2, 1) and (3, 3, 0) waveforms have similar amplitudes which are significantly smaller than the (2, 2, 0). Furthermore, these three components provide an excellent approximation to the full waveform. The overall network SNR of the signal is set to 25 which gives SNRs of 14.5, 18.6, and 8.3 in H1, L1, and V1, respectively. This signal has a network SNR of 4.4 in both the (3, 3, 0) and (2, 2, 1) waveform components. Finally, we can calculate the overlap between the full waveform and the approximate waveform,

$$O = \frac{\langle h|h' \rangle}{|h||h'|}. \quad (18)$$

[Note that in the above, we do not maximize the overlap over phase or time but require them to be identical in the full waveform and the approximate one]. The overlap between the full waveform and the (2, 2, 0) component is 0.965, while the overlap with the 3-component waveform is 0.997. Thus, our approximate waveform is only




FIG. 1. The gravitational waveform emitted by a $40 - 10M_{\odot}$ black hole binary, with aligned spin of 0.5 and in-plane spin of 0.4 on the larger black hole and zero spin on the smaller, oriented at an angle $\theta_{JN} = 0.6$. In the plot we show the full waveform and, in addition, the leading contributions to the waveform: the $(2, 2, 0)$ component which is the leading contribution to the $(2, 2)$ multipole, the $(2, 2, 1)$ component which is the leading-order precession correction to the $(2, 2)$ multipole and the $(3, 3, 0)$ component which is the leading contribution to the $(3, 3)$ multipole. The $(2, 2)$ multipole is the most significant and the $(3, 3)$ is the next largest. In addition, we show the expected detector sensitivities for the advanced LIGO and Virgo observatories for the fourth gravitational wave observing run (O4) [85].

distinguishable from the full waveform at a SNR of over 30, which is larger than any observed binary black hole SNR observed in O1-O3 [4] (see e.g. [59,84] and Sec. III for details of the distinguishability criteria).

B. Orthogonalization of waveform components

In Sec. III, we argue that identifying power in the leading-precession or higher-multipole waveform components can play a crucial role in improving parameter estimates, as only a subset of the parameter space will be consistent with the observation of additional features. In addition, nonobservation allows us to exclude regions of parameter space which predict the presence of an observable feature. So far, we have made the simplifying assumption that the different waveform components, specifically the precession harmonics and higher multipole waveforms, are orthogonal to the leading $(2, 2, 0)$ waveform. In many cases, the assumption of orthogonality between waveform components is reasonable, particularly between the $(2, 2, 0)$ waveform and the higher multipoles [69]. However, this does break down in certain regions of parameter space, most notably at higher masses where there are fewer waveform cycles in the observable band of the detector [70].

To obtain SNR in higher multipoles or precession which is orthogonal to the leading waveform component, we must project the waveform component for the mode k onto the

space orthogonal to the leading mode. Since the relative phase between the waveform components depends upon the physical parameters of the system it is a free parameter. Hence the projection of mode k which is orthogonal to the leading component (with an arbitrary phase) is

$$h_{k,\perp} = h_k - \frac{(h_o|h_k)}{|h_o|^2} h_o - \frac{(e^{i\pi/2}h_o|h_k)}{|h_o|^2} e^{i\pi/2} h_o, \quad (19)$$

where the waveform $e^{i\pi/2}h_o$ is simply h_o rotated through 90° . Let us define the complex overlap between the waveform mode k and the dominant mode as

$$o_k = \frac{(h_o|h_k)}{|h_o||h_k|} + i \frac{(e^{i\pi/2}h_o|h_k)}{|h_o||h_k|} \quad (20)$$

then the orthogonal SNR in the mode k is reduced by a factor $\sqrt{1 - |o(k)|^2}$,

$$\rho_{k,\perp} = \rho_k \sqrt{1 - |o_k|^2}. \quad (21)$$

In addition, the total SNR in the signal is

$$\rho^2 = \rho_o^2 + \sum_{k \neq o} \{\rho_k^2 + 2\text{Re}[o_k] \rho_o \rho_k\}. \quad (22)$$

The overlap between the two modes will always lead to a reduction in the perpendicular SNR of the mode k .

However, the total SNR can be increased or decreased depending upon the phase of the overlap between the two signal components. See [70] for a discussion of this in the context of precession.

C. The two gravitational wave polarizations

We have, in Eq. (11), presented an expression for the gravitational waveform emitted by a precessing binary, restricting to the leading-order precession effects as well as the most significant signal multipoles. We observe a hierarchical decomposition of the signal in terms of the precession parameter \bar{b} and higher-multipole amplitude α_{33} . It is tempting to identify the viewing angle, encoded in $\tau = \tan(\theta_{JN}/2)$, as an additional expansion parameter, and separate terms in Eq. (11) which appear with higher powers of τ . However, for a single detector, it is not possible to distinguish the two gravitational wave polarizations encoded in h_k and h_k^* . When we generalize to a network of detectors, τ does become an appropriate expansion parameter.

In Eq. (1) we gave the observed signal in a gravitational wave detector, as a function of both the gravitational wave signal h and the detector's antenna response F . In many cases, including here, it is more natural to work with the left and right *circular* polarizations of the gravitational wave. The detector response to the circular polarizations is F and F^* , respectively, and the observed waveform can be expressed as

$$h_X(t) = \text{Re} \left[\sum_{k \in \text{modes}} h_k [F^X \mathcal{A}_k^R + (F^X)^* \mathcal{A}_k^L] \right], \quad (23)$$

where, for our purposes, the sum over modes is restricted to (2, 2, 0), (2, 2, 1) and (3, 3, 0). The amplitudes of these modes can be read off from Eq. (11) as

$$\begin{aligned} \mathcal{A}_{22,0}^R &= \frac{d_o}{d_L} \frac{e^{2i\phi_o}}{(1 + \tau^2)^2}, & \mathcal{A}_{22,0}^L &= \frac{d_o}{d_L} \frac{\tau^4 e^{2i\phi_o}}{(1 + \tau^2)^2}, \\ \mathcal{A}_{22,1}^R &= \frac{d_o}{d_L} \frac{4\tau e^{2i\phi_o}}{(1 + \tau^2)^2}, & \mathcal{A}_{22,1}^L &= \frac{d_o}{d_L} \frac{-4\tau^3 e^{2i\phi_o}}{(1 + \tau^2)^2}, \\ \mathcal{A}_{33,0}^R &= \frac{d_o}{d_L} \frac{4\tau e^{3i\phi_o}}{(1 + \tau^2)^3}, & \mathcal{A}_{33,0}^L &= \frac{d_o}{d_L} \frac{4\tau^5 e^{3i\phi_o}}{(1 + \tau^2)^3}. \end{aligned} \quad (24)$$

When expressing the waveform amplitudes in terms of left and right circular polarizations, they separate in powers of $\tau = \tan(\theta_{JN}/2)$, so that the waveform is right circularly polarized at $\tau = 0$ (a face-on signal), and $\mathcal{A}_k^L = 0$. It is therefore tempting to introduce τ as an expansion parameter, in a similar way to b and α_{33} . However, the analogy is not exact. The astrophysical population of binaries is expected to be randomly oriented, so that $\cos\theta_{JN}$ is uniformly distributed between -1 and 1 . Thus, there will be systems taking all possible values of τ , including

edge-on systems for which $\tau \approx 1$ and face-off systems for which $\tau \rightarrow \infty$. For a face-away signal, it is natural to change to a coordinate

$$\gamma := \cot\left(\frac{\theta_{JN}}{2}\right) = \tan\left(\frac{\pi - \theta_{JN}}{2}\right) \quad (25)$$

so that $\gamma = 1/\tau$, and $\gamma = 0$ corresponds to a left circularly polarized waveform. Therefore, we can always perform an expansion in the smaller of τ and γ .³ Next, we note that the amplitude of gravitational wave emission in the (2, 2, 0) component is strongest for $\tau \approx 0$ (and $\gamma \approx 0$). Consequently, there will be an observation bias towards sources which are close to face-on or face-away [86], and indeed the majority of signals are expected to be observed with $\theta_{JN} < 60^\circ$, for which $\tau = 1/\sqrt{3}$ —in which case the amplitude of the second circular polarization is reduced by a factor of 9. Thus, it is reasonable to include the orientation angle as our third expansion parameter, and keep only leading terms in τ .

Since a single gravitational wave interferometer is only sensitive to one polarization of the signal, we require a network of detectors⁴ to measure the polarization content of the signal. Therefore, we extend our analysis to a gravitational wave detector network. To do so, we define the multidetector inner product as the sum over individual detector contributions

$$(\mathbf{a}|\mathbf{b}) := \sum_{X \in \text{dets}} (a_X|b_X)_X, \quad (26)$$

where the subscript X on the inner product denotes the fact that the PSD varies between detectors. The expected network SNR is simply the quadrature sum over detectors of the individual detector SNRs. The expected SNR of each polarization and each waveform component in the network of detectors is

$$(\rho_k^{R,L})^2 = \boldsymbol{\sigma}^2 |\mathbf{F}|^2 \alpha_k^2 |\mathcal{A}_k^{R,L}|^2, \quad (27)$$

where $\boldsymbol{\sigma}^2$ and $|\mathbf{F}|^2$ denote vector dot products over the space of detectors. Thus, the network SNR depends upon the detector sensitivities, σ , the network response \mathbf{F} , the relative significance of the waveform component, α_k , and its overall amplitude, $\mathcal{A}_k^{R,L}$.

To measure the left circular polarization, we are interested in the power orthogonal to the right circular polarization. Therefore, we need to calculate the overlap between the two circular polarizations. To do so, it is convenient to

³In what follows, we will restrict to $\tau \leq 1$, with the understanding that the calculation can be easily repeated for $\tau > 1$ by switching to the variable γ .

⁴Alternatively a single gravitational wave observatory comprising multiple interferometers, such as the Einstein Telescope, can be used to infer the polarization content.

introduce the concept of the *dominant polarization* frame [87,88]. The detector response functions \mathbf{F} depend upon the unknown polarization angle of the source. In many cases, it is more convenient to fix a preferred polarization frame when considering the network response and then include the polarization angle ψ in the description of the waveform. To this end, we introduce the weighted network response

$$\mathbf{w} = \boldsymbol{\sigma} \mathbf{F} e^{-2i\psi}, \quad (28)$$

where ψ is the polarization angle. Then \mathbf{w} is simply a vector describing the sensitivity of each detector to the gravitational wave signal, as encoded by the product of the detector's sensitivity, $\boldsymbol{\sigma}$ and the response to the gravitational wave, \mathbf{F} . We fix the polarization angle by working in the dominant polarization frame for which the network is maximally sensitive to the $+$ polarization and, consequently, minimally sensitive to the \times polarization [87,88]. In the dominant polarization frame, \mathbf{w} satisfies

$$\mathbf{w}_+ \cdot \mathbf{w}_\times = 0 \Leftrightarrow \mathbf{w} \cdot \mathbf{w} = \mathbf{w}^* \cdot \mathbf{w}^*. \quad (29)$$

We characterize the network by its overall sensitivity to the dominant polarization, which is given by $|\mathbf{w}_+|$. The sensitivity to the second polarization is $|\mathbf{w}_\times|$. Following [89], we define the relative sensitivity between to $+$ and \times linear polarizations which is given by

$$\alpha_{\text{net}} = \frac{|\mathbf{w}_\times|}{|\mathbf{w}_+|}. \quad (30)$$

In many cases, the detector network is significantly more sensitive to a single polarization of gravitational waves. For example [87,89], the typical sensitivity to the second polarization for the advanced LIGO-Virgo network is 0.3. As additional detectors are added to the network, both the overall sky coverage and sensitivity to the second polarization improves.

In general, the two circular polarizations are not orthogonal. The complex overlap between the left and right circular polarizations is

$$\begin{aligned} o_L &= \frac{(\mathbf{F}\mathbf{h}_k | \mathbf{F}^* \mathbf{h}_k)}{|\mathbf{F}\mathbf{h}_k|^2} + \frac{i(\mathbf{F}\mathbf{h}_k e^{i\pi/2} | \mathbf{F}^* \mathbf{h}_k)}{|\mathbf{F}\mathbf{h}_k|^2} \\ &= \left[\frac{1 - \alpha_{\text{net}}^2}{1 + \alpha_{\text{net}}^2} \right] (\cos 4\psi + i \sin 4\psi), \end{aligned} \quad (31)$$

where, to obtain the result, we have used the form of \mathbf{w} given in Eq. (28) as well as the definition of α_{net} from Eq. (30). For a single detector or network sensitive to only one polarization ($\alpha_{\text{net}} = 0$), the two polarizations are completely degenerate, as expected. For a network with equal sensitivity to the two polarizations ($\alpha_{\text{net}} = 1$), the two circular polarizations are orthogonal. For a typical signal in the advanced LIGO-Virgo network, $\alpha_{\text{net}} \approx 0.3$ and the

overlap between the two circular polarizations is $o_L \approx 0.9$. This places us in a very different situation than for precession and higher multipoles where the overlap is typically small.

When attempting to identify the presence of the second circular polarization, we must identify the power that is orthogonal to the leading polarization. This is obtained by projecting a left-circular signal onto the space orthogonal to that spanned by right-circular signals, in the same way as we did in Eq. (19) to obtain $h_{k,\perp}$. The fact that there is significant overlap between the polarizations has two major impacts. First, the observable power in the second polarization is significantly reduced,

$$\begin{aligned} \rho_{L,\perp} &= \rho_L \sqrt{1 - |o_L|^2} \\ &= \left[\frac{2\alpha_{\text{net}}}{1 + \alpha_{\text{net}}^2} \right] \rho_L. \end{aligned} \quad (32)$$

Second, the fact that the overlap is large means that the overall power in the signal can vary considerably as the polarization angle changes. Specifically, the network SNR is

$$\begin{aligned} \rho_{\text{net}}^2 &= (\mathbf{h} | \mathbf{h}) = \boldsymbol{\sigma}^2 |\mathbf{F}|^2 \sum_k \alpha_k^2 \left\{ (\mathcal{A}_k^L)^2 + (\mathcal{A}_k^R)^2 \right. \\ &\quad \left. + 2\mathcal{A}_k^L \mathcal{A}_k^R \left[\frac{1 - \alpha_{\text{net}}^2}{1 + \alpha_{\text{net}}^2} \right] \cos 4\psi \right\}. \end{aligned} \quad (33)$$

As before, we have made the approximation that α_k is the same for each detector in the network. This is a reasonable approximation—the overall sensitivity of the detectors is captured by σ . Provided the *shape* of the PSD is similar between detectors, then the relative importance of the different waveform components will also be similar.

D. The observed waveform

We have now identified three expansion parameters which enable us to identify the dominant contribution to the waveform, and the leading subdominant contributions. In particular, we parametrize the precession contribution through $b = \tan(\beta/2)$, the higher multipoles through $\alpha_{\ell m}$, their amplitude relative to the (2,2) mode, and the second polarization through the binary orientation $\tau = \tan(\theta_{JN/2})$.

Keeping only the leading terms and the most significant subleading terms in each of these parameters, we obtain the waveform as

$$\begin{aligned} h &\approx \frac{d_o}{d_L} \left[\frac{e^{2i(\phi_o + \psi)}}{(1 + \tau^2)^2} h_{22,0} + \frac{\tau^4 e^{2i(\phi_o - \psi)}}{(1 + \tau^2)^2} h_{22,0}^* \right. \\ &\quad \left. + \frac{4\tau e^{2i(\phi_o + \psi)}}{(1 + \tau^2)^2} h_{22,1} + \frac{4\tau e^{3i\phi_o + 2i\psi}}{(1 + \tau^2)^3} h_{33,0} \right]. \end{aligned} \quad (34)$$

Here, we have chosen to absorb the (unknown) polarization angle into the expression for the waveform, and subsequently work in a fixed polarization frame. The waveform is comprised of four terms. The first is the dominant contribution—it is the right circularly polarized waveform for the leading contribution of the (2, 2) mode. The other contributions are all subdominant in different ways. The second term is the left circularly polarized contribution. This is downweighted by the fourth power of $\tau = \tan(\theta/2) \leq 1$ and, additionally, the observable power is reduced due to the significant overlap between the left and right circular polarized signals. The third term is the leading-precession correction which is downweighted by the precession amplitude $b = \tan(\beta/2)$ as well as τ . The final term is the most-significant higher-multipole contribution to the waveform, which is down-weighted by the reduced amplitude of the higher multipole, encoded in α_{33} . We note that the formalism is equally applicable to predominantly left circularly polarized signals, for which we simply convert to $\gamma = \cot(\theta_{JN}/2)$ and cases where a different multipole, for example (4, 4), is the second most significant.

The expected SNR in the leading mode is given by

$$\rho_o = \frac{d_o}{d_L} \frac{|\sigma \mathbf{F}|}{(1 + \tau^2)^2}. \quad (35)$$

The expected SNR in the different subdominant waveform components are

$$\rho_L = \rho_o \left[\frac{2\alpha_{\text{net}}\tau^4}{1 + \alpha_{\text{net}}^2} \right], \quad (36)$$

$$\rho_{33} = \rho_o \left[\frac{2\alpha_{33}\tau}{1 + \tau^2} \right], \quad (37)$$

$$\rho_{\text{prec}} = \rho_o [4\bar{b}\tau]. \quad (38)$$

The values of b , α_{33} and α_{net} are given, for different binary parameters and network configurations in [69,70,89], respectively. In the expressions above, we have explicitly included the overlap between the two polarizations, which gives rise to the $2\alpha_{\text{net}}/(1 + \alpha_{\text{net}}^2)$ factor, but neglected the overlaps with the precession and higher multipoles signals as these are typically smaller. Additionally, we have not included the impact of the additional waveform components on the recovered SNR consistent with the leading mode.

III. EXTRACTING ASTROPHYSICAL PARAMETERS FROM AN OBSERVED SIGNAL

Once a gravitational wave signal has been observed, the challenge is to extract the astrophysical parameters of the source. Over the years, numerous methods have been developed to obtain parameter estimates, typically using Bayesian methods and densely sampling the parameter

space [14–27,90]. Here, we take a different approach and attempt to identify the *primary* feature that enables the measurement of one, or a combination of, the astrophysical parameters of a source. By doing so, we build up an intuitive understanding of how the parameters can be extracted from the observed waveform. We consider a quasicircular (noneccentric) binary described by fifteen parameters; the masses, m_1 and m_2 , of the black holes, their spins, denoted by the vectors \mathbf{S}_1 and \mathbf{S}_2 , the orientation of the binary given by the phase ϕ_o , inclination angle θ_{JN} and source polarization ψ , and the location of the system relative to the earth, given by the sky location (α, δ) , distance d_L , and arrival time t_o of the signal.

In Sec. II, we demonstrated that the waveform can be expressed as a dominant component, with three subdominant contributions which are the leading-order corrections. The signal can be decomposed in terms of three expansion parameters: $\tau = \tan \theta_{JN}/2$, where θ_{JN} is the angle between the binary’s total angular momentum and the line of sight; $b = \tan \beta/2$, where β is the opening angle between the orbital and total angular momentum and governs the observability of precession effects; $\alpha_{\ell,m}$, the sensitivity of the network to the leading subdominant multipole, (ℓ, m) , relative to the (2, 2) multipole. By expressing the waveform in this way, we are able to identify the impact that the observability (or otherwise) of these three features has upon the parameter recovery. In some cases, the next-order corrections will be observable but, as we argue later, they are unlikely to dramatically impact the parameter recovery. Each of the features above enables us to break a degeneracy between parameters. Identification of the next-order terms merely allows us to refine the measurements, but does not lead to an ability to measure entirely new features.

Throughout this section, we provide examples using signal shown in Fig. 1; the gravitational waveform emitted during the merger of a $40M_\odot$ — $10M_\odot$ black hole binary with aligned spins of 0.5 on the larger black hole and 0 on the smaller black hole. The system is placed at a distance to provide a total network SNR of 25 in the LIGO-Virgo at projected O4 sensitivity [85]. We vary the distance, viewing angle and in-plane spins of the system to investigate the impact of observability of precession, higher multipoles and the second circular polarization on parameter recovery.

A. Parameter measurement accuracy

Given an observed signal s in a gravitational wave detector, the likelihood ratio for the data to contain a signal h , rather than just Gaussian noise n , is given by

$$\Lambda(\vec{\lambda}, s) = \frac{\exp\left[-\frac{1}{2}(s - h(\vec{\lambda})|s - h(\vec{\lambda}))\right]}{\exp\left[-\frac{1}{2}(s|s)\right]}, \quad (39)$$

where the inner product $(a|b)$ was previously introduced in Eq. (15). For a network of detectors, the likelihood is the

product of likelihoods for individual detectors. In order to calculate the posterior distribution for the parameters $\vec{\lambda}$, we use Bayes formula

$$p(\vec{\lambda}|s) \propto \Lambda(\vec{\lambda}, s)\pi(\vec{\lambda}), \quad (40)$$

where $\pi(\vec{\lambda})$ is the prior distribution on the parameters $\vec{\lambda}$.

Let us return to Eq. (39) and consider the case where the data s is well-approximated by a signal with parameters $\hat{\lambda}$, in the sense that $(s|h(\vec{\lambda})) \approx (h(\hat{\lambda})|h(\vec{\lambda}))$. When considering simulated signals, it is natural to identify $h(\hat{\lambda})$ with the known signal. More generally, as discussed in Sec. IV A, it is straightforward to identify the peak likelihood and identify this as $\hat{\lambda}$. Then, the peak likelihood is related to the expected SNR, defined in Eq. (14), through

$$\Lambda(\hat{\lambda}, s) \approx \exp[\rho_h^2/2]. \quad (41)$$

We can also explore the features of the likelihood as the parameters $\vec{\lambda}$ are varied. Substituting into Eq. (39), we see that the likelihood depends upon the difference between the waveforms

$$\Lambda(\vec{\lambda}|\hat{\lambda}) \propto \exp\left[-\frac{1}{2}|h(\hat{\lambda}) - h(\vec{\lambda})|^2\right] \quad (42)$$

Next, following [59], we relate this to the similarity between waveforms, as characterized by the match, M , defined as

$$M(h_1, h_2) = \max_{dt, d\phi} \frac{(h_1|h_2)}{|h_1||h_2|}, \quad (43)$$

where dt and $d\phi$ denote the time and phase offset between the two waveforms, respectively. In particular, we can reexpress the likelihood as

$$\Lambda(\vec{\lambda}|\hat{\lambda}) \propto \exp\left[-\frac{\rho_h^2}{2}(1 - M^2)\right]. \quad (44)$$

As a final approximation, we assume that the match varies quadratically with the difference in parameters $\delta\vec{\lambda}$. This is true at leading order, but at low SNR and for large-dimensional parameter spaces, this approximation breaks down [66]. Nonetheless, the quadratic approximation can be useful in investigating the properties of the signal. To use it, we construct the waveform metric, g_{ab} , defined through

$$M(\delta\vec{\lambda}) \approx 1 - g_{ab}\delta\lambda^a\delta\lambda^b \quad \text{where } \delta\vec{\lambda} = \hat{\lambda} - \vec{\lambda}. \quad (45)$$

Then, the likelihood is approximated as

$$\Lambda(\vec{\lambda}|\hat{\lambda}) \propto \exp[-\rho_h^2(g_{ab}\delta\lambda^a\delta\lambda^b)]. \quad (46)$$

The eigenvectors and eigenvalues of the matrix g_{ab} provide, respectively, the principal directions and the measurement accuracy in these directions. Specifically, the approximate contour containing a fraction p of the posterior distribution, for a signal with SNR ρ_h is given by

$$g_{ab}\delta\lambda^a\delta\lambda^b = \frac{\chi_k^2(1-p)}{2\rho_h^2}, \quad (47)$$

where k is the dimension of the parameter space under consideration and $\chi_k^2(1-p)$ is the chi-square value with k degrees of freedom for which there is a $(1-p)$ probability of obtaining that value or larger. We will typically be interested in generating 90% contours, and will be working in 2, 3, or 4 dimensions, in which case the thresholds are $g_{ab}\delta\lambda^a\delta\lambda^b \leq 2.3, 3.1, 3.9\rho_h^{-2}$ for $k = 2, 3, 4$ respectively.

B. The chirp waveform

We begin by restricting attention to the dominant component of the waveform, arising from the right circularly polarized,⁵ leading contribution to the (2, 2) harmonic and neglect subdominant contributions from the second polarization, higher multipoles and precession. As is well-known, the amplitude and phase evolution of the waveform can be used to extract the masses and (aligned) spins of the black holes [91,92].

The amplitude and phase evolution of the binary-merger waveform is given, at large separations, by the post-Newtonian expansion [93] of Einstein's equations and, close to and at merger, from numerical simulations of binary systems which are combined to provide full models of the gravitational wave signal [31–51]. In the inspiral stage, the leading-order evolution of the waveform is given by the chirp mass

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} = M\eta^{3/5}, \quad (48)$$

where m_1 and m_2 are the masses of the two components, M is the total mass and η the symmetric mass ratio given by

$$M = m_1 + m_2; \quad \eta = \frac{m_1 m_2}{M^2}. \quad (49)$$

Corrections to the phasing arise at subsequent post-Newtonian orders (powers of v/c), with the first correction at 1PN depending upon the mass ratio, η , and the coefficient at 1.5PN order depending also upon the components of the black hole spins which are aligned with the orbital angular momentum. Consequently, the chirp mass is

⁵As before we assume $\tau < 1$ so that the waveform is preferentially right circular polarized. The calculation is equally applicable to left circular waveforms under the replacement $\tau \rightarrow \gamma$.




FIG. 2. The posterior distribution for chirp mass and mass ratio. The contours show the posterior probability density function (pdf), the dashed blue contour shows the 90% region obtained from the pdf and the solid orange ellipse shows the approximate 90% region. In this, and all following plots, the pdf is plotted so that the value at the peak is unity.

the best-measured mass parameter, while the mass ratio and binary spins are typically less well-constrained.

In Fig. 2, we show the posterior probability distribution for our fiducial source ($m_1 = 40M_\odot$ and $m_2 = 10M_\odot$ with SNR 25) when varying only the masses. The posterior distribution is generated by calculating the match across the mass space and substituting into Eq. (46) to obtain the likelihood. The ellipse represents the approximate 90% confidence interval obtained from calculating the metric over the two-dimensional mass space. The metric provides a good approximation to the likelihood, although the fact that the lower probability contours are curved (rather than elliptical) shows that the simple quadratic approximation is beginning to break down.

In Fig. 3, we show the posterior probability distribution for the components of the black hole spins aligned with the orbital angular momentum. We denote the aligned spin components as

$$\chi_{1z} = \frac{\mathbf{S}_1 \cdot \hat{\mathbf{L}}}{m_1^2} \quad \text{and} \quad \chi_{2z} = \frac{\mathbf{S}_2 \cdot \hat{\mathbf{L}}}{m_2^2}. \quad (50)$$

We keep the other parameters of the signal fixed while allowing the aligned spins to vary. As expected, there is a clear degeneracy between the inferred spins. And, as with the mass space, the metric approximation provides a good description of the likelihood distribution. On the figure, we have also plotted lines of constant effective spin,

$$\chi_{\text{eff}} = \frac{m_1 \chi_{1z} + m_2 \chi_{2z}}{m_1 + m_2}, \quad (51)$$

which is typically used to describe a binary's in-plane spin. Since our fiducial system has aligned spin only on the




FIG. 3. The posterior distribution for the aligned spins, with other parameters kept fixed. The contours show the posterior pdf, the dashed blue contour shows the 90% region obtained from the pdf and the solid orange ellipse shows the approximate 90% region. Lines of constant χ_{eff} , χ_{hu} , and χ_{align} are plotted as gray dotted lines, gray dash-dotted lines and black dashed lines respectively.

larger black hole, χ_{eff} does not accurately describe the spin degeneracy, as can be seen in Fig. 3.

We therefore use an alternative effective spin parameter throughout the remainder of this paper, which accurately describes the spin degeneracy shown in Fig. 3, and attribute this spin to both black holes equally (i.e. $\chi_{1z} = \chi_{2z} = \chi_{\text{align}}$).⁶ We use

$$\chi_{\text{align}} = \frac{m_1^\alpha \chi_{1z} + m_2^\alpha \chi_{2z}}{m_1^\alpha + m_2^\alpha}, \quad (52)$$

where $\alpha = \frac{4}{3}$. This was chosen since it accurately describes the spin-degeneracy for low-mass systems (including the one considered here), it is similar to χ_{hu} which describes the number of orbits before merger [94], as can be seen in Fig. 3, and it has the nice property that it is equal to χ_{1z} when $\chi_{1z} = \chi_{2z}$.

Having restricted to a single aligned spin parameter, we can calculate the posterior distribution across the remaining three-dimensional parameter space of masses and (aligned) spins. Figure 4 shows the likelihood on two-dimensional slices through the parameter space. The likelihood distribution at each point in the \mathcal{M} - η space is obtained both by maximizing the match over the aligned spin values and evaluating the likelihood using the maximized match. The distribution for other pairs of parameters is calculated similarly. In addition, we plot ellipses corresponding to the 90% contours using either the two-dimensional metric (fixing the third parameter) or a three-dimensional

⁶We do not make this restriction on the simulated signals, only on our parameter recovery.




FIG. 4. The posterior distribution for the chirp mass, symmetric mass ratio and effective spin. Each plot shows a two-dimensional slice through the parameter space, where the value of the third parameter is chosen to maximize the value of the likelihood. The blue dashed contour shows the 90% region and the orange ellipse shows the approximate 90% region based on calculating the parameter space metric, as discussed in detail in Sec. III. The black ellipse shows the region based on fixing the third parameter to equal the simulated value.

metric projected into the two-dimensional space under consideration. The three-dimensional metric accurately reproduces the likelihood distribution. The two-dimensional metric significantly underestimates the parameter uncertainties due to correlations between the parameters, particularly mass ratio and spin [40]. For example, in Fig. 2, the symmetric mass ratio is bounded between $\eta \in [0.145, 0.175]$ (mass ratio between 3.5 and 4.5 to 1) while allowing the spin to vary increases the range to $\eta \in [0.125, 0.195]$ (a mass ratio between 3:1 and 6:1).

In all cases, we observe that the quadratic approximation given by the metric provides a good fit to the posterior distribution. In particular, the principal directions of the metric match those of the distribution and the surfaces of constant probability are reasonably well-described by ellipses. However, there are some discrepancies, most notably the shape of contours in the $\mathcal{M} - \chi$ space and the asymmetry of contours in the $\eta - \chi$ space.

The overall amplitude, phase and time of arrival of the signal also carry physical information. The amplitude of the observed gravitational wave scales with the mass of the system and is inversely proportional to the distance. Furthermore, the signal amplitude varies with the orientation of the binary relative to the line of sight, and the binary's location relative to the detector network. These facts are encoded in our expression for the SNR in the leading mode, Eq. (35), which we restate here:

$$\rho_o = \frac{d_o}{d_L} \frac{|\mathbf{F}\boldsymbol{\sigma}|}{(1 + \tau^2)^2},$$

where, as before $|\mathbf{F}\boldsymbol{\sigma}|$ denotes a sum over the detectors in the network. The variation of masses and spins impacts $\boldsymbol{\sigma}$, the location of the source impacts F , while the orientation is encoded in τ . All of these combine to limit the accuracy with which the distance to the source can be measured. The phase of the SNR also provides information about the

signal. In particular, looking at Eq. (34), we see that the phase of the waveform is given by $2(\phi + \psi)$. Thus, the measurement of the SNR phase enables us to determine a combination of coalescence phase and polarization angle.

Finally, we note that the observed gravitational wave signal is redshifted as it travels to the detector. This reduces the frequency of the waveform by an overall factor of $(1 + z)$. For black hole binary systems, the frequency content of the observed waveform also scales with the total mass of the binary. Consequently, when the distance/redshift to the system is not known, we are only able to infer the redshifted mass $\mathcal{M} = \mathcal{M}^{\text{source}}(1 + z)$, and not the source mass. In the remainder of the paper, all results are shown in terms of the redshifted chirp mass \mathcal{M} .

C. Observation in a network of detectors

In a network of detectors, we independently measure SNR, signal phase and time of arrival in each of the detectors. In addition to improving the accuracy of mass and spin measurements by increasing the observed SNR, this also enables measurement of the sky location of the source and the second gravitational wave polarization.

1. Source localization

Localization of a transient gravitational wave source in a network of detectors is primarily achieved through timing: the relative time delays between the observed signal at the different detectors can be inverted to provide a sky region from which the source originates [62,95]. The timing accuracy in a single detector is given by $\sigma_t = (2\pi\rho\sigma_f)^{-1}$ where σ_f is the frequency bandwidth of the system and ρ is the SNR. With two detectors, timing alone can restrict the source to a ring on the sky, although it is often possible to identify a most likely region on the ring based upon the relative amplitude and phase of the signal in the two detectors [96,97]. With three detectors, the source can be

localized by timing to two regions of the sky, one above and the other below the plane formed by the three detectors. Observation in three detectors also affords three measurements of the signal amplitude and phase. In many cases, consistency with a gravitational wave signal comprised of two polarizations enables the identification of a preferred sky region [96,98].

For the purposes of this paper, we are not interested in a detailed discussion of source localization. Nonetheless, uncertainty in the sky location of the source will impact the inference of other parameters. Most notably, an accurate estimate of the detector response \mathbf{F} in Eq. (35) enables an accurate measurement of the distance to the source d_L . Similarly, the sensitivity of the network to the second polarization, encoded in α_{net} determines the expected SNR in the left circular polarization or, inverting the problem, measurement of α_{net} and the SNR in the left circular polarization enables inference of the binary orientation from Eq. (36), as we discuss in detail below.

Figure 5 shows the inferred localization for a simulated signal. The signal has a total SNR of 25 in the LIGO-Virgo network, using the expected sensitivity of the fourth observing run [10]. For the given sky location and detector sensitivity, that translates to a SNR of 14.7 in H1, 18.5 in L1, and 8.3 in V1. Using only timing information, the event is localized to an area of 135 deg^2 . By requiring a consistent amplitude and phase across the detectors, the localization area improved to 80 deg^2 . This localization is poorer than achieved for some high SNR events, such as GW170817 [99] and GW190814 [11] due to the selected sky location, close to the plane of the detectors. For this event, large changes in sky position led to relatively small impact on the time of arrival of the source.

2. Binary orientation and distance

Observation of a signal in a network of detectors enables the measurement of the second gravitational wave polarization. In Sec. II C we obtained expressions for the expected SNR in the left and right circular polarizations, with the ratio between them, which depends upon α_{net} and τ , given in Eq. (36), and repeated below:

$$\rho_L = \rho_o \left[\frac{2\alpha_{\text{net}}\tau^4}{1 + \alpha_{\text{net}}^2} \right].$$

For our example signal, the sensitivity to the second polarization is $\alpha_{\text{net}} = 0.35$, which varies between 0.3 and 0.65 over the localization region (as shown in Fig. 5). Thus, sensitivity to the second polarization is reduced by a factor of $2\alpha_{\text{net}}/(1 + \alpha_{\text{net}}^2) \approx 0.6$ relative to the leading polarization. Then, measurement of ρ_o and ρ_L , coupled with a knowledge of α_{net} provide an estimate of the binary orientation, encoded in τ .

In Fig. 6, we show how the binary orientation can be restricted based upon the measurement of the SNR in the

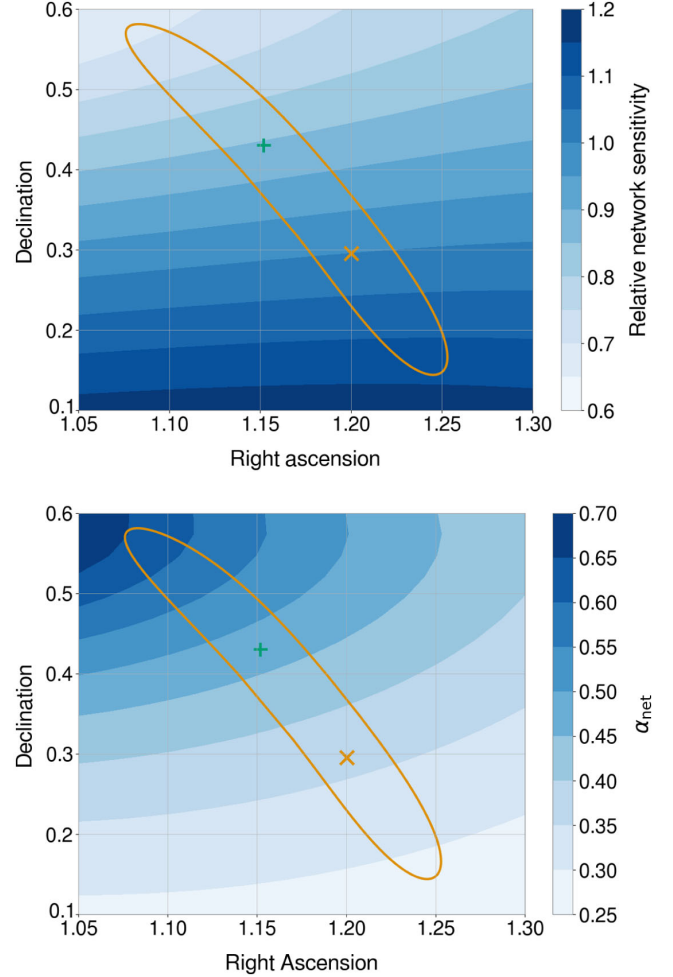


FIG. 5. The 90% source localization ellipse for our simulated signal. The localization is obtained using time of arrival and consistency in amplitude and phase of the gravitational wave signal in the detectors [98]. The orange \times shows location of the simulated signal, the green $+$ shows the “mirror” location (reflection in the plane of the detectors). The top panel shows the localization ellipse overlaid on the network sensitivity while the bottom shows the sensitivity to the second polarization, α_{net} .

two polarizations. As is clear from the equation above, the masses, spins and overall network sensitivity will not impact the estimation of τ . Nonetheless, they do impact the inferred distance. Consequently, for simplicity of presentation, we consider the case where the masses, spins and sky location of the source are fixed. Then the measured SNR in the right circular polarization provides a measurement of $(1 + \cos \theta_{JN})^2 d_L^{-1}$. Similarly, measurement of the SNR in the left circular polarization provides a measurement of $(1 - \cos \theta_{JN})^2 d_L^{-1}$. In the figure, we show two example signals, with binaries inclined at 35° and 67° respectively. For each, we show the region in distance– θ_{JN} space consistent with the observed SNRs. For an expected SNR $\hat{\rho}$ the measured squared SNR will be noncentrally χ^2 distributed with a noncentrality parameter $\hat{\rho}^2$ and two




FIG. 6. Inference of the distance to and orientation of the binary based upon the observed SNR in the two circular polarizations. The blue contours show the probability density from the observed SNR in the right polarization, the green contours show the pdf from SNR in the left polarization and the red regions show the combined pdf. For both examples, the SNR of the source is 25. For the left figure, the binary is inclined at 35° , giving an SNR of 0.15 in the left circular polarization. For the right figure, the binary is inclined at 67° giving an SNR of 3 in the second polarization.

degrees of freedom [69,83]. Thus, for any measured SNR, we can infer region in the distance-inclination space that would give an expected SNR consistent with the observation. The fractional uncertainty in SNR is proportional to ρ^{-1} and, consequently, the observation of the (lower SNR) left circular polarization provides a significantly weaker constraint than the right circular polarization.

For the binary at 35° , there is negligible power observable in the second polarization and the binary is consistent with being face-on. However, the binary orientation cannot be accurately measured and can only be restricted to lie in the range $\theta_{JN} \lesssim 50^\circ$. The system inclined at 67° , has an SNR of 3 in the left circular polarization so that the system is no longer consistent with a circular polarized gravitational wave. The binary orientation can now be restricted to lie in the range $40^\circ \lesssim \theta_{JN} \lesssim 70^\circ$. In these examples, the likelihood is shown as a function of distance and binary orientation. As discussed in [89], it is more appropriate to use a distance prior which is uniform in (comoving-) volume and an orientation distribution flat in $\cos\theta_{JN}$. These distributions will add further weight to large distances and face-on systems making it more difficult to identify inclined sources.

Finally, we note that the phase of the second polarization has a different dependence on the polarization angle ψ , as can be seen in Eq. (34). Thus, observation of both polarizations enables measurement of both the coalescence phase ϕ and the polarization ψ .

D. Higher-order multipoles

As discussed in detail in Sec. II, all gravitational wave signals will contain contributions from multipole moments other than the (2, 2). For the majority of signals, we do not expect to observe these multipoles, as their amplitude will

be too small. However, if it is possible to observe additional harmonics or place limits on the power contained in them, then we can further constrain the range of parameters consistent with an observed signal. The SNR in the (3, 3, 0) waveform component, relative to the (2, 2, 0) component is given in Eq. (37). It scales linearly with $2\tau(1 + \tau^2)^{-1} = \sin\theta_{JN}$ and α_{33} , the relative significance of the (3, 3, 0) harmonic. The value of α_{33} increases with the mass of the binary and is higher for binaries with more unequal components (see Fig. 2 in [69] for details). Therefore, the (3, 3, 0) component is most significant in unequal mass binaries viewed away from face-on (or face-off).

Given an observed SNR in the (3, 3, 0) waveform, we can obtain a region in the distance-inclination plane which is consistent with the observed signal, overlaid on the constraints from the two polarizations of the (2, 2, 0) waveform. For concreteness, we use the same system as before, a binary with masses of $40M_\odot$ and $10M_\odot$, which is inclined at $\theta_{JN} = 35^\circ$. This gives a SNR of 4.4 in the (3, 3, 0) waveform. In Fig. 7, we show how measurement of the SNR in the (3, 3, 0) waveform can be used to restrict the distance and orientation of the binary. Since this is a binary with a significant mass ratio, the (3, 3, 0) waveform plays a much more significant role in determining the orientation of the binary than the second polarization, which has negligible SNR. The observation of the (3, 3, 0) waveform clearly shows that the binary is not face-on, with $\theta_{JN} \gtrsim 15^\circ$. For the events GW190412 and GW190814, it was observation of power in the (3, 3, 0) waveform which enabled measurement of the binary orientation.

In Fig. 8, we show the region in mass-ratio and binary orientation that is consistent with a given observed value of ρ_{33} . The range of θ_{JN} at $q = 4$ corresponds to that in Fig. 7. However, when we allow mass ratio to vary, the allowed




FIG. 7. The restriction of the distance to and inclination of the binary based upon the observed SNR in the two circular polarizations and the (3, 3, 0) harmonic. The blue contours show the probability density from the observed SNR in the right polarization, the green from power orthogonal to the right polarization, the orange from the (3, 3, 0) harmonic and the red show the combined distribution. In calculating the (3, 3, 0) harmonic contours, we have kept the mass ratio fixed to the true value, allowing it to vary will broaden this distribution. The SNR of the source is 25, and it is inclined at an angle of 35° , giving an SNR of 0.15 in the second polarization and 4.3 in the (3, 3, 0) harmonic.

region encompasses close-to-equal-mass systems which are significantly inclined or unequal mass systems which are close-to face on. Since the mass-ratio is already restricted by the observed leading-order waveform, as discussed in Sec. III B, the measurement of the higher-multipole SNR can be used to restrict the binary orientation, as shown on the figure. It is straightforward to add additional multipoles




FIG. 8. The restriction of the mass ratio to and inclination of the binary based upon the observed SNR in the (3, 3, 0) multipole. The orange band gives the posterior based only on measurement of the (3, 3, 0) multipole, the gray band is the region of mass-ratio space consistent with the (2, 2, 0) waveform and the red contours give the region consistent with both measurements.

to this analysis, and the relative power will typically have a different dependence on θ_{JN} . However, additional multipoles will likely refine the measurements but probably not significantly improve them.

The measurement of the phase of the (3, 3, 0) waveform can be used to extract measurements of both the signal's polarization and phase angle. Looking at Eq. (34), we see that the phase of the (3, 3, 0) waveform differs from the (2, 2, 0) by the coalescence phase ϕ . Thus, observation of both waveform components allows for measurement of the phase and, consequently, also the polarization.

E. Precession

Black hole spins which are mis-aligned with the orbit leads to precession of the orbital plane [79] which manifests as amplitude and phase modulations of the signal. As discussed in Sec. II A, precession leads to a splitting of the gravitational wave multipoles. In particular, the (2, 2) multipole is split into five, where the (2, 2, 0) harmonic is the leading term and the (2, 2, 1) harmonic is the first-order precession correction. The SNR in the (2, 2, 1) precession harmonic, relative to the leading (2, 2, 0) harmonic, is given in Eq. (38) as $4\tau\bar{b}$, where \bar{b} is the average value of $b = \tan(\beta/2)$ and β is the opening angle between the total and orbital angular momenta. To leading order during the inspiral phase,

$$\tan \beta = \frac{S_{\perp}}{L + S_{\parallel}}, \quad (53)$$

where S_{\perp} and S_{\parallel} are the perpendicular and parallel components of the spins and L is the orbital angular momentum. The effective precession spin parameter, χ_p , is obtained by averaging the in-plane spins of the system over a precession cycle, so that $S_{\perp} \approx m^2 \chi_p$. Thus, measurement of precession SNR allows us to infer a combination of the precession spin and binary orientation.

In Fig. 9 we show the region of $\chi_p - \theta_{JN}$ parameter space which is consistent with a precession SNR of 4.⁷ In this case, the signal can clearly be identified as precessing and, therefore, both the in-plane spin and binary orientation are bounded away from zero. Nonetheless, there remains a broad range of parameter space consistent with the observation, ranging from maximal in-plane spins for binaries inclined at 15° to edge-on binaries with $\chi_p \approx 0.1$.

In contrast to the left circular polarization and higher multipoles, the observation of precession is unlikely to lead to a significant improvement in the measurement of the binary orientation, distance or phase. The reason for this is

⁷In the figure, the simulated value is slightly offset from the center of the inferred region. This is due to the fact that there is a small amount of power in the left circular polarization in the (2, 2, 1) harmonic which we do not account for when inferring χ_p and θ_{JN} from ρ_p .




FIG. 9. The restriction of the in-plane spin and the binary orientation based upon the observed SNR in precession, ρ_p . The purple contours show the probability density from a system with $\chi_p = 0.4$ inclined at 35° , giving an observed SNR in precession of 4.0. Since the precession spin is otherwise unconstrained, measurement of ρ_p provides minimal restriction to the binary orientation—the inclination is restricted to be above 15° . However, if the orientation is already constrained, e.g. from the observation of higher multipoles, as indicated by the gray band, then the range of permitted values of χ_p can be significantly reduced, as indicated by the red region.

that the amplitude and phase of the precession SNR depend upon the in-plane spins, encoded in the precession spin χ_p and the precession phase α_o . Thus, measurement of the precession SNR enables a measurement of χ_p while measurement of the phase of the precession SNR enables us to extract the precession phase α_o , as can be seen from Eq. (34).⁸ Of course, if the binary orientation has already been restricted through observation of a second gravitational wave polarization or higher multipoles, then this can lead to significant restrictions on χ_p , as is shown in Fig. 9. As an example, the event GW190814 [11] was observed to SNR ≈ 6 in the (3, 3) multipole but minimal SNR in precession, enabling the inference of a very low spin for the primary.

In-plane spins will impact the phase evolution of the (2, 2, 0) waveform component. This can be seen in Eq. (A11), where the waveform acquires an additional phase of $\exp[2i(\alpha - \alpha_o - \epsilon)]$ relative to the nonprecessing signal. Under the approximation that the opening angle β is small and approximately constant, we can simplify Eq. (6) to obtain

$$\epsilon \approx \alpha(1 - \beta^2/2) + \text{const.} \quad (54)$$

⁸See [100] for a discussion of the measured precession SNR in existing gravitational wave events and [101] for a discussion of the measurability of the precession phase.




FIG. 10. The posterior distribution for the precession and aligned spins of the system, keeping other parameters fixed. The contours are generated by calculating the match, the orange contour is generated in the $\chi_p^2 - \chi_{\text{align}}^2$ space, but plotted against χ_p .

Therefore, the phase is approximately quadratic in β . Furthermore, for small values, the opening angle is linearly dependent upon χ_p . Thus, the phasing due to precession will, to leading order, scale with χ_p^2 . To investigate the impact of this, we should reexamine the accuracy with which the masses and aligned spins can be measured when we also allow for precession.

In Fig. 10, we show the posterior distributions for the precession and aligned spin components, keeping the masses fixed. In addition, we construct the metric in the two-dimensional $\chi_{\text{align}} - \chi_p^2$ space. The metric accurately reconstructs the posterior, whereas working in $\chi_{\text{align}} - \chi_p$ coordinates does not accurately capture the degeneracy. From Fig. 10, it is clear that the value of χ_p is essentially undetermined from the phasing of the leading waveform component (the reason that the contour does not extend to $\chi_p = 1$ is that the total spin $\sqrt{\chi_p^2 + \chi_{\text{align}}^2}$ is required to be less than 1). Nonetheless, we must include the degeneracy between χ_p and the other mass and spin parameters when calculating posterior distributions for the parameters.

F. Summary

In the preceding sections, we have laid out how the gravitational wave parameters are encoded into the gravitational wave signal and how, upon observing the signal and various specific features, we are able to extract estimates of the different parameters. Here, we provide a brief summary of the above discussions.

The amplitude and phase evolution of the (leading harmonic of) the gravitational waveform enables us to measure:

- (1) The redshifted chirp mass, $\mathcal{M}(1+z)$ of the signal.
- (2) The symmetric mass ratio η .

- (3) A combination χ_{align} of the aligned spins χ_{1z} and χ_{2z} . The individual spins are typically unmeasurable.
- (4) The time of coalescence of the system, as measured at the detector.

There is significant degeneracy between these parameters, most notably the mass ratio and aligned spins.

When the system is observed in a network of detectors

- (5) The right ascension of the source.

- (6) The declination of the source.

When the system is observed in at least three observatories, we typically localize to a region with an area of a few square degrees. Signals observed in two detectors are only localized to (a fraction of) a ring in the sky with areas of hundreds of square degrees.

Using the detector antenna response, we are able to identify

- (7) The amplitude of the gravitational wave signal, which enables inference of a combination of distance to the source and its orientation, $(1 \pm \cos \theta_{JN})^2/d_L$.

- (8) The phase of the dominant circular gravitational wave polarization ($\phi \pm \psi$).

If we are able to identify either the second polarization or power in higher multipoles, or both, this enables measurement of

- (9) The binary orientation, and consequently a more accurate distance measurement.

- (10) A second phase measurement, which enables the separate inference of the coalescence phase ϕ_o and polarization ψ .

We note that for an aligned-spin system, this comprises all of the parameters when we simplify to a single, effective spin parameter; two masses, one spin parameter, four extrinsic parameters, sky location and time of arrival. Thus, in principle, with a network of detectors, all of the parameters can be measured. Those which are observed with the least accuracy tend to be the second combination of effective spin and mass ratio, and those which require an observation of the second polarization or higher multipoles.

With measurement of power in precession we can measure

- (11) The precession spin χ_p , provided the orientation is already constrained and, if not, then a combination of χ_p and θ_{JN} .

- (12) The precession phase, α_o (also denoted ϕ_{JL}).

IV. SIMPLE-PE IMPLEMENTATION

The intuitive understanding of how the binary parameters are encoded in the observed gravitational wave signal, given in Sec. III F, can be used to develop a simple, computationally cheap parameter-estimation routine. Here, we introduce the `simple-pe` [102,103] algorithm that has been developed for this purpose.

The outline of the method is as follows. First, we identify the peak of the likelihood in the mass and spin space in the network of detectors, maximizing the time of arrival,

amplitude and phase of the signal independently in each detector. The values of masses and spins at the peak are used as central values for those parameters. The arrival times, relative amplitudes and phases of the signal in each observatory are used to obtain an estimate for the sky location of the source. Around this peak, we construct posterior distributions for the masses and spins of the binary, using the expected accuracies and known degeneracies presented in Sec. III B. We matched filter the data at the peak of the likelihood to identify the SNR in the second-polarization, higher-multipole and precession waveforms. Then, based upon the expected SNR in each of these features as a function of the masses, spins and binary orientation, we identify regions of parameter space that are consistent with the observed SNRs. In particular, the SNR in the second polarization can be used to restrict the orientation, both mass ratio and orientation are constrained by the SNR in higher multipoles and in-plane spins restricted using the SNR in the leading-precession correction to the waveform. Finally, we infer the distance distribution based on the masses, network sensitivity and binary orientation.

A. Find the maximum likelihood in mass-spin space

The first step is to identify the peak of the likelihood, or equivalently the maximum SNR, across the mass and spin space. In Sec. III B, we have shown that the masses and aligned spin can be inferred from the phasing of the dominant waveform component. In Sec. II we have argued that the higher multipoles and precession waveform components could contribute a non-negligible amount of SNR to the overall signal. Thus, when finding the peak of the SNR or likelihood, we must consider whether it is necessary to incorporate the power in either of these features.

Figure 11 shows the SNR of the (2, 2, 0), (2, 2, 1), and (3, 3, 0) waveform components when matched-filtered against the same simulated signal which we have considered previously—masses of 40 and $10M_\odot$ with aligned spin components of 0.5 and 0, respectively, and $\chi_p = 0.5$. We compute the SNR for the L1 detector, across a range of masses, keeping the values of χ_{align} and χ_p fixed. The simulated signal has an SNR of 18.5. As expected, from Sec. II D, the SNR in the (2, 2, 0) component is the largest and has a value of almost 18 by itself. Interestingly, though, the peak of the SNR, occurs at masses offset from the simulated values. The offset in the peak SNR is largely caused by precession. For this signal, the overlap between the two precession harmonics (2, 2, 0) and (2, 2, 1) is ≈ 0.2 and, therefore, one obtains a higher SNR in the (2, 2, 0) harmonic at values of the masses where it picks up some of the power in the (2, 2, 1) harmonic. This can be clearly seen from the SNR distribution for the (2, 2, 1) harmonic, which decreases significantly from $\rho_p = 4$ at the simulated values to $\rho_p = 3$ at the (2, 2, 0) peak. The structure of the SNR in




FIG. 11. The SNR distribution for across the mass space for the L1 detector for the modes (2, 2, 0), (2, 2, 1), and (3, 3, 0). The system masses are marked with a \times . The system has a precessing spin χ_p of 0.4. The peak of the matched-filter SNR for each mode is marked with a $+$ (the peak for the (2, 2, 1) component is off the plot). The SNR of the (2, 2, 0) waveform peaks away from the simulated mass values due to the presence of power in precession. The combined SNR of the (2, 2, 0), (2, 2, 1), and (3, 3, 0) components peaks closer to the simulated value. The remaining discrepancy arises from using χ_p and χ_{align} , rather than the full spin vectors, to describe the binary spins.

the (3, 3, 0) waveform is very similar to the leading mode, varying between $\rho_{33} \approx 2.6$ to 2.9 across the region of interest. It will therefore have a limited contribution to the offset of the peak. When combining the SNRs in the three waveform components, we find that the peak is approximately in the correct location. Interestingly, it is not in the exact location, and this arises because we are using χ_{align} and χ_p to describe the signal and these parameters do not perfectly describe the simulated waveform. The simulated signal has spin on only the larger black hole while we assign the same (aligned)-spin value χ_{align} to both components when identifying the signal. While this has limited impact for nonprecessing systems, the difference can be greater in precessing systems as the final black hole spin and spin orientation will be impacted by the component spins.

In performing parameter estimation on the signal, we identify the location of the peak SNR as follows. First, we find the peak SNR of the (2, 2, 0) waveform across the mass and aligned spin space, using a fixed value of χ_p . To do so, we filter the data from each detector independently and sum the maximum in quadrature (over time and phase) of the

SNR in each detector. We use the `scipy.optimize` routine [104] with an initial guess offset from the peak—in this example, we offset $\mathcal{M} = 16.6M_\odot$, $\eta = 0.15$, and $\chi_{\text{align}} = 0.4$, although we have varied the starting point to demonstrate that it has minimal impact on the optimization result. For our example signal, we obtain values of $\mathcal{M} = 16.87M_\odot$, $\eta = 0.180$, and $\chi_{\text{align}} = 0.44$ for the peak of the (2, 2, 0) SNR, which is consistent with the peak shown in Fig. 4. For a real signal, we would use the parameters returned by the search and, indeed, gravitational wave searches have now implemented a similar maximization procedure to obtain mass and spin measurements more accurately, see e.g. [105].

Based on the discussion above, and the plots in Fig. 11, it is clear that we can obtain an improved estimate of the peak location using the two-harmonic SNR [70] which incorporates the power in the (2, 2, 1) harmonic. To do so, we perform a second optimization step, again over the chirp mass, mass ratio and aligned spin space, to identify the peak of the two-harmonic waveform. We use the previously identified peak to seed the second maximization. Although we have included the precession correction, we *do not* vary the precession spin χ_p at this stage, as we find that it is not helpful in identifying the peak. This is to be expected, given the significant degeneracy between the precession and aligned spins shown in Fig. 10. The precession spin is better constrained by the amplitude of the (2, 2, 1) harmonic, which we consider later. At present, the optimization routine is not able to accurately identify the peak of the two-harmonic SNR. We are continuing to investigate the reason for this. Consequently to obtain an accurate peak, we currently construct a dense grid of points around the (2, 2, 0) peak and filter them against the two precession harmonics to find the peak SNR. We make use of the parameter-space metric to identify the eigendirections and generate a grid of points which covers the 3σ uncertainty region around the (2, 2, 0) peak. This method identifies the peak with good accuracy, but does slow down our analysis, as filtering the grid is computationally intensive. The peak of two-harmonic SNR is identified to occur at $\mathcal{M} = 16.69M_\odot$, $\eta = 0.171$ and $\chi_{\text{align}} = 0.421$.

B. Obtain the sky location

There are several rapid sky localization analyses, most notably BAYESTAR [96], that can return the sky position of the signal quickly and accurately. Our goal here is not to present a new localization method, as that problem is already well-addressed [62,95]. Nonetheless, we do require an estimate of the sky location of the source, and its uncertainty, for several purposes. Most notably, knowing the sky position enables us to estimate the network sensitivity, and this is critical to obtaining a good estimate of the distance to the source, from Eq. (35). In addition, the sky location is used to calculate sensitivity to the second

polarization, α_{net} , and the observed power in the second polarization.

To estimate the sky location, we use a simple chi-squared minimization, as described in the Appendix of [62], to identify the preferred location. For the simulated signals discussed in this paper, the signal is observed in the LIGO-Virgo network, so that timing alone provides two locations (above and below the plane of the detectors). We generate a localization for each, as described in Sec. III C 1—in this case they overlap, so we obtain a single sky patch. Across the sky patch, we calculate the network sensitivity $|\mathbf{F}\boldsymbol{\sigma}|$ and the sensitivity to the second polarization α_{net} . For the analysis presented here, we use the mean and variance of the network sensitivity and the mean value of α_{net} in reconstructing the source parameters.

C. Generate samples in the mass and spin space

Starting from the maximum likelihood point, we calculate the approximate uncertainties in the masses and spins using the parameter space metric, introduced in Sec. III A. As has been discussed in detail [59,91,106,107], the choice of parameters used in the metric expansion can have a significant impact on the domain of validity of the quadratic approximation in Eq. (46),

$$M(\delta\vec{\lambda}) \approx 1 - g_{ab}\delta\lambda^a\delta\lambda^b \quad \text{where } \delta\vec{\lambda} = \hat{\lambda} - \vec{\lambda}.$$

We use the chirp mass \mathcal{M} , symmetric mass ratio η , aligned χ_{align} and in-plane spins, parametrized by χ_p^2 —the reason for using χ_p^2 is discussed in Sec. III E. These provide a good basis for estimating the parameter accuracy.

In many cases, the metric is calculated by taking derivatives of the waveform in various parameters [108,109], to obtain the leading-order variation. Here, we instead choose to take finite differences when evaluating the metric. This means that any higher-order terms, which are of interest for the scale of variations being considered, will be appropriately included in the metric. Since we are particularly interested in identifying the high-likelihood region, e.g. the 90% confidence interval for the physical parameters, this provides a natural scale at which to evaluate the metric. (A similar method has previously been introduced in [110]). For the four-dimensional parameter space under consideration, the 90% region is given by

$$g_{ab}\delta\lambda^a\delta\lambda^b = \frac{3.9}{\rho_h^2}. \quad (55)$$

To generate the metric, we begin with the four basis vectors in the directions $(\mathcal{M}, \eta, \chi_{\text{align}}, \chi_p^2)$. We scale each basis vector to obtain the mismatch required for the observed SNR of the system, as given in Eq. (55). In principle, the mismatch will be symmetric in steps $\pm\delta\lambda^a$ but, as can be seen in Fig. 4 this is not exactly true for finite steps. Therefore, we take the average mismatch between




FIG. 12. Degeneracy between mass ratios and spins for a signal with SNR of 10. The black ellipse shows the predicted degeneracy obtained from a metric calculated with variations along η and χ_{align} . The orange ellipse shows the final result after with variations taken along the eigendirections of the degeneracy. The final metric matches the degeneracy well, as shown in the pdf, while the initial metric significantly underestimates it.

positive and negative variations. These immediately provide the diagonal elements of the metric g_{ab} . To obtain the off-diagonal elements, we calculate the mismatch for steps in directions $\pm\delta\lambda^a \pm \delta\lambda^b$, where a and b run over the four parameters. We average over the mismatches to obtain an estimate of the off-diagonal terms in g_{ab} . Thus, by calculating a small number of matches, we obtain an expression for the metric g_{ab} .

Typically the co-ordinate directions are *not* a good choice of basis for calculating the metric, since the degenerate directions do not lie along the physical parameters, as can be seen in Fig. 4. Thus, while the uncertainties in individual parameters will be well-approximated by calculating the metric along the physical parameters, the correlations will typically be less well-estimated. To improve the accuracy of the metric, we iteratively update it to use coordinates which lie along the eigendirections in the parameter space. Specifically, we first calculate the metric using variations $\delta\lambda^a$ along each physical parameter. We then generate the eigendirections of the initial metric and test whether they do, indeed, describe the principal axes of the parameter degeneracy ellipse. If not, then we re-normalize them to have the desired mismatch [given in Eq. (55)], and recalculate the metric in these new coordinates. We stop when the metric is no longer changing significantly, specifically, we test whether the eigendirections of the old metric remain eigen-directions of the updated metric (within a given tolerance, i.e. that the vectors reproduce the desired mismatch and are orthogonal). Figure 12 shows the importance of this iterative process for a low SNR system. For higher SNRs, the effect is less significant.

The metric provides an approximate likelihood over the mass and spin parameter space, using Eq. (46). Since the likelihood is approximated as a multivariate Gaussian, we can very quickly generate large numbers of samples across the mass and spin parameter space. To do so, we generate the requested number of points drawn from four normal distributions, and then use the metric to project these to the physical parameter space.

D. Generate samples in distance and orientation

We assume that sources are distributed uniformly in volume and that their orientations are uniformly distributed. While the latter assumption should be generically valid, at small or large distances sources will not be uniformly distributed—they will either follow the local density of galaxies or cosmological effects, including a redshift dependence on the merger rate, become important. Nonetheless, it is standard to generate parameter estimates using uniform volume distribution and then reweight for astrophysical or cosmological distributions later [4,111].

A uniform in volume distribution of sources leads to a preference for the observation of face-on (or face-away) sources due to their greater gravitational wave amplitude [see, e.g. Eq. (35)] [86]. In Sec. III C 2, we have argued that the majority of observed sources will have significantly greater SNR in one of the two circular polarizations. Therefore, we wish to obtain the distribution for $\cos\theta_{JN}$ for a signal observed with a fixed SNR in the right or left circular polarization. This is given as

$$p(\cos\theta_{JN}) \propto \int d_L^2 dd_L d \cos\theta_{JN} \times \delta\left(\rho_o - |\sigma\mathbf{F}| \frac{d_o}{d_L} \left(\frac{1 \pm \cos\theta_{JN}}{2}\right)^2\right), \quad (56)$$

where δ denotes the Dirac delta function. It is straightforward to marginalize over the distance distribution and identify that the (originally flat in $\cos\theta_{JN}$) distribution becomes

$$p(\cos\theta_{JN}) \propto (1 \pm \cos\theta_{JN})^6, \quad (57)$$

where the positive/negative corresponds to right/left polarization respectively. We use this distribution to generate samples in θ_{JN} corresponding to left and right circularly polarized signals.

Given the sky location of the source and the (complex) SNR observed in each detector, it is straightforward to obtain the SNR in the right and left circular polarizations, $\rho_{R,L}$. As described in [98], we achieve this by first rotating to the dominant polarization frame, and calculating the network's sensitivity to the $+$ and \times polarizations, w_+ and w_\times respectively. We then use the given sky location and project the (complex) SNR observed in each detector onto the space of circularly polarized signals using,

$$P_{\text{circ}}^{ij} = \left[\frac{(w_+^i \pm iw_\times^i)(w_+^j \pm iw_\times^j)}{|w_+|^2 + |w_\times|^2} \right], \quad (58)$$

where i and j run over the detectors in the network and the $+/-$ gives the SNR in the left and right polarizations respectively. The relative SNR in each of the circular polarizations can be used to appropriately weight the probability that the signal is predominantly either left or right circularly polarized. Since the likelihood is proportional to $\exp[\rho_{R,L}^2/2]$, this provides the appropriate normalization factor to weight the number of samples drawn from the left and right circular polarization distributions for θ_{JN} . In many cases, the signal will be preferentially right (or left) circularly polarized, in which case the majority of samples will correspond to face-on (or face-away) orientation. If α_{net} is small, it is often impossible to distinguish between left and right circular polarizations. In this case, there will be large numbers of samples for both face-on and face-away signals, with a minimum at edge on since these systems emit the weakest gravitational wave signal.

Given the binary orientation θ_{JN} , we can obtain the distance from Eq. (35), which we repeat below,

$$\rho_o = \frac{d_o}{d_L} \frac{|\sigma\mathbf{F}|}{(1 + \tau^2)^2}.$$

The inferred distance depends upon the network response $|\sigma\mathbf{F}|$ through the detector response encoded in \mathbf{F} and the masses and spins through the detector sensitivity to the signal encoded in σ . We incorporate both of these effects when estimating the distance. For the network sensitivity, we simply draw samples from a Gaussian, based upon the previously measured mean and variance. To incorporate variations in the mass and spin space, we must recalculate σ for each sample, given the values of mass and spin. Since σ is a slowly varying function across mass and spin (it will typically vary by at most tens of percent over the mass-spin posterior), we first interpolate across the space and then use the interpolation function to evaluate σ at each of the samples. Finally, the observed SNR has measurement uncertainty, which is well-modeled by a noncentral χ^2 distribution with two degrees of freedom and a noncentrality parameter $\rho_{R,L}$. Thus, given a value of θ_{JN} and σ for each sample, we randomly sample ρ_o and $|\mathbf{F}|$ from the appropriate distributions and use Eq. (35) to calculate the distance.

E. Restrict the parameters using additional waveform components

In the previous subsections, we have described a method to obtain samples in masses, spins, distance and binary orientation. This provides a good initial estimate of the binary parameters but we have additional information which can still be used to improve the parameter estimates.

In particular, we have not yet used the measured SNRs in the second polarization, higher multipoles or precession.

The measured SNR in the second polarization, precession and higher multipoles can be used to improve measurement of the binary parameters, as discussed in detail in Secs. III C 2, III D, and III E. For each of the samples, we calculate the *expected* SNR in each of these features, using the given masses, spins, distance and orientation. In particular, the SNR in the second polarization is given by Eq. (36) and depends upon the orientation θ_{JN} and α_{net} . The expected SNR in higher multipoles is given in Eq. (37) and depends upon the orientation and relative significance of the higher multipoles, encoded in $\alpha_{\ell m}$, which is primarily determined by the mass ratio [69]. The expected SNR in precession is given in Eq. (38) and depends upon the binary orientation and the opening angle between the orbital and total angular momenta. The opening angle is largely determined by the in-plane spins, χ_p , but also varies with mass ratio and aligned spin. Those samples where the expected SNR in these waveform features matches the observed SNR are given a higher weighting than those where the expected and observed SNRs differ significantly.

We calculate the observed SNR in the second polarization, precession and higher multipoles by matched filtering the waveform components, evaluated at the maximum likelihood point identified in Sec. IV A, against the data. In detail, the precession and higher multipole SNRs are obtained by matched filtering the $h_{22,1}$ and $h_{33,0}$ waveforms against the data from the network of detectors and evaluating the SNR in each detector at the time where the SNR for $h_{22,0}$ is maximum. To obtain the SNR in the second polarization, we project the SNR in the (2, 2, 0) component into the space orthogonal to the right/left circular polarization, as described in Sec. II C. For each sample, we calculate the likelihood of obtaining the observed SNR in these waveform components—it is given by a noncentral chi-squared distribution with two degrees of freedom, where the noncentrality parameter is the expected SNR at the parameter values of the sample and the distribution is evaluated at the observed SNR. Thus, points in the parameter space which accurately predict the observed SNR in the second polarization, higher multipoles and precession are preferred to those which predict either too much or too little SNR in these features. We assign weights to each of the samples, based upon the product of the probabilities for obtaining the given SNR in each of the three features, and then use these weights to importance sample the points to produce our final result.

In the above, we have used only the SNRs calculated at the maximum likelihood point. This has the benefit of significantly reducing computational cost, as we do not need to recompute the likelihood at every sample. From Fig. 11, we see that the distribution of the higher multipole SNR has a similar structure to the leading waveform

component, although in the example show the peak is somewhat offset. Furthermore, in the region of mass and spin space consistent with the (2, 2, 0) waveform (shown, e.g., in Fig. 4) the SNR in the (3, 3, 0) waveform component varies by around 10%. Therefore, it is a reasonable approximation to take this to be constant. For the second polarization, the variation of SNR with an incorrect estimate of θ_{JN} (which will lead to an over/underestimation of the SNR in the second polarization) should be independent of the masses and spins. For precession, the SNR distribution in mass and spin space is significantly different than the (2, 2, 0) waveform. Nonetheless, the variation in precession SNR over the mass and spin space of interest is $\approx 30\%$ and so we can reasonably treat this as constant. Finally, for speed, we precompute slowly varying quantities such as \bar{b} and α_{33} coarsely over the parameter space and then interpolate them when calculating the expected SNR in each waveform component.

F. Summary

The `simple-pe` algorithm generates a set of discrete samples that approximate the mass, spin, distance and orientation posterior distribution for the observed gravitational wave signal. The mass and spin distributions are approximated as multidimensional Gaussian distributions centered around the point with the maximum SNR in the (2, 2) multipole (incorporating two precession harmonics). While the sky location is inferred from the SNRs in the network, the information is only used in estimating the overall network sensitivity to the two gravitational wave polarizations, and their variation over the localization region. The distance and orientation are generated assuming uniformly distributed events in both volume and binary orientation, and weighting for observability. Correlations between the inferred distance/orientation and mass/spin parameters arise through the requirement that the binary parameters accurately predict the observed power in the second polarization, higher multipoles and precession.

Finally, we must discuss priors. So far, we have been evaluating the likelihood, either through the metric expansion or via observed SNRs in different waveform components. We have explicitly considered priors on distance and orientation, where we chose a distribution uniform in volume and in $\cos\theta_{JN}$. This means that, over the mass and spin space, we have effectively imposed flat priors in the parameters that we are using to generate samples, namely $\mathcal{M}, \eta, \chi_{\text{align}}$ and χ_p^2 . In the majority of analyses, priors are chosen to be uniform in component masses, spin magnitudes, and orientations, see e.g. Refs. [2, 4, 112, 113]. While, for many systems, the masses are well-constrained and the differing prior causes little impact, this is not the case for equal mass systems. To transform to a prior which is uniform in component masses requires

$$p(\mathcal{M}, \eta) \propto \frac{\mathcal{M}}{\sqrt{1-4\eta}} \quad (59)$$

which blows up at equal mass ($\eta = \frac{1}{4}$). To address this, we limit the denominator to 50, effectively truncating the prior at $\eta = 0.2499$. This leads to a minor under-sampling of near-equal mass systems (with a mass ratio between 1 and 1.04). Similarly, we transform from uniform priors on χ_{align} and χ_p^2 to uniform priors on component spin magnitudes and orientations using the transformations given in [114]. These priors are implemented by drawing a large number of samples from the underlying distributions and then using rejection sampling to retain samples with a distribution matching the desired priors, see, e.g. Appendix C in Ref. [18].

The entire analysis presented above runs in a few minutes on a single CPU.⁹ It can be neatly separated into two parts: in the first, we matched-filter the data to identify the maximum SNR point, the sky location and power in each of the waveform components: second polarization, higher multipoles, and precession. For the estimation of the binary parameters, we make use of these values and then calculate distributions in masses, spins, orientation and distance. This second step requires the PSD of the detector data to obtain the metric, but not the data themselves.

V. SIMPLE-PE RESULTS

We use four simulated signals to demonstrate an application of the `simple-pe` algorithm described in Sec. IV. The goal of the examples is to demonstrate the accuracy of the parameter estimation across the mass, spin, distance, and orientation space. In particular, we are interested in seeing that the quadratic approximation from the metric is appropriate for generating samples in the mass and spin space. Furthermore, we would like to investigate how the observed SNR in the additional waveform components can impact parameter recovery. To do so, we choose signals with many parameters in common, and vary only those which will impact the SNR in the second polarization, higher multipoles and precession. The common parameters of the systems are provided in Table I. In all cases, the signals have masses of $40M_\odot$ and $10M_\odot$, with aligned spin components of 0.5 on the more massive black hole and zero on the smaller system. The binaries are located at a fixed sky location (the same as shown in Fig. 5).

The binary orientation and in-plane spins are varied to give examples with negligible or significant power in both higher multipoles (HM) and precession (prec). The power in the second polarization also varies but, in none of the

⁹In the current implementation, filtering the two-harmonic SNR over a grid to find the peak takes around 15 minutes and dominates the computational time. Other techniques, such as relative binning [28], may reduce this computational cost, and we leave an investigation future work.

TABLE I. Common parameters of the simulated signals.

Source-frame primary mass, [M_\odot]	40.0
Source-frame secondary mass, [M_\odot]	10.0
Source-frame chirp mass, \mathcal{M} [M_\odot]	16.65
Symmetric mass ratio, η	0.16
yth component of the primary spin	0
zth component of the primary spin	0.5
Secondary spin magnitude	0
Effective spin, χ_{eff}	0.4
Aligned spin, χ_{align}	0.42
Right ascension, α [rad]	1.2
Declination, δ [rad]	0.3
Polarization, ψ [rad]	0.5
Network alignment factor, α_{net}	0.364
GPS merger time, [s]	1677672000.0

cases we have considered is it sufficient to enable confident observation of the second polarization. We fix the distance to the signals to ensure a SNR of 25 in the LIGO-Virgo network with expected O4 sensitivity. The set of varying parameters for each signal are given in Table II, along with the SNR in each of the waveform components.

In Fig. 13, we show the `simple-pe` posterior obtained for signal (4), with negligible power in the second polarization, but significant power in both higher multipoles and precession. The posteriors obtained are in good agreement with the simulated values, with both the θ_{JN} and χ_p distributions clearly peaked away from zero. In addition, we see clear correlations between several pairs of parameters. As expected, the mass ratio and aligned spin are (anti) correlated due to their impact on the phase evolution of the waveform. The distance and binary orientation are also anticorrelated based upon the measured SNR in the (2, 2, 0) waveform component—to achieve a given SNR an inclined signal must be at a smaller distance. Finally, the binary orientation and precessing spin are (anti)correlated as larger values of either leads to a larger SNR in precession. We also observe a slight correlation between orientation and mass ratio. This arises from the necessity to obtain the correct SNR in higher multipoles—the SNR increases with θ_{JN} but decreases for increasing η (more equal mass binaries).

TABLE II. Parameters that vary for the simulated signals.

Signal	(1)	(2)	(3)	(4)
	HM	Prec	HM-prec	
Effective precession spin, χ_p	0.01	0.1	0.7	0.4
Opening angle, β	0.003	0.06	0.4	0.2
Angle between L and N , ι	0.005	0.6	0.1	0.6
Angle between J and N , θ_{JN}	0.006	0.6	0.4	0.6
Luminosity distance, d_L [Mpc]	982	825	859	814
SNR in (3,3) multipole, ρ_{33}	<0.1	4.3	2.8	4.2
SNR in precession, ρ_p	<0.1	1.02	4.4	4.4
SNR in second polarization	<0.1	0.1	<0.1	0.2

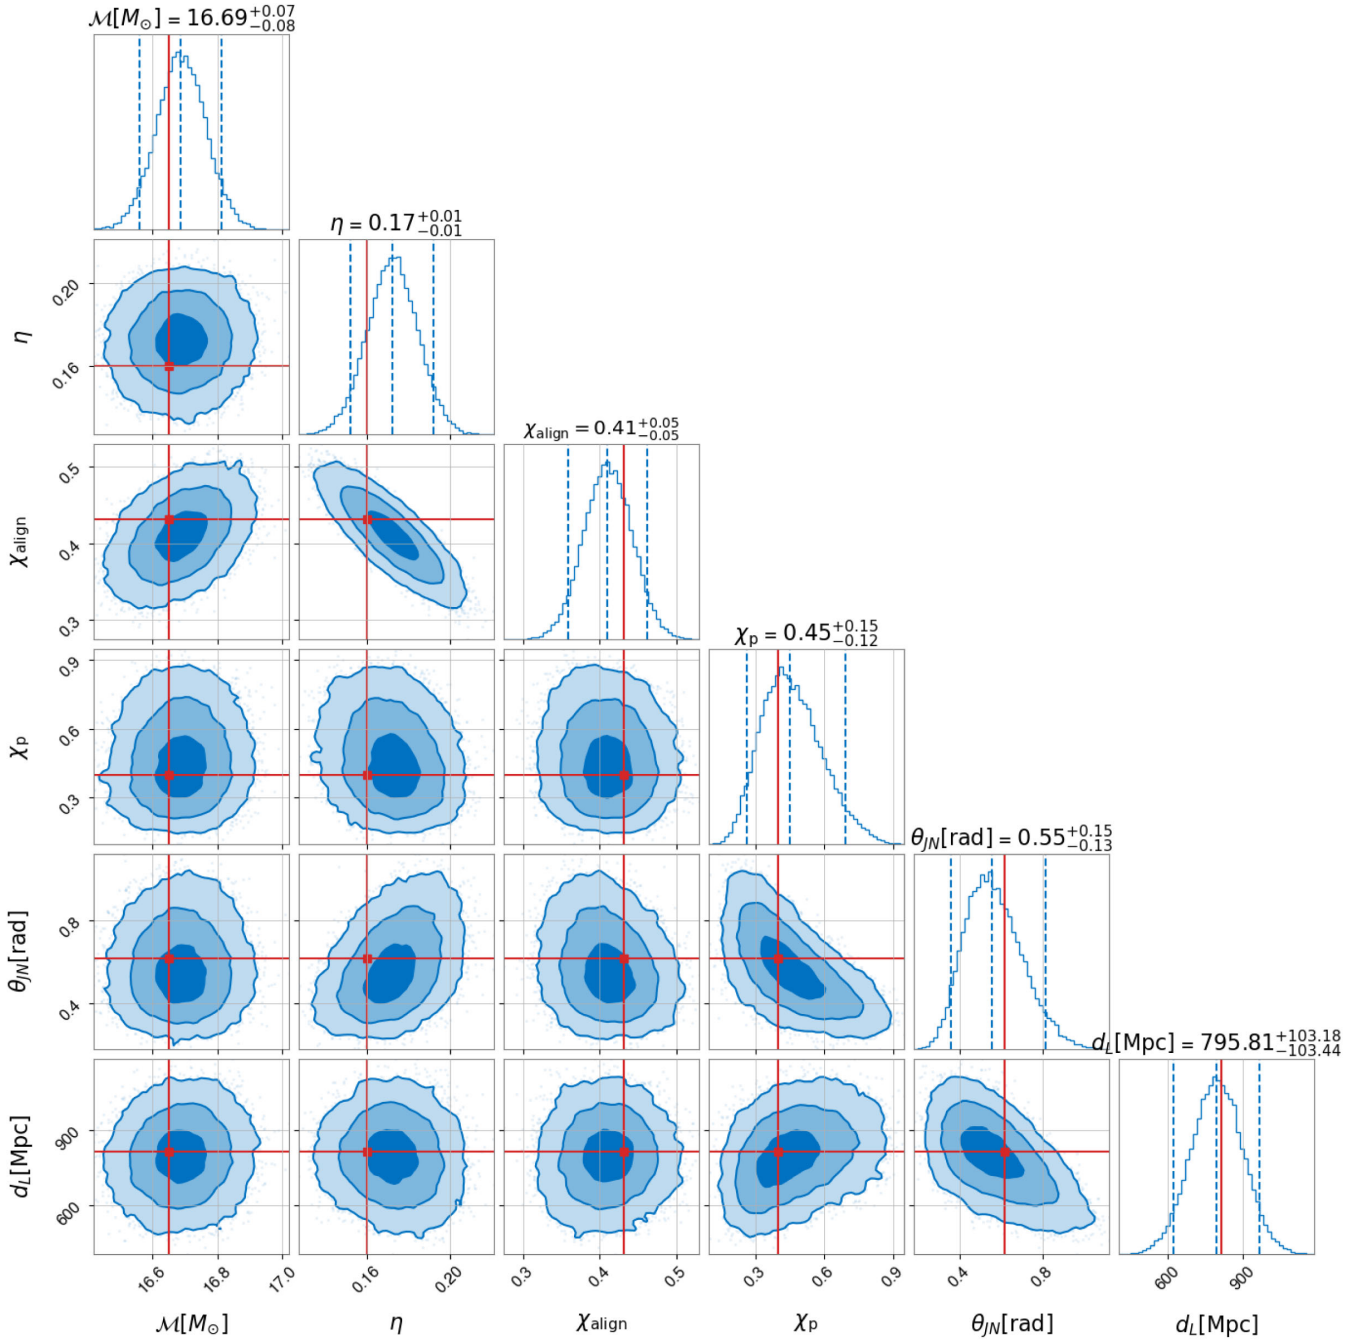


FIG. 13. The posterior distribution for masses (chirp mass \mathcal{M} and symmetric mass ratio η) spins (aligned χ_{align} and precessing χ_p), distance d_L and binary orientation θ_{JN} . The shaded regions in the two-dimensional figures show the 1, 2, and 3 σ regions for the parameters and the one-dimensional distributions show the median and 90% symmetric confidence interval. The simulated values are shown as red lines. For all of the parameters, the inferred values are in good agreement with the simulation. The binary was deliberately generated to have negligible power in the second polarization, and significant power in both higher multipoles and precession, signal (4) in Table II.

Next, we consider how the parameter recovery varies for the three other systems listed in Table II. While the distributions of all parameters do change somewhat, we focus on the binary orientation, θ_{JN} , and precessing spin, χ_p . These parameters have the greatest impact on the

observed SNR in precession and higher multipoles and we therefore expect them to be most impacted by the changes. Of course, the distance will also vary (as is clear from Table II), but this is primarily due to our requirement of a fixed SNR in the signal. Although the mass ratio

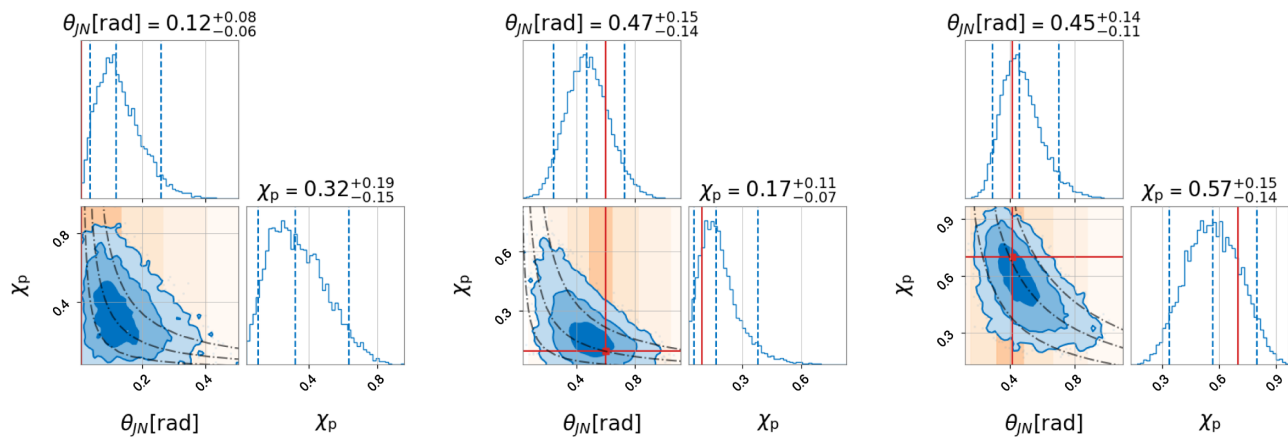


FIG. 14. The posterior distribution for the binary’s orientation θ_{JN} and precession spin χ_p for 3 simulated signals; Left: a binary with insignificant power in precession and higher-order multipoles, signal (1). Middle: a binary with significant power in higher-order multipoles, signal (2), and Right: a binary with significant power in precession and some power in higher-order multipoles, signal (3). The blue shaded regions in the two-dimensional figures show the 1, 2, and 3 σ regions for the parameters. The black dash-dotted lines show lines of constant ρ_p , and the orange shaded regions show lines of constant ρ_{33} ; the darker region encompasses the injected value. The one-dimensional distributions show the median and 90% symmetric confidence interval. The simulated values are shown as red lines. We note that the priors on θ_{JN} and χ_p vanish at zero leading to zero posterior support there, and consequently posteriors peaked away from the injected values, as shown in the left-hand panel.

remains fixed, one might expect its recovery to vary when higher multipoles and precession are observed (as was the case for GW190412, [12]). However, since the signal is already identified as unequal mass from the (2, 2, 0) waveform, they have little impact on the mass ratio distribution. For all 3 signals—one with negligible power in both higher multipoles and precession, one with significant power in higher multipoles, and another with significant power in precession—we see that the recovered θ_{JN} and χ_p are in good agreement with the simulated values. For all 3 cases, the simulated values lie within the 1 σ confidence interval. Furthermore, the inferred ranges of θ_{JN} and χ_p can be explained by the SNR in higher-order multipoles and precession. In each case, the binary orientation is restricted by the measured (3, 3, 0) SNR. The precession SNR increases with both χ_p and θ_{JN} as, to a good approximation it scales with $\chi_p \cdot \theta_{JN}$. Thus, the measured precession SNR limits the permitted range in the two-dimensional space. As expected a higher observed (3, 3) multipole SNR leads to a larger inferred θ_{JN} while a larger precession SNR leads to larger values of both θ_{JN} and χ_p .

The events we have investigated were deliberately chosen to lie in a region of the parameter space where higher-order multipoles, precession and power in the second polarization are all likely to be observed. In particular, we selected a high SNR signal with unequal masses, observed well away from face-on. For this example, even when there is no power in higher multipoles or precession, the lack of power can be used to restrict the physical parameters, as shown in 14. The majority of

gravitational wave observations are expected to be low SNR and close to equal mass. For these systems, it is very unlikely that additional waveform features will be identifiable. In this case, lack of power in precession, higher multipoles and second polarization have limited impact on the inferred parameters. In Appendix B.

A. Comparison with other samplers

In Fig. 15, we show the posterior distributions for a signal with power in the second polarization, and significant power in higher multipoles and precession obtained using `simple-pe` (as shown in Fig. 13) with that obtained using `DYNESTY` [115] (generated through the `Billby` infrastructure [16,18,20]). `DYNESTY` was chosen as it is commonly used for gravitational wave follow-up analyses (see e.g. Refs. [4,7]). We see that in general, there is good agreement between the two sampling techniques, with the `simple-pe` results algorithm presented here consistent with `DYNESTY` for all parameters. The `simple-pe` posteriors are somewhat broader than those obtained with the full Bayesian parameter estimation routine. Most notably, the distance and precession spin posteriors are significantly broader. The distance uncertainty arises from three effects; a degeneracy between distance and orientation, a degeneracy between distance and mass, and a variation of distance over the sky region. For this signal, it is the variation of detector sensitivity over the sky patch that gives the greatest contribution to the distance uncertainty. While our localization routine returns a reasonable approximation to the sky patch, it does give a significantly larger region than `DYNESTY` and this explains the difference in distance

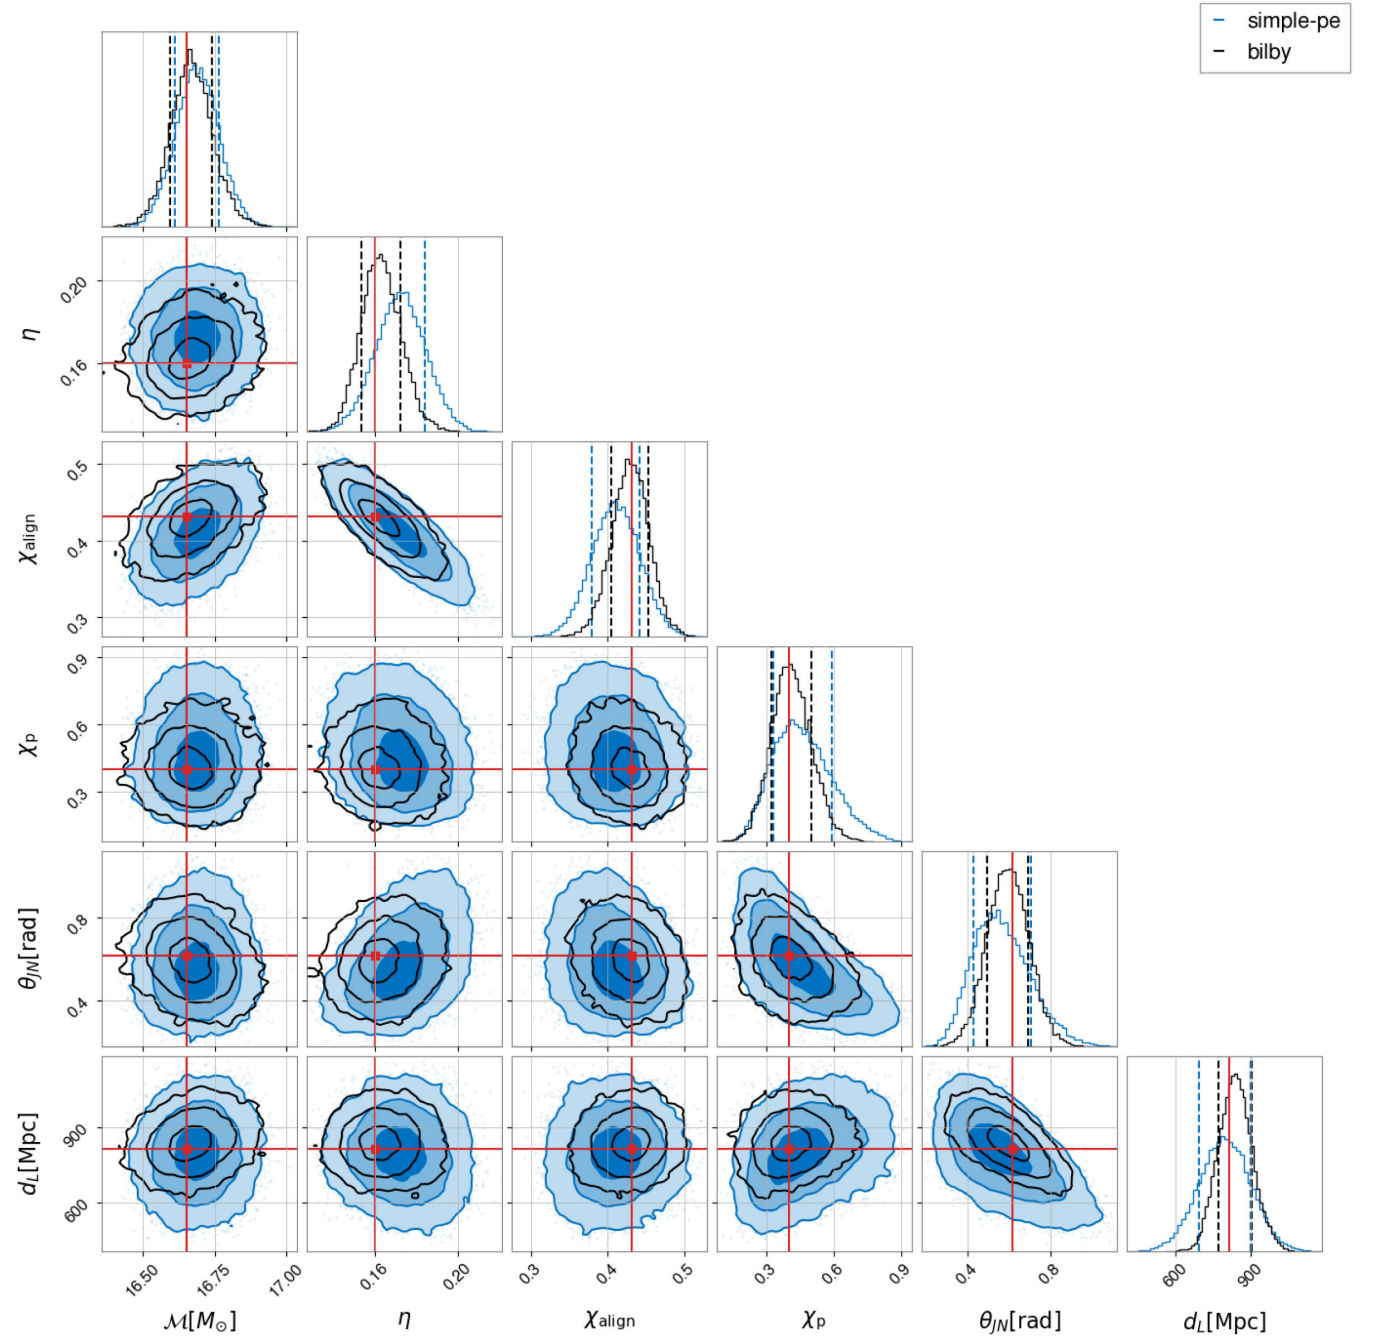


FIG. 15. Same as Fig. 13 but with the posteriors obtained with DYNESTY (generated through the Bilby infrastructure [16,18,20]) overlaid in black. DYNESTY took ~ 20 hours to complete when parallelized over 400 CPUs while `simple-pe` took ~ 20 minutes on a single CPU.

inference. The precessing spin is also less well-measured by the `simple-pe` result. Similar results are obtained for the three other systems discussed in Sec. V, namely that the `simple-pe` posteriors are comparable with, although typically broader than those obtained with DYNESTY. Importantly, DYNESTY took ~ 20 hours to complete when parallelized over 400 CPUs while the work presented here took ~ 20 minutes on a single CPU.

VI. DISCUSSION

In this paper, we have presented a detailed discussion of the primary way in which different physical parameters impact the observed gravitational wave signal emitted during a black hole binary coalescence. Using these insights, we have developed a parameter estimation method, `simple-pe`, that uses this information to provide

estimates of the physical parameters of the system. By restricting focus to the *observable* impact on the gravitational waveform, we are able to obtain rapid parameter estimates which naturally incorporate degeneracies between physical parameters. We have presented parameter estimation results for a small set of signals and, in Sec. VA, shown that our results are in good agreement with those obtained by the Bilby [16] sampler which is used in interpreting gravitational wave observations [4]. The `simple-pe` results are obtained in a fraction of the time and recover broadly the same results as the more intensive samplers. There have been numerous other analyses developed which provide rapid parameter estimation for gravitational wave observations [17,20–26,29,30,90,96]. However, many of these use waveforms that exclude either precession or higher multipoles or both. Thus, while they are able to provide reasonable estimates of masses, aligned spins and sky locations, they are unable to probe some of the interesting astrophysics that will only be uncovered through higher multipoles (which are one of the best ways to clearly identify unequal mass systems) or precession (which is the only way to probe in-plane spins). Furthermore, the `simple-pe` analysis has the advantage of providing clear interpretation as to *why* and *how* the various parameters are measured and the accuracy achieved. For example, in Fig. 14, we can understand the inferred orientation and in-plane spins based upon the higher multipole and precession SNRs.

Here, we have restricted attention to a limited number of simulated signals, primarily to describe the methods employed and the properties of the observed signal which enable measurement of various parameters. In the main paper, we have focused on an example where additional waveform features are observable and can therefore be used to improve parameter recovery. In Appendix B, we present an example for a low SNR signal with close to equal masses for which higher multipoles and precession have minimal impact. This does not constitute a comprehensive test of the analysis. In a future paper, we plan to present an in depth investigation of `simple-pe` results using both a large set of simulations and also existing GW observations [4].

In future observing runs, the rate of observed binary mergers will increase as instrumental sensitivity improves. For example, in the O4 run [10], several events per week are expected and in O5 and beyond we could regularly be observing multiple events per day. Thus, it becomes ever more important to obtain detailed estimates of the binary properties in a short time-frame—comparable to the time between events. The `simple-pe` analysis provides a tool that can become part of the process for quickly understanding the observations. We are able to provide good parameter estimates in a matter of minutes. The analysis, inevitably, introduces some approximations, for example by considering only the most significant higher multipole and precession harmonic. In addition, in the current implementation, we do not extract the phase, polarization

or precession phase angles, although we have briefly discussed how this could be done. Finally, the `simple-pe` analysis does not, currently, include calibration uncertainties which can broaden the recovered parameter widths [116]. Thus, in its current implementation, `simple-pe` will not provide the most accurate estimate of the binary parameters. Nonetheless, a fast and reasonably accurate estimate of the binary parameters is useful to rapidly identify interesting events for prioritized analysis and provide parameter estimates to guide follow-up electromagnetic observations, where knowledge of masses, mass ratio and binary orientation can impact the expected signal [117]. Furthermore, there may be applications, for example, studies of the black hole population [111] or searches for lensed events [118], where the accuracy provided by `simple-pe` is sufficient to obtain the results. Finally, where full parameter estimation is required, the initial parameter estimates provided by `simple-pe` can be used to inform the detailed parameter estimation routines. For example, the posteriors obtained by `simple-pe` could be importance sampled by Bilby to produce a final set of posterior samples, as has been done previously to obtain faster results for higher multipoles and eccentricity [13,119].

There are a number of additional features which can lead to an observable effect in the emitted gravitational waveform, including matter effects, eccentricity and mode asymmetries. For binaries containing one or two neutron stars, the waveform will carry information about the structure of the star, this arises as additional post-Newtonian corrections during the inspiral and also a difference in the merger and post-merger signal from that of a black hole binary [120,121]. For binaries on non-circular orbits, the emitted gravitational wave signal carries an imprint of the orbital eccentricity [122]. For several of the binaries observed in O3, there is tantalizing evidence of eccentricity [13,123]. Finally, for binaries where the orbit precesses, there are asymmetries between the gravitational wave signal emitted above and below the orbital plane. These correspond to the “bobbing” of the binary and the final kicks given to the system following merger [124,125]. Detailed studies of these additional features, and their impact on parameter measurements, have been pursued by several groups. Studies of neutron star structure include [126,127]; eccentricity had been studied in [13,128] and mode asymmetries are discussed in Refs. [125,129].

All three of these features are amenable to the approach presented in this paper. The presence of neutron-star structure will impact the overall frequency and phase evolution of the system, and is likely to be primarily measured from the leading waveform component. Mode asymmetries lead to a difference between, e.g. the (2, 2) and (2, -2) multipoles. By separating into the leading (symmetric) and subdominant (antisymmetric) components, we obtain an additional waveform component

whose amplitude will help constrain the in-plane spins. Eccentricity impacts the overall evolution of the waveform, and it is well-known that the eccentricity is degenerate with mass (both higher eccentricity and higher mass lead to a faster merger). However, eccentricity also leads to additional harmonics in the waveform, at multiples of the orbital frequency [130]. Therefore, it seems likely that a combination of techniques will be required to incorporate eccentricity. Finally, we note that most waveform models include only a subset of the physical effects we have discussed, as each additional effect expands the dimensionality of the parameter space to be simulated. Our approach of identifying the leading effect on the waveform of each new (astro)physical phenomenon has the potential to help guide development of future waveform models. Additionally, even if complete models are not available, we could incorporate all of their leading-order effects into the `simple-pe` analysis to obtain a more complete picture of each observed binary.

ACKNOWLEDGMENTS

We thank Tom Dent, Ilya Mandel, Frank Ohme, Javier Roulet, Rory Smith, Eric Thrane, Veronica Villa, and Barak Zackay for useful discussions and suggestions. We thank Ben Farr and Javier Roulet for detailed comments on an earlier draft of the paper. We thank Mark Hannam and Vivien Raymond for continued discussions and insights throughout this project. S. F., R. G., C. H., C. M., and S. U. thank STFC for support through the Grants No. ST/V005618/1 and No. ST/N005430/1. C. H. thanks the UKRI Future Leaders Fellowship for support through Grant No. MR/T01881X/1. S. F. was supported by a Leverhulme Trust International Fellowship. The authors are also grateful for computational resources provided by LIGO Laboratory and supported by National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459 and provided by Cardiff University, and funded by STFC Grant No. ST/I006285/1. `simple-pe` is programmed in Python and utilizes modules from NumPy [131], SciPy [104], PyCBC [132] and PESummary [133]. Plots were prepared with Matplotlib [134], and PESummary [133]. As part of this paper, the `parallel-bilby` parameter estimation software [20], which made use of the DYNESTY nesting sampling package [115], was used for comparisons.

APPENDIX A: MULTIPOLE DECOMPOSITION OF THE GRAVITATIONAL WAVEFORM

In the body of the text, we derived the expression for the gravitational waveform as a sum over multipole moments (ℓ, m) with an additional “splitting” of the harmonics due to precession, given in Eq. (9), which we repeat here,

$$h = \sum_{\ell, m, n} {}^{-2}Y_{\ell, m}(\theta, \phi) D_{n, m}^{\ell}(\alpha, \beta, \epsilon) h_{\ell, n}^{\text{NP}}(t, \vec{\lambda}). \quad (\text{A1})$$

Our aim is to explicitly extract the waveform dependence upon the opening angle β and to show that we can decompose the waveform into components with a natural hierarchy in the parameter $b = \tan(\beta/2)$.

We begin by expanding the spherical harmonics using

$${}^{-2}Y_{\ell, m}(\theta_{JN}, -\alpha_o) = \sqrt{\frac{2\ell+1}{4\pi}} d_{m, 2}^{\ell}(\theta_{JN}) e^{-im\alpha_o} \quad (\text{A2})$$

and

$$D_{n, m}^{\ell}(\alpha, \beta, \epsilon) = e^{im\alpha} d_{n, m}^{\ell}(-\beta) e^{-in\epsilon}, \quad (\text{A3})$$

to obtain

$$h = \sum_{\ell} \sum_{m, n=-\ell}^{\ell} \sqrt{\frac{2\ell+1}{4\pi}} \times h_{in}^{\text{NP}} d_{m, 2}^{\ell}(\theta_{JN}) d_{n, m}^{\ell}(-\beta) e^{-in\epsilon} e^{im(\alpha-\alpha_o)}. \quad (\text{A4})$$

Next, we impose the restriction that the gravitational wave emission is symmetric above and below the plane of the binary so that $h_{\ell, n} = (-1)^{\ell} h_{\ell, -n}^*$.¹⁰ This allows us to express the waveform in terms of multipoles with $n \geq 0$ as

$$h = \sum_{\ell} \sum_{n=0}^{\ell} \sum_{m=-\ell}^{\ell} \sqrt{\frac{2\ell+1}{4\pi}} d_{m, 2}^{\ell}(\theta_{JN}) e^{im(\alpha-\alpha_o)} \times [h_{in}^{\text{NP}} d_{n, m}^{\ell}(-\beta) e^{-in\epsilon} + (-1)^{\ell} (h_{in}^{\text{NP}})^* d_{-n, m}^{\ell}(-\beta) e^{in\epsilon}]. \quad (\text{A5})$$

Next, we collect together the β terms as we want to obtain an expansion in terms of the parameter $b = \tan(\beta/2)$. To do so, we use the discrete symmetries of the Wigner d -matrices, $d_{n, m}^{\ell}(-\beta)$,

$$d_{n, m}^{\ell} = (-1)^{m-n} d_{m, n}^{\ell} = (-1)^{m-n} d_{-n, -m}^{\ell} \quad (\text{A6})$$

and a relabeling of $m \rightarrow -m$ in the second term of Eq. (A5) to extract the factor $d_{n, m}^{\ell}(-\beta)$ from both terms. In addition, for consistency with previous work, we relabel the indices by switching n and m . This enables us to use (ℓ, m) to label the multipole moments of the waveform in the frame aligned with the orbital angular momentum and $n \in [-\ell, \ell]$ to denote the precession harmonics.¹¹ Combining these steps leads to a final expression for the waveform in terms of nonprecessing multipoles as

¹⁰This is a reasonable approximation but, particularly for large in-plane spins, there is an asymmetry in the emitted radiation [43]. The formalism presented here can be extended to that more general case, but we leave that to future work.

¹¹Following [70], we will later introduce $k \in [0, 2\ell + 1]$ as the harmonic labeling precession, and show that the harmonics have a power-law structure with index k .

$$\begin{aligned}
h &= \sum_{\ell} \sum_{m=0}^{\ell} \sum_{n=-\ell}^{\ell} \sqrt{\frac{2\ell+1}{4\pi}} d_{m,n}^{\ell}(-\beta) \\
&\times [d_{n,2}^{\ell}(\theta_{JN})(h_{lm}^{\text{NP}} e^{-im\epsilon} e^{in(\alpha-\alpha_o)}) \\
&+ (-1)^{\ell-m} d_{n,-2}^{\ell}(\theta_{JN})(h_{lm}^{\text{NP}} e^{-im\epsilon} e^{in(\alpha-\alpha_o)})^*]. \quad (\text{A7})
\end{aligned}$$

In the following subsections, we examine the structure of the waveform in Eq. (A7) by considering the (2, 2) and (3, 3) multipoles explicitly, as well as the general case of the (ℓ, ℓ) and $(\ell, \ell - 1)$ multipoles.

1. The (2,2) two-harmonic waveform

The two-harmonic waveform introduced in [70] can be recovered by restricting to the (2, 2) multipole in Eq. (A7) and then further restricting to the two leading-order terms in b . We briefly summarize the calculation below.

The Wigner d -matrices for $\ell = m = 2$ are given by

$$\begin{aligned}
d_{2,n}^2(-\beta) &:= C_{2,n} \cos^{2+n}(-\beta/2) \sin^{2-n}(\beta/2) \\
&= C_{2,n} \frac{(-b)^{2-n}}{(1+b^2)^2}, \quad (\text{A8})
\end{aligned}$$

where $C_{2,\pm 2} = 1$, $C_{2,\pm 1} = 2$, $C_{2,0} = \sqrt{6}$, and $b = \tan(\beta/2)$. This enables us to express the waveform amplitudes as a power series in b . Similarly, it is convenient to express the Wigner d -matrices for θ_{JN} appearing in Eq. (A7) in terms of

$$\tau := \tan(\theta_{JN}/2). \quad (\text{A9})$$

Restricting to $\ell = m = 2$ in Eq. (A7) and substituting the explicit form of the Wigner d -matrices for both $-\beta$ and θ_{JN} gives

$$\begin{aligned}
h_{22} &= \sum_{n=-2}^2 \sqrt{\frac{5}{4\pi}} \frac{(C_{2,n})^2 b^{2-n}}{(1+b^2)^2} \\
&\times \left[\frac{\tau^{2-n}}{(1+\tau^2)^2} (h_{22}^{\text{NP}}(t) e^{-2i\epsilon} e^{in(\alpha-\alpha_o)}) \right. \\
&\left. + \frac{(-\tau)^{2+n}}{(1+\tau^2)^2} (h_{22}^{\text{NP}}(t) e^{-2i\epsilon} e^{in(\alpha-\alpha_o)})^* \right]. \quad (\text{A10})
\end{aligned}$$

This shows the desired structure, with each of the precession harmonics appearing with a factor of b^{2-n} . Using the fact that $b < 1$, we can restrict to the two leading-precession harmonics (b^0 and b^1). Specifically, we introduce

$$\begin{aligned}
h_{22,0} &= \sqrt{\frac{5}{4\pi}} \frac{1}{(1+b^2)^2} h_{22}^{\text{NP}}(t) e^{-2i\epsilon} e^{2i(\alpha-\alpha_o)}, \\
h_{22,1} &= \sqrt{\frac{5}{4\pi}} \frac{b}{(1+b^2)^2} h_{22}^{\text{NP}}(t) e^{-2i\epsilon} e^{i(\alpha-\alpha_o)}. \quad (\text{A11})
\end{aligned}$$

Both waveforms accumulate a secular phase of $2i\epsilon$ relative to the nonprecessing waveform, as first noted in [79]. The amplitude of the second precession harmonic is reduced by a factor of b relative to the first and the frequency is reduced by the precession frequency Ω_p , where $\dot{\Omega}_p = \alpha$.

The two-harmonic waveform is given as

$$h_{22} \approx \frac{(h_{22,0} + \tau^4 h_{22,0}^*)}{(1+\tau^2)^2} + \frac{4\tau(h_{22,1} - \tau^2 h_{22,1}^*)}{(1+\tau^2)^2}, \quad (\text{A12})$$

reproducing the expression provided in [70].

2. The (3, 3) waveform

Let us now obtain a similar expression for the (3, 3) multipole, following a similar procedure as in the previous section. As before, we wish to evaluate the multipole amplitude explicitly in terms of b and τ . The required Wigner d -matrices are

$$\begin{aligned}
d_{3,n}^3(-\beta) &:= C_{3,n} \cos^{3+n}(\beta/2) \sin^{3-n}(-\beta/2) \\
&= C_{3,n} \frac{(-b)^{3-n}}{(1+b^2)^3}, \quad (\text{A13})
\end{aligned}$$

where $C_{3,\pm 3} = 1$ and $C_{3,\pm 2} = \sqrt{6}$ and, as before, $b = \tan(\beta/2)$. This gives

$$\begin{aligned}
h_{33} &= \sum_{n=-3}^3 \sqrt{\frac{7}{4\pi}} C_{3,n} \frac{(-b)^{3-n}}{(1+b^2)^3} \\
&\times [d_{n,2}^3(\theta_{JN})(h_{33}^{\text{NP}} e^{-3i\epsilon} e^{in(\alpha-\alpha_o)}) \\
&+ d_{n,-2}^3(\theta_{JN})(h_{33}^{\text{NP}} e^{-3i\epsilon} e^{in(\alpha-\alpha_o)})^*]. \quad (\text{A14})
\end{aligned}$$

Thus, we see that the terms appear with an amplitude factor b^{3-n} . For the (2, 2) harmonic, we restricted to the two most significant precession contributions (b^0 and b^1). We will do the same for the (3, 3) multipole. For the leading-order precession contribution b^0 , we can evaluate the $d_{3,\pm 2}^3(\theta_{JN})$ terms Eq. (A13). For the first-order precession correction, we need to evaluate $d_{2,\pm 2}^3$ which are given by

$$\begin{aligned}
d_{2,2}^3(\theta) &= \frac{1 - 5\tau^2}{(1+\tau^2)^3}, \\
d_{2,-2}^3(\theta) &= \frac{5\tau^4 - \tau^6}{(1+\tau^2)^3}, \quad (\text{A15})
\end{aligned}$$

where $\tau = \tan(\theta/2)$.

We introduce the precession contributions to the (3, 3) multipole as¹²

¹²We have chosen the normalization of the (3, 3) waveform to match the one used in [69]. This differs from the more ‘‘natural’’ normalization where the prefactor would be $\sqrt{7/4\pi}$. This choice of normalization has no impact on the results in the paper.

$$h_{33,k} = \frac{1}{4} \sqrt{\frac{21}{2\pi}} \frac{b^k}{(1+b^2)^3} (h_{33}^{\text{NP}} e^{-3i\epsilon} e^{(3-k)i(\alpha-\alpha_o)}). \quad (\text{A16})$$

We can then express the (3, 3) multipole, to first order in b as

$$h_{33} \approx \frac{4\tau(h_{33,0} + \tau^4 h_{33,0}^*)}{(1+\tau^2)^3} + \frac{-4\tau(1-5\tau^2)h_{33,1} + 4\tau(\tau^2-5)h_{33,1}^*}{(1+\tau^2)^3}. \quad (\text{A17})$$

We are also interested in isolating the $h_{33,0}$ and $h_{33,1}$ waveform harmonics. To do so, we follow [70] and generate a waveform at different orientations and combine them to produce a waveform containing only the leading-precession harmonic for the (3, 3) multipole. For the (2, 2) multipole, only the 0 and 4 precession harmonics are nonzero for face-on systems so the simplest way to generate the leading-precession harmonic is simply to generate a face-on system—this will be correct up to $O(b^4)$, and even that contribution can be easily removed by subtracting two waveforms with different phase and polarization (see [70] for details). For the (3, 3) multipole, the leading-precession harmonic vanishes for a face-on system. Indeed, it is often stated that the (3, 3) multipole vanishes for face-on systems, actually this is not true if the system is precessing as the b^1 and b^5 precession harmonics are nonzero. Thus, to $O(b^4)$ the face on (3, 3) multipole gives the b^1 precession harmonic. Specifically,

$$h_{33,1} \approx h_{33}(\theta_{JN} = 0, \alpha_o = 0, \phi = 0, \psi = 0). \quad (\text{A18})$$

This will be correct to $O(b^4)$. The contribution from the $k = 5$ harmonic can be removed by taking the average of the above waveform and one generated with $\alpha_o = \frac{\pi}{4}$, $\phi = -\frac{\pi}{6}$.

To generate the leading-precession harmonic for the (3, 3) multipole, we require the waveform for $\theta_{JN} = \frac{\pi}{2}$ ($\tau = 1$). In this case, the amplitude of the left and right circular polarizations, given by the τ and τ^5 terms in Eq. (A17), are equal and the signal is linearly polarized, with power only in the + polarization. The same is true for the b^2 , b^4 , and b^6 harmonics, while the b^1 and b^5 harmonics have power only in the \times polarization and the b^3 harmonic vanishes. Thus, the + polarized waveform at $\theta_{JN} = \frac{\pi}{2}$ provides the leading-precession harmonic, accurate to b^2 . We can further improve the accuracy by noting that the precession phase α_o appears with a different factor for the different modes. In particular, generating the waveform at $\alpha_o = 0, \frac{\pi}{2}$ leads to the a phase shift of π in the leading harmonic and 2π in the b^2 harmonic. Therefore, to accuracy of b^4 we can write

$$h_{33,0} \approx \frac{1}{2} \left[h_{33+} \left(\theta_{JN} = \frac{\pi}{2}, \alpha_o = 0, \phi = 0 \right) + h_{33+} \left(\theta_{JN} = \frac{\pi}{2}, \alpha_o = \frac{\pi}{2}, \phi = \frac{\pi}{6} \right) \right]. \quad (\text{A19})$$

Again, it is not difficult to obtain the waveform accurate to all powers in b by generating waveforms with $\alpha_o = \pm \frac{\pi}{4}$ and appropriately combining them to remove the b^4 contribution.

3. The waveform for $\ell = m$ multipoles

Here, we briefly sketch the derivation for a generic (ℓ, ℓ) multipole. The calculation is essentially identical to the one above for the (3, 3) multipole. In particular, we first evaluate the $d_{\ell,n}^{\ell}(-\beta)$ term using the explicit form of the Wigner d -matrices

$$\begin{aligned} d_{\ell,n}^{\ell}(-\beta) &:= C_{\ell,n} \cos^{\ell+n}(\beta/2) \sin^{\ell-n}(-\beta/2) \\ &= C_{\ell,n} \frac{b^{\ell-n}}{(1+b^2)^{\ell}}, \end{aligned} \quad (\text{A20})$$

where $C_{\ell,n}$ is a real number (we give the value where necessary). Thus we can write

$$\begin{aligned} h_{\ell\ell} &= \sum_{n=-\ell}^{\ell} \sqrt{\frac{2\ell+1}{4\pi}} C_{\ell,n} \frac{(-b)^{\ell-n}}{(1+b^2)^{\ell}} \\ &\times [d_{n,2}^{\ell}(\theta_{JN})(h_{\ell\ell}^{\text{NP}} e^{-i\ell\epsilon} e^{in(\alpha-\alpha_o)}) \\ &+ d_{n,-2}^{\ell}(\theta_{JN})(h_{\ell\ell}^{\text{NP}} e^{-i\ell\epsilon} e^{in(\alpha-\alpha_o)})^*]. \end{aligned} \quad (\text{A21})$$

As before, we immediately see that the (ℓ, ℓ) multipole has the same decomposition in terms of b , with a clearly identified leading-order term when $n = \ell$. Restricting to this leading-precession contribution we note that $C_{\ell,\pm\ell} = 1$. In addition, we can evaluate the $d_{2,\pm\ell}^{\ell}(\theta_{JN}) = d_{\ell,\pm 2}^{\ell}(\theta_{JN})$ term using Eq. (A20), where we note the values $C_{22} = 1$, $C_{32} = \sqrt{6}$, $C_{42} = 2\sqrt{7}$.

Then, to leading order in the precession parameter b ,

$$h_{\ell\ell} \approx \frac{2(2\tau)^{\ell-2} (h_{\ell\ell,0} + \tau^4 h_{\ell\ell,0}^*)}{(1+\tau^2)^{\ell}} \quad (\text{A22})$$

with the normalization again chosen to match that introduced in [69] and

$$h_{\ell\ell,0} \propto \frac{1}{(1+b^2)^{\ell}} (h_{\ell\ell}^{\text{NP}} e^{-i\ell\epsilon} e^{i\ell(\alpha-\alpha_o)}). \quad (\text{A23})$$

The proportionality constant is determined through the desired normalization for $h_{\ell\ell}$.

As with the (2, 2) multipole, if we were to keep subdominant precession harmonics, their amplitudes would

be suppressed by powers of b while their frequencies would be reduced by multiples of the precession frequency.

4. The waveform for $m = \ell - 1$ multipoles

The other multipoles which can contribute significantly to the observed gravitational wave signal are the (2, 1) and

(3, 2) multipoles [42], so we also look briefly at the multipoles with $m = \ell - 1$. Now, the Wigner d -matrix can be written as

$$d_{\ell-1,n}^{\ell}(-\beta) = \frac{A_{\ell,n}(-b)^{\ell-n+1} + B_{\ell,n}(-b)^{\ell-n-1}}{(1+b^2)^{\ell}},$$




FIG. 16. The posterior distribution for masses (chirp mass \mathcal{M} and mass ratio $q = m_2/m_1$) spins (aligned χ_{align} and precessing χ_p), distance d_L and binary orientation θ_{JN} for a binary with masses $39M_{\odot}$ and $32M_{\odot}$ at a SNR of 12. The shaded regions in the two-dimensional figures show the 1, 2, and 3 σ regions for the parameters and the one-dimensional distributions show the median and 90% symmetric confidence interval. The simulated values are shown as red lines. At this low SNR, we observe broad posteriors for many parameters, in particular orientation and precessing spin.

where $A_{\ell n}$ and $B_{\ell n}$ are real numbers and, since the power of b must be non-negative, $A_{\ell, -\ell} = B_{\ell, \ell} = 0$.

As for the other harmonics, we see that the form of the $h_{\ell, (\ell-1)}$ waveform can be factorized into terms with a prefactor of b^k . Since the (3, 2) and (2, 1) multipoles are subdominant [69], we restrict to the leading order, b^0 , contributions. These arise from the $|n| = m = \ell - 1$ terms, for which $A_{\ell, -(\ell-1)} = B_{\ell, (\ell-1)} = 1$. Therefore, we obtain

$$h_{\ell(\ell-1)} = \sqrt{\frac{2\ell+1}{4\pi}} \frac{1}{(1+b^2)^\ell} \quad (\text{A24})$$

$$\begin{aligned} & [d_{\ell-1,2}^\ell(\theta_{JN})(h_{\ell(\ell-1)}^{\text{NP}} e^{i(\ell-1)(\alpha-\alpha_o-\epsilon)}) \\ & - d_{\ell-1,-2}^\ell(\theta_{JN})(h_{\ell(\ell-1)}^{\text{NP}} e^{i(\ell-1)(\alpha-\alpha_o-\epsilon)})^*]. \quad (\text{A25}) \end{aligned}$$

Finally, we note that for the $(\ell, \ell-1)$ multipoles the subdominant precession term would include two additional harmonics—one at a frequency Ω_p above the dominant harmonic and one a frequency Ω_p below.

APPENDIX B: SIMPLE-PE EXAMPLE FOR LOW SNR EVENT

In this appendix, we show the results for a `simple-pe` analysis of a low SNR event with close to equal masses. As discussed in Sec. V, such a system will have limited SNR in higher-order multipoles, precession and the second GW polarization for the majority of possible orientations and in-plane spin configurations. Therefore, we expect that the lack of SNR in these features will do little to restrict the inferred properties of the binary, most notably the orientation and in-plane spins. This is borne out in the results shown in Fig. 16. The binary is simulated with masses $39M_\odot$ and $32M_\odot$ (in the detector frame), giving a redshift of $z = 0.33$ and a network SNR of 12. There is minimal power in precession, higher-order modes and the second polarization with all features having $\text{SNR} < 0.5$. This leads to broad distributions in χ_p , mass ratio q and orientation θ_{JN} . Furthermore, there is only peak preference for a left (rather than right) circular polarization, which leads to the observed bimodality in orientation. While a detailed investigation of `simple-pe` results for GW observations is beyond the scope of this paper, we note that the widths of parameter posteriors are broadly consistent with observed signals [4].

-
- [1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
 - [2] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt *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] R. Abbott *et al.* (LIGO Scientific and VIRGO Collaborations), GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, [arXiv:2108.01045](https://arxiv.org/abs/2108.01045).
 - [4] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, [arXiv:2111.03606](https://arxiv.org/abs/2111.03606).
 - [5] 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).
 - [6] S. Olsen, T. Venumadhav, J. Mushkin, J. Roulet, B. Zackay, and M. Zaldarriaga, New binary black hole mergers in the LIGO-Virgo O3a data, *Phys. Rev. D* **106**, 043009 (2022).
 - [7] A. H. Nitz, S. Kumar, Y.-F. Wang, S. Kastha, S. Wu, M. Schäfer, R. Dhurkunde, and C. D. Capano, 4-OGC: Catalog of gravitational waves from compact-binary mergers, *Astrophys. J.* **946**, 59 (2023).
 - [8] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
 - [9] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin *et al.*, Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2014).
 - [10] B. P. Abbott *et al.* (KAGRA, LIGO Scientific, and Virgo Collaborations), Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, *Living Rev. Relativity* **23**, 3 (2020).
 - [11] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW190814: Gravitational waves from the coalescence of a 23 solar mass black hole with a 2.6 solar mass compact object, *Astrophys. J. Lett.* **896**, L44 (2020).
 - [12] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW190412: Observation of a binary-black-hole coalescence with asymmetric masses, *Phys. Rev. D* **102**, 043015 (2020).
 - [13] I. M. Romero-Shaw, P. D. Lasky, E. Thrane, and J. C. Bustillo, GW190521: Orbital eccentricity and signatures of dynamical formation in a binary black hole merger signal, *Astrophys. J. Lett.* **903**, L5 (2020).

- [14] 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).
- [15] C. M. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, PyCBC Inference: A Python-based parameter estimation toolkit for compact binary coalescence signals, *Publ. Astron. Soc. Pac.* **131**, 024503 (2019).
- [16] G. Ashton *et al.*, Bilby: A user-friendly Bayesian inference library for gravitational-wave astronomy, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [17] J. Lange, R. O’Shaughnessy, and M. Rizzo, Rapid and accurate parameter inference for coalescing, precessing compact binaries, [arXiv:1805.10457](https://arxiv.org/abs/1805.10457).
- [18] I. M. Romero-Shaw *et al.*, Bayesian inference for compact binary coalescences with Bilby: Validation and application to the first LIGO–Virgo gravitational-wave transient catalogue, *Mon. Not. R. Astron. Soc.* **499**, 3295 (2020).
- [19] G. Ashton and C. Talbot, Bilby-MCMC: An MCMC sampler for gravitational-wave inference, *Mon. Not. R. Astron. Soc.* **507**, 2037 (2021).
- [20] R. J. E. Smith, G. Ashton, A. Vajpeyi, and C. Talbot, Massively parallel Bayesian inference for transient gravitational-wave astronomy, *Mon. Not. R. Astron. Soc.* **498**, 4492 (2020).
- [21] H. Gabbard, C. Messenger, I. S. Heng, F. Tonolini, and R. Murray-Smith, Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy, *Nat. Phys.* **18**, 112 (2022).
- [22] S. R. Green, C. Simpson, and J. Gair, Gravitational-wave parameter estimation with autoregressive neural network flows, *Phys. Rev. D* **102**, 104057 (2020).
- [23] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, Real-Time Gravitational Wave Science with Neural Posterior Estimation, *Phys. Rev. Lett.* **127**, 241103 (2021).
- [24] M. Dax, S. R. Green, J. Gair, M. Pürrer, J. Wildberger, J. H. Macke, A. Buonanno, and B. Schölkopf, Neural Importance Sampling for Rapid and Reliable Gravitational-Wave Inference, *Phys. Rev. Lett.* **130**, 171403 (2023).
- [25] H. Shen, E. A. Huerta, E. O’Shea, P. Kumar, and Z. Zhao, Statistically-informed deep learning for gravitational wave parameter estimation, *Mach. Learn. Sci. Technol.* **3**, 015007 (2022).
- [26] S. R. Green and J. Gair, Complete parameter inference for GW150914 using deep learning, *Mach. Learn. Sci. Technol.* **2**, 03LT01 (2021).
- [27] M. J. Williams, J. Veitch, and C. Messenger, Nested sampling with normalizing flows for gravitational-wave inference, *Phys. Rev. D* **103**, 103006 (2021).
- [28] B. Zackay, L. Dai, and T. Venumadhav, Relative binning and fast likelihood evaluation for gravitational wave parameter estimation, [arXiv:1806.08792](https://arxiv.org/abs/1806.08792).
- [29] L. Pathak, A. Reza, and A. S. Sengupta, Rapid reconstruction of compact binary sources using meshfree approximation, [arXiv:2210.02706](https://arxiv.org/abs/2210.02706).
- [30] T. Islam, J. Roulet, and T. Venumadhav, Factorized parameter estimation for real-time gravitational wave inference, [arXiv:2210.16278](https://arxiv.org/abs/2210.16278).
- [31] A. Bohé *et al.*, Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors, *Phys. Rev. D* **95**, 044028 (2017).
- [32] R. Cotesta, A. Buonanno, A. Bohé, A. Taracchini, I. Hinder, and S. Ossokine, Enriching the symphony of gravitational waves from binary black holes by tuning higher harmonics, *Phys. Rev. D* **98**, 084028 (2018).
- [33] R. Cotesta, S. Marsat, and M. Pürrer, Frequency domain reduced order model of aligned-spin effective-one-body waveforms with higher-order modes, *Phys. Rev. D* **101**, 124040 (2020).
- [34] S. Ossokine *et al.*, Multipolar effective-one-body waveforms for precessing binary black holes: Construction and validation, *Phys. Rev. D* **102**, 044055 (2020).
- [35] S. Babak, A. Taracchini, and A. Buonanno, Validating the effective-one-body model of spinning, precessing binary black holes against numerical relativity, *Phys. Rev. D* **95**, 024010 (2017).
- [36] 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).
- [37] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from non-precessing black-hole binaries. I. New numerical waveforms and anatomy of the signal, *Phys. Rev. D* **93**, 044006 (2016).
- [38] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jiménez 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).
- [39] L. London, S. Khan, E. Fauchon-Jones, C. García, M. Hannam, S. Husa, X. Jiménez-Forteza, C. Kalaghatgi, F. Ohme, and F. Pannarale, First Higher-Multipole Model of Gravitational Waves from Spinning and Coalescing Black-Hole Binaries, *Phys. Rev. Lett.* **120**, 161102 (2018).
- [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] S. Khan, K. Chatziioannou, M. Hannam, and F. Ohme, Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects, *Phys. Rev. D* **100**, 024059 (2019).
- [42] S. Khan, F. Ohme, K. Chatziioannou, and M. Hannam, Including higher order multipoles in gravitational-wave models for precessing binary black holes, *Phys. Rev. D* **101**, 024056 (2020).
- [43] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, Surrogate models for precessing binary black hole simulations with unequal masses, *Phys. Rev. Res.* **1**, 033015 (2019).
- [44] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, Surrogate model of hybridized numerical relativity binary black hole waveforms, *Phys. Rev. D* **99**, 064045 (2019).

- [45] G. Pratten, S. Husa, C. García-Quiros, M. Colleoni, A. Ramos-Buades, H. Estelles, and R. Jaume, Setting the cornerstone for a family of models for gravitational waves from compact binaries: The dominant harmonic for non-precessing quasicircular black holes, *Phys. Rev. D* **102**, 064001 (2020).
- [46] C. García-Quirós, M. Colleoni, S. Husa, H. Estellés, G. Pratten, A. Ramos-Buades, M. Mateu-Lucena, and R. Jaume, Multimode frequency-domain model for the gravitational wave signal from nonprecessing black-hole binaries, *Phys. Rev. D* **102**, 064002 (2020).
- [47] G. Pratten *et al.*, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, *Phys. Rev. D* **103**, 104056 (2021).
- [48] H. Estellés, A. Ramos-Buades, S. Husa, C. García-Quirós, M. Colleoni, L. Haegel, and R. Jaume, Phenomenological time domain model for dominant quadrupole gravitational wave signal of coalescing binary black holes, *Phys. Rev. D* **103**, 124060 (2021).
- [49] H. Estellés, S. Husa, M. Colleoni, D. Keitel, M. Mateu-Lucena, C. García-Quirós, A. Ramos-Buades, and A. Borchers, Time-domain phenomenological model of gravitational-wave subdominant harmonics for quasicircular nonprecessing binary black hole coalescences, *Phys. Rev. D* **105**, 084039 (2022).
- [50] H. Estellés, M. Colleoni, C. García-Quirós, S. Husa, D. Keitel, M. Mateu-Lucena, M. d. L. Planas, and A. Ramos-Buades, New twists in compact binary waveform modeling: A fast time domain model for precession, *Phys. Rev. D* **105**, 084040 (2022).
- [51] E. Hamilton, L. London, J.E. Thompson, E. Fauchon-Jones, M. Hannam, C. Kalaghatgi, S. Khan, F. Pannarale, and A. Vano-Vinuales, Model of gravitational waves from precessing black-hole binaries through merger and ring-down, *Phys. Rev. D* **104**, 124027 (2021).
- [52] C. Devine, Z. B. Etienne, and S. T. McWilliams, Optimizing spinning time-domain gravitational waveforms for Advanced LIGO data analysis, *Classical Quantum Gravity* **33**, 125025 (2016).
- [53] T. D. Knowles, C. Devine, D. A. Buch, S. A. Bilgili, T. R. Adams, Z. B. Etienne, and S. T. McWilliams, Improving performance of SEOBNRv3 by $\sim 300\times$, *Classical Quantum Gravity* **35**, 155003 (2018).
- [54] B. Gadre, M. Pürrer, S. E. Field, S. Ossokine, and V. Varma, A fully precessing higher-mode surrogate model of effective-one-body waveforms, [arXiv:2203.00381](https://arxiv.org/abs/2203.00381).
- [55] S. Khan and R. Green, Gravitational-wave surrogate models powered by artificial neural networks, *Phys. Rev. D* **103**, 064015 (2021).
- [56] L. M. Thomas, G. Pratten, and P. Schmidt, Accelerating multimodal gravitational waveforms from precessing compact binaries with artificial neural networks, *Phys. Rev. D* **106**, 104029 (2022).
- [57] C. Cutler and E. E. Flanagan, Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform?, *Phys. Rev. D* **49**, 2658 (1994).
- [58] E. Poisson and C. M. Will, Gravitational waves from inspiraling compact binaries: Parameter estimation using second post-Newtonian wave forms, *Phys. Rev. D* **52**, 848 (1995).
- [59] 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).
- [60] B. Farr *et al.*, Parameter estimation on gravitational waves from neutron-star binaries with spinning components, *Astrophys. J.* **825**, 116 (2016).
- [61] 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).
- [62] S. Fairhurst, Source localization with an advanced gravitational wave detector network, *Classical Quantum Gravity* **28**, 105021 (2011).
- [63] L. Wen and Y. Chen, Geometrical expression for the angular resolution of a network of gravitational-wave detectors, *Phys. Rev. D* **81**, 082001 (2010).
- [64] C. L. Rodriguez, B. Farr, W. M. Farr, and I. Mandel, Inadequacies of the Fisher information matrix in gravitational-wave parameter estimation, *Phys. Rev. D* **88**, 084013 (2013).
- [65] W. Zhao and L. Wen, Localization accuracy of compact binary coalescences detected by the third-generation gravitational-wave detectors and implication for cosmology, *Phys. Rev. D* **97**, 064031 (2018).
- [66] 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).
- [67] K. Grover, S. Fairhurst, B. F. Farr, I. Mandel, C. Rodriguez, T. Sidery, and A. Vecchio, Comparison of gravitational wave detector network sky localization approximations, *Phys. Rev. D* **89**, 042004 (2014).
- [68] C. Kalaghatgi, M. Hannam, and V. Raymond, Parameter estimation with a spinning multimode waveform model, *Phys. Rev. D* **101**, 103004 (2020).
- [69] C. Mills and S. Fairhurst, Measuring gravitational-wave higher-order multipoles, *Phys. Rev. D* **103**, 024042 (2021).
- [70] 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).
- [71] S. Fairhurst, R. Green, M. Hannam, and C. Hoy, When will we observe binary black holes precessing?, *Phys. Rev. D* **102**, 041302 (2020).
- [72] R. Green, C. Hoy, S. Fairhurst, M. Hannam, F. Pannarale, and C. Thomas, Identifying when precession can be measured in gravitational waveforms, *Phys. Rev. D* **103**, 124023 (2021).
- [73] G. Pratten, P. Schmidt, R. Buscicchio, and L. M. Thomas, Measuring precession in asymmetric compact binaries, *Phys. Rev. Res.* **2**, 043096 (2020).
- [74] N. V. Krishnendu and F. Ohme, Interplay of spin-precession and higher harmonics in the parameter estimation of binary black holes, *Phys. Rev. D* **105**, 064012 (2022).
- [75] M. L. Chan, C. Messenger, I. S. Heng, and M. Hendry, Binary neutron star mergers and third generation detectors: Localization and early warning, *Phys. Rev. D* **97**, 123014 (2018).

- [76] K. S. Thorne, Multipole expansions of gravitational radiation, *Rev. Mod. Phys.* **52**, 299 (1980).
- [77] L. E. Kidder, Using full information when computing modes of post-Newtonian waveforms from inspiralling compact binaries in circular orbit, *Phys. Rev. D* **77**, 044016 (2008).
- [78] D. Brown, S. Fairhurst, B. Krishnan, R. A. Mercer, R. K. Kopparapu, L. Santamaria, and J. T. Whelan, Data formats for numerical relativity waves, [arXiv:0709.0093](https://arxiv.org/abs/0709.0093).
- [79] 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).
- [80] A. Lundgren and R. O’Shaughnessy, Single-spin precessing gravitational waveform in closed form, *Phys. Rev. D* **89**, 044021 (2014).
- [81] M. Boyle, R. Owen, and H. P. Pfeiffer, A geometric approach to the precession of compact binaries, *Phys. Rev. D* **84**, 124011 (2011).
- [82] C. McIsaac, C. Hoy, and I. Harry, A search technique to observe precessing compact binary mergers in the advanced detector era, [arXiv:2303.17364](https://arxiv.org/abs/2303.17364).
- [83] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries, *Phys. Rev. D* **85**, 122006 (2012).
- [84] L. Lindblom, B. J. Owen, and D. A. Brown, Model waveform accuracy standards for gravitational wave data analysis, *Phys. Rev. D* **78**, 124020 (2008).
- [85] LIGO Scientific and Virgo Collaborations, Noise curves for use in simulations pre-O4, <https://dcc.ligo.org/LIGO-T2200043/public> (2022).
- [86] B. F. Schutz, Networks of gravitational wave detectors and three figures of merit, *Classical Quantum Gravity* **28**, 125023 (2011).
- [87] S. Klimenko, G. Vedovato, M. Drago, G. Mazzolo, G. Mitselmakher, C. Pankow, G. Prodi, V. Re, F. Salemi, and I. Yakushin, Localization of gravitational wave sources with networks of advanced detectors, *Phys. Rev. D* **83**, 102001 (2011).
- [88] I. W. Harry and S. Fairhurst, A targeted coherent search for gravitational waves from compact binary coalescences, *Phys. Rev. D* **83**, 084002 (2011).
- [89] S. A. Usman, J. C. Mills, and S. Fairhurst, Constraining the inclinations of binary mergers from gravitational-wave observations, *Astrophys. J.* **877**, 82 (2019).
- [90] V. Tiwari, C. Hoy, S. Fairhurst, and D. MacLeod, VARAHA: A fast non-Markovian sampler for estimating gravitational-wave posteriors, *Phys. Rev. D* **108**, 023001 (2023).
- [91] 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).
- [92] C. Hanna *et al.*, Fast evaluation of multidetector consistency for real-time gravitational wave searches, *Phys. Rev. D* **101**, 022003 (2020).
- [93] L. Blanchet, Gravitational radiation from post-Newtonian sources and inspiralling compact binaries, *Living Rev. Relativity* **9**, 4 (2006).
- [94] J. Healy and C. O. Lousto, Hangup effect in unequal mass binary black hole mergers and further studies of their gravitational radiation and remnant properties, *Phys. Rev. D* **97**, 084002 (2018).
- [95] S. Fairhurst, Triangulation of gravitational wave sources with a network of detectors, *New J. Phys.* **11**, 123006 (2009).
- [96] L. P. Singer and L. R. Price, Rapid Bayesian position reconstruction for gravitational-wave transients, *Phys. Rev. D* **93**, 024013 (2016).
- [97] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, Detecting binary compact-object mergers with gravitational waves: Understanding and Improving the sensitivity of the PyCBC search, *Astrophys. J.* **849**, 118 (2017).
- [98] S. Fairhurst, Localization of transient gravitational wave sources: Beyond triangulation, *Classical Quantum Gravity* **35**, 105002 (2018).
- [99] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Phys. Rev. Lett.* **119**, 161101 (2017).
- [100] C. Hoy, S. Fairhurst, M. Hannam, and V. Tiwari, Understanding how fast black holes spin by analyzing data from the second gravitational-wave catalogue, *Astrophys. J.* **928**, 75 (2022).
- [101] J. Roulet, S. Olsen, J. Mushkin, T. Islam, T. Venumadhav, B. Zackay, and M. Zaldarriaga, Removing degeneracy and multimodality in gravitational wave source parameters, *Phys. Rev. D* **106**, 123015 (2022).
- [102] S. Fairhurst, C. Hoy, R. Green, C. Mills, and S. A. Usman, <https://stephen-fairhurst.docs.ligo.org/simple-pe/>, MIT license (2023).
- [103] S. Fairhurst, C. Hoy, R. Green, C. Mills, and S. A. Usman, <https://git.ligo.org/stephen-fairhurst/simple-pe>, MIT license (2023).
- [104] P. Virtanen *et al.* (SciPy 1.0 Contributors), SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [105] T. Dal Canton, A. H. Nitz, B. Gadre, G. S. Cabourn Davies, V. Villa-Ortega, T. Dent, I. Harry, and L. Xiao, Real-time search for compact binary mergers in Advanced LIGO and Virgo’s third observing run using PyCBC live, *Astrophys. J.* **923**, 254 (2021).
- [106] F. Ohme, A. B. Nielsen, D. Keppel, and A. Lundgren, Statistical and systematic errors for gravitational-wave inspiral signals: A principal component analysis, *Phys. Rev. D* **88**, 042002 (2013).
- [107] B. Farr, E. Ochsner, W. M. Farr, and R. O’Shaughnessy, A more effective coordinate system for parameter estimation of precessing compact binaries from gravitational waves, *Phys. Rev. D* **90**, 024018 (2014).
- [108] B. J. Owen, Search templates for gravitational waves from inspiraling binaries: Choice of template spacing, *Phys. Rev. D* **53**, 6749 (1996).
- [109] B. J. Owen and B. S. Sathyaprakash, Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement, *Phys. Rev. D* **60**, 022002 (1999).
- [110] H.-S. Cho, E. Ochsner, R. O’Shaughnessy, C. Kim, and C.-H. Lee, Gravitational waves from black hole-neutron

- star binaries: Effective Fisher matrices and parameter estimation using higher harmonics, *Phys. Rev. D* **87**, 024004 (2013).
- [111] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), The Population of Merging Compact Binaries Inferred Using Gravitational Waves Through GWTC-3, *Phys. Rev. X* **13**, 011048 (2023).
- [112] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run, *Phys. Rev. X* **11**, 021053 (2021).
- [113] LIGO Scientific Collaboration and Virgo Collaboration, GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run—parameter estimation data release, Zenodo, [10.5281/zenodo.5117703](https://zenodo.org/record/5117703) (2021).
- [114] T. Callister, A thesaurus for common priors in gravitational-wave astronomy, [arXiv:2104.09508](https://arxiv.org/abs/2104.09508).
- [115] J. S. Speagle, DYNESTY: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [116] S. Vitale, C.-J. Haster, L. Sun, B. Farr, E. Goetz, J. Kissel, and C. Cahillane, Physical approach to the marginalization of LIGO calibration uncertainties, *Phys. Rev. D* **103**, 063016 (2021).
- [117] B. D. Metzger, Kilonovae, *Living Rev. Relativity* **20**, 3 (2017).
- [118] R. Abbott *et al.* (LIGO Scientific and VIRGO Collaborations), Search for lensing signatures in the gravitational-wave observations from the first half of LIGO–Virgo’s third observing run, *Astrophys. J.* **923**, 14 (2021).
- [119] E. Payne, C. Talbot, and E. Thrane, Higher order gravitational-wave modes with likelihood reweighting, *Phys. Rev. D* **100**, 123017 (2019).
- [120] S. De, D. Finstad, J. M. Lattimer, D. A. Brown, E. Berger, and C. M. Biwer, Tidal Deformabilities and Radii of Neutron Stars from the Observation of GW170817, *Phys. Rev. Lett.* **121**, 091102 (2018); **121**, 259902(E) (2018).
- [121] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW170817: Measurements of Neutron Star Radii and Equation of State, *Phys. Rev. Lett.* **121**, 161101 (2018).
- [122] P. C. Peters, Gravitational radiation and the motion of two point masses, *Phys. Rev.* **136**, B1224 (1964).
- [123] V. Gayathri, J. Healy, J. Lange, B. O’Brien, M. Szczepańczyk, I. Bartos, M. Campanelli, S. Klimentko, C. O. Lousto, and R. O’Shaughnessy, Eccentricity estimate for black hole mergers with numerical relativity simulations, *Nat. Astron.* **6**, 344 (2022).
- [124] M. Boyle, L. E. Kidder, S. Ossokine, and H. P. Pfeiffer, Gravitational-wave modes from precessing black-hole binaries, [arXiv:1409.4431](https://arxiv.org/abs/1409.4431).
- [125] A. Ramos-Buades, P. Schmidt, G. Pratten, and S. Husa, Validity of common modeling approximations for precessing binary black holes with higher-order modes, *Phys. Rev. D* **101**, 103014 (2020).
- [126] G. Pratten, P. Schmidt, and N. Williams, Impact of Dynamical Tides on the Reconstruction of the Neutron Star Equation of State, *Phys. Rev. Lett.* **129**, 081102 (2022).
- [127] D. Finstad, L. V. White, and D. A. Brown, Prospects for a precise equation of state measurement from Advanced LIGO and Cosmic Explorer, [arXiv:2211.01396](https://arxiv.org/abs/2211.01396).
- [128] A. M. Knee, I. M. Romero-Shaw, P. D. Lasky, J. McIver, and E. Thrane, A Rosetta Stone for eccentric gravitational waveform models, *Astrophys. J.* **936**, 172 (2022).
- [129] P. Kolitsidou, J. Thompson, and M. Hannam, Systematic biases in binary-black-hole parameter measurements due to the neglect of antisymmetric multipoles (to be published).
- [130] T. Damour, A. Gopakumar, and B. R. Iyer, Phasing of gravitational waves from inspiralling eccentric binaries, *Phys. Rev. D* **70**, 064028 (2004).
- [131] C. R. Harris *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).
- [132] A. Nitz, I. Harry, D. Brown, C. M. Biwer, J. Willis, T. D. Canton, C. Capano, L. Pekowsky, T. Dent, A. R. Williamson *et al.*, gwastro/pycbc: pyCBC release v1.15.2, Zenodo, [10.5281/zenodo.3596447](https://zenodo.org/record/3596447) (2019).
- [133] C. Hoy and V. Raymond, PESummary: The code agnostic parameter estimation summary page builder, *SoftwareX* **15**, 100765 (2021).
- [134] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).