

Reconstructing phenomenological distributions of compact binaries via gravitational wave observations

Daniel Wysocki,^{*} Jacob Lange, and Richard O’Shaughnessy
Rochester Institute of Technology, Rochester, New York 14623, USA

 (Received 18 May 2018; published 13 August 2019)

Gravitational wave (GW) measurements will provide insight into the population of coalescing compact binaries throughout the universe. We describe and demonstrate a flexible parametric method to infer the event rate as a function of compact binary parameters, accounting for Poisson error and selection biases. Using synthetic data based on projections for LIGO and Virgo’s third observing run (O3), we discuss how well GW measurements could constrain the mass and spin distribution of coalescing neutron stars and black holes (BHs) in the near future, within the context of several phenomenological models described in this work. We demonstrate that only a few tens of events can enable astrophysically significant constraints on the spin magnitude and orientation distribution of BHs in merging binaries. We discuss how astrophysical priors or other measurements can inform the interpretation of future measurements. Using publicly available results, we estimate the event rate versus mass for binary black holes (BBHs). To connect to previously published work, we provide estimates including reported O2 BBH candidates, making several unwarranted but simplifying assumptions for the sensitivity of the network and completeness of the reported set of events. Consistent with prior work, we find BHs in binaries likely have low natal spin. With available results and a population favoring low spin, we cannot presently constrain the typical misalignments of the binary black hole population. All of the tools described in this work are publicly available and ready-to-use to interpret real or synthetic LIGO data, and to synthesize projected data from future observing runs.¹

DOI: [10.1103/PhysRevD.100.043012](https://doi.org/10.1103/PhysRevD.100.043012)

I. INTRODUCTION

The Advanced Laser Interferometer Gravitational Wave Observatory (LIGO) [1] and Virgo [2,3] detectors have and will continue to discover gravitational waves (GW) from coalescing binary black holes (BBHs) and neutron stars. Several tens of binary black holes and potentially neutron stars are expected to be seen in O3, LIGO’s next observing run, alone; and several hundreds more detections are expected over the next five years [4,5]. Already, the properties of the sources responsible—the inferred event rates, masses, and spins—have confronted other observations of black holes’ masses and spins [5], challenged previous formation scenarios [5,6], and inspired new models [7–10] and insights [11,12] into the evolution of massive stars and the observationally accessible gravitational waves they emit [13,14]. Over the next several years, our understanding of the lives and deaths of massive stars over cosmic time will be transformed by the identification and interpretation of the population(s) responsible for coalescing binaries

[6,15,16], because measurements will enable robust tests to distinguish between formation scenarios [17] with present [18] and future instruments [19,20].

During the first few years of discovery, substantial theoretical modeling challenges and the rapid pace of events suggest that GW observations could soon outpace theory. In this work, we introduce a flexible, concrete, and production-ready approach to infer a compact binary merger rate and compact binary distribution, in the context of an (arbitrary) parametrized phenomenological model. We extend or employ previously proposed models [21,22]. We are motivated by how constraints on these phenomenological models enable us to address broad astrophysical questions—the mass and spin distribution of neutron stars and black holes, as imparted at their birth; the dominant formation mechanism for compact binaries, such as the role of dynamical versus isolated formation channels for binary black holes. To that end, we provide concrete demonstrations of how a few GW measurements will provide insights that enable sharp discrimination between proposed astrophysical alternatives, or measurements of their parameters. We use simple phenomenological arguments and calculations to characterize the information that these first few hundred observations should provide. Conversely, we provide simple approaches to extend our

^{*}dw2081@rit.edu.

¹<https://git.ligo.org/daniel.wysocki/bayesian-parametric-population-models/> and <https://git.ligo.org/daniel.wysocki/synthetic-PE-posteriors>

phenomenological approach in sophistication and complexity as several thousand compact binary mergers provide sharp constraints on their underlying properties. This approach complements inferences that work within a concrete model family as explored in other proof-of-concept investigations (see, e.g., [16,17,23–28] and references therein).

GW measurements probe only a selection-biased part of the compact binary distribution. Previously reported estimates of the overall compact binary event rate rely on extrapolation away from the observed population, using some fixed model for the compact binary mass distribution [5]. In fact, the compact binary mass distribution and inferred event rate are strongly coupled. This paper provides the first self-consistent approach to infer both the compact binary event rate and parameter distribution; then it describes and explains the expected correlation in an accessible way.

Several recent studies have explored how well GW measurements can constrain the mass and spin distribution of binary black holes [5,21,23,29–38]. Our approach is novel insofar as it reconstructs both the strongly correlated event rate and the parameter distribution, making our method a robust tool to assess astrophysical formation scenarios. In our modeling, we focus on measuring the black hole (BH) spin magnitude and misalignment distribution, as a method to probe the formation scenarios for binary BHs. As first described in [17], GW provide a unique opportunity to distinguish between isolated and dynamic formation mechanisms: measurements of the spin properties of the BHs [6,18,32,36,39,40]. The presence of a component of the BH spins in the plane of the orbit leads to precession of that plane. If suitably massive and significantly spinning, such binaries will strongly precess within the LIGO sensitive band. If BBHs are the end points of isolated binary star systems, they would be expected to contain BHs with spins preferentially aligned with the orbital angular momentum [40,41], and therefore rarely be strongly precessing. If, however, BBHs predominantly form as a result of gravitational interactions inside dense populations of stellar systems, the relative orientations of the BH spins with their orbits will be random, and some gravitational wave signals may be very strongly precessing. At this early stage, observations cannot firmly distinguish between these two scenarios, or more broadly other possible BBH formation mechanisms [6]. These include the evolution of isolated pairs of stars [7,8,25,42–44], dynamic binary formation in dense clusters [9], and pairs of primordial BHs [10]; see, e.g., [6] and references therein. Loosely speaking, however, the isolated evolution and globular cluster formation scenarios are the most well-developed and verifiable using independent observational constraints. More broadly, precise measurements of their properties will provide unique clues into how BHs and massive stars evolve [18–20,33,36,39,45].

This paper is organized as follows. In Sec. II we describe our techniques to infer compact binary populations,

building upon inferences about parameters of individual events. Unlike prior work, we simultaneously reconstruct the event rate, mass distribution, and spin (vector) distribution. In Sec. III, we demonstrate our population inference strategy with two examples. In the first, we perform a full end-to-end analysis of synthetic GW data generated from a synthetic population of astrophysically distributed sources. In the second, using a tool to mimic how well we could constrain parameters of a candidate GW signal, we perform a large-scale investigation into how well GW measurements could constrain the mass and spin distribution of binary black holes. We find that the mass and spin distribution can be tightly constrained with only a few tens of events. By virtue of explicitly exploiting only some of the available information, our estimates are necessarily conservative. In Sec. IV, we apply our method to the currently reported binary black hole population. For simplicity, assuming the reported events to date represent a fair sample of the results of LIGO's first two observing runs (O1 and O2), we corroborate previous results, finding black hole spins are likely small and that the black hole mass spectrum may have an upper bound. Due to small BH spins, except for GW151226, we can extract no information about typical BBH spin-orbit misalignments. We emphasize our demonstration uses a nonfinal sample for LIGO's O2 survey: depending on that survey's results, applying our methods to final O2 results could produce substantially different astrophysical conclusions. In Sec. V we briefly discuss the accuracy to which population parameters can be determined, and the surprisingly significant role of waveform systematics in the near future. After summarizing our conclusions in Sec. VI, we supply three Appendixes. In Appendix A, we describe a robust, extensible procedure for generating synthetic posterior distributions for proposed GW events. This open-source procedure could be widely used to assess the viability of GW measurements to distinguish between proposed astrophysical channels. A subsequent short Appendix B describes how to generate synthetic populations of selection-biased GW sources using this procedure. Next, in Appendix C, following on and extending previous work, we use toy models for both the measurement process and source population to illustrate how well GW observations will constrain the mass and spin distribution of compact binaries, likely providing robust insights into compact object formation (e.g., BH natal spins and maximum masses) and binary formation mechanisms (e.g., dynamical over isolated).

II. METHOD

A. Characterizing and inferring parameters of individual binary black holes

A coalescing compact binary in a quasicircular orbit can be completely characterized by its intrinsic parameters, namely its individual masses m_i and spins S_i , and its seven

extrinsic parameters: right ascension, declination, luminosity distance, coalescence time, and three Euler angles characterizing its orientation (e.g., inclination, orbital phase, and polarization). In this work, we will also use the total mass $M = m_1 + m_2$ and mass ratio q defined in the following way:

$$q = m_2/m_1, \quad \text{where } m_1 \geq m_2. \quad (1)$$

We will also refer to two other commonly used mass parametrizations: the chirp mass $\mathcal{M}_c = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$ and the symmetric mass ratio $\eta = m_1 m_2 / (m_1 + m_2)^2$. With regard to spin, we define an effective spin [46–48], which is a combination of the spin components along the orbital angular momentum direction \hat{L} , in the following way:

$$\chi_{\text{eff}} = (\mathbf{S}_1/m_1 + \mathbf{S}_2/m_2) \cdot \hat{L}/M, \quad (2)$$

where \mathbf{S}_1 and \mathbf{S}_2 are the spins on the individual BH. We will also characterize BH spins using the dimensionless spin variables

$$\chi_i = \mathbf{S}_i/m_i^2. \quad (3)$$

We will express these dimensionless spins in terms of Cartesian components $\chi_{i,x}, \chi_{i,y}, \chi_{i,z}$, expressed relative to a frame with $\hat{z} = \hat{L}$ and (for simplicity) at the orbital frequency corresponding to the earliest time of astrophysical interest (e.g., an orbital frequency of ≈ 10 Hz).

When necessary, compact binary parameters are inferred through the use of Bayesian analysis via Rapid parameter Inference on gravitational wave sources via Iterative Fitting (RIFT) [49], which reproduces the results of standard Monte Carlo techniques described in [50,51] and references therein. For any event, fully characterized by parameters x , we can compute the (Gaussian) likelihood function $p(d|x)$ for detector network data d containing a signal by using waveform models and an estimate of the (approximately Gaussian) detector noise on short timescales (see, e.g., [50–52] and references therein). In this expression x is shorthand for the set of 15 parameters needed to fully specify a quasicircular BBH. The posterior probability distribution is therefore $p(x|d) \propto p(d|x)p(x)$, where $p(x)$ is the prior probability of finding a merger with different masses, spins, and orientations somewhere in the universe. These parameters x can and are often described with alternate coordinate systems. We sometimes refer to the source luminosity distance d_L or equivalently its source redshift z , and to the detector-frame or redshifted masses $m_{i,z} = m_i(1+z)$. (To distinguish from the detector-frame masses, we will sometimes refer to m_i as the source-frame binary masses.) LIGO-Virgo analyses have adopted a fiducial prior $p_{\text{ref}}(x)$ that is uniform in orientation, in luminosity distance cubed, in redshifted mass, in spin direction (on the sphere), and, importantly for us, in spin magnitude [50,51]. Using

standard Bayesian tools [50,51], one can produce a sequence of independent, identically distributed samples $x_{n,s}$ ($s = 1, 2, \dots, S$) from the posterior distribution $p(x|d)$ for each event n ; that is, each $x_{n,s}$ is drawn from a distribution proportional to $p(d_n|x_n)p_{\text{ref}}(x_n)$. Typical calculations of this type provide $\lesssim 10^4$ samples [50,51] from which the posterior probability distribution is inferred.

For other examples involving purely synthetic observing scenarios, we perform this procedure with a familiar Fisher matrix approximation for the form of $p(d|x)$ as a function of x [53–55]; see Appendix A for details.

B. Population inference

We use Bayesian inference to constrain the mass and spin distributions of the astrophysical population of BBHs. To do this, we assume that the distribution is one of a family of distributions, parametrized by Λ and scaled by some overall rate $\mathcal{R} = dN/(dt dV_c)$, which is constant in comoving volume V_c . Each BBH in the population has properties denoted by $\lambda \equiv (m_1, m_2, \chi_1, \chi_2)$.

Ultimately we are interested in determining the likelihood of the astrophysical BBH population having a given merger rate \mathcal{R} and obeying a given parametrization Λ , given the data for N detections, $\mathcal{D} = (d_1, \dots, d_N)$. This likelihood, $\mathcal{L}(\mathcal{R}, \Lambda) \equiv p(\mathcal{D}|\mathcal{R}, \Lambda)$, is that of an inhomogeneous Poisson process

$$\mathcal{L}(\mathcal{R}, \Lambda) \propto e^{-\mu(\mathcal{R}, \Lambda)} \prod_{n=1}^N \int d\lambda \ell_n(\lambda) \mathcal{R} p(\lambda|\Lambda), \quad (4)$$

where $\mu(\mathcal{R}, \Lambda)$ is the expected number of detections under a given population parametrization Λ with overall rate \mathcal{R} and where $\ell_n(\lambda) = p(d_n|\lambda)$ is the likelihood of data d_n given binary parameters λ . A derivation for μ is given in Sec. II C.

Using Bayes' theorem, $p(\mathcal{R}, \Lambda|\mathcal{D}) \propto p(\mathcal{R}, \Lambda)\mathcal{L}(\mathcal{R}, \Lambda)$, one may obtain a posterior distribution on \mathcal{R} and Λ , after assuming some prior $p(\mathcal{R}, \Lambda)$. To avoid computing the normalization constant, we instead draw samples from the posterior distribution via Goodman and Weare's affine invariant Markov chain Monte Carlo (MCMC) ensemble sampler [56], as implemented in the PYTHON package EMCEE [57].

C. Estimate for VT

Current LIGO-Virgo search sensitivity is well approximated by a familiar approximation: a source will typically be detected if the estimated signal to noise (SNR) of the second-most-sensitive detector is greater than 8; see, e.g., [58] and references therein. Using this approximation, one can directly evaluate the characteristic volume within which a source will be detected [59]; for nonspinning BH binaries, this estimate is in reasonable agreement with detailed calculations of search sensitivity [5]. In this work,

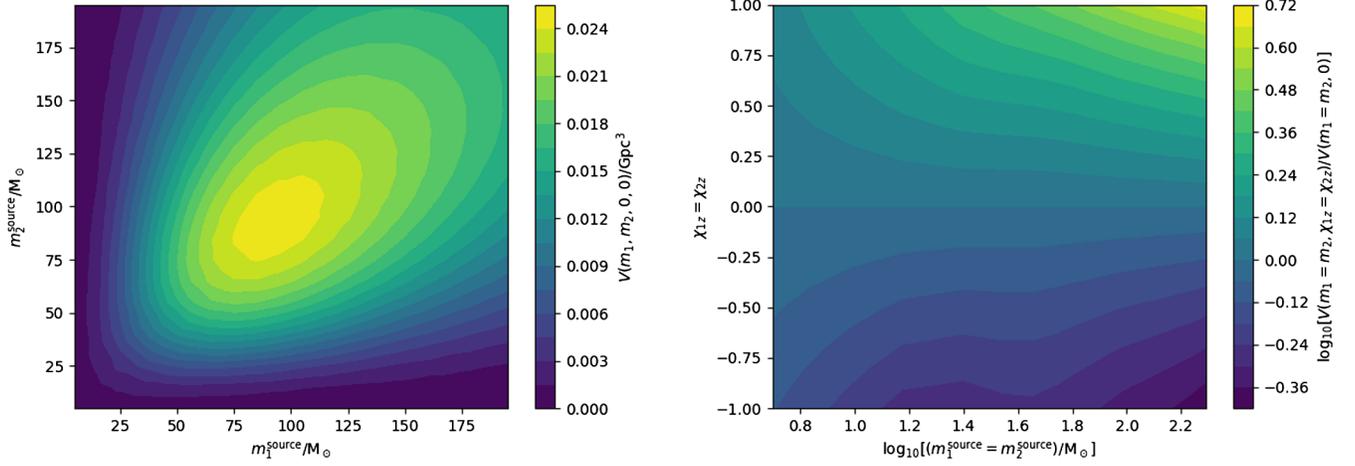


FIG. 1. Estimated sensitive comoving volume (V) versus mass and spin. *Left*: Sensitive comoving volume V at O1 sensitivity for nonspinning BBHs, in cubic giga-parsecs. *Right*: Sensitive comoving volume for equal-mass, equal-spin, nonprecessing BBHs, relative to the zero-spin case. Note that V is strictly increased (decreased) if $\chi_{i,z} > 0$ (< 0), with higher mass making the effect more pronounced.

we therefore adopt the same approximation. Specifically, we estimate the orientation-averaged sensitive 3-volume V to which a search is sensitive by the integral [6,60]

$$V(\lambda) = \int P(< D(z)/D_h(\lambda)) \frac{dV_c}{dz} \frac{dz}{1+z}, \quad (5)$$

where $D(z)$ is the luminosity distance for redshift z ; $D_h(m_1(1+z), m_2(1+z))$ is the horizon distance to which the source can be seen; V_c is the comoving volume; z is the redshift of the merger event; and the cumulative distribution $P(>w) = \int_{w>w(\Omega, l, \psi)} d\Omega d\psi d\cos i$ is a cumulative distribution for $w = 8/\rho$ where ρ is the signal to noise ratio [59–61]. Using this definition for V , we expect that for a uniform comoving merger rate \mathcal{R} (e.g., in units of $\text{Gpc}^{-3} \text{yr}^{-1}$), and after observing at this sensitivity for a time T , the average number of detections will be

$$\mu(\mathcal{R}, \Lambda) = \int (VT)(\lambda) \mathcal{R} p(\lambda|\Lambda) d\lambda, \quad (6)$$

where $p(\lambda|\Lambda)$ is the probability density function for a random binary in the Universe to have intrinsic parameters λ . In this expression, Λ denotes the parameters that characterize the distribution from which all coalescing binaries are drawn. To calculate the horizon distance D_h and hence V for each combination of candidate binary parameters, we use the IMRPHENOMD gravitational waveform approximation [62,63].

The procedure described above allows us to estimate V for any nonprecessing binary. Figure 1 shows this estimate as a function of the component masses, based on a single LIGO detector operating at O1 sensitivity. Motivated by LIGO observations to date, however, we assume black holes will not be rapidly spinning. In these circumstances, spin has at best a modest impact on the sensitive volume;

further complications due to precession would be expected to be smaller still [64,65].

Though we pursue a semianalytic estimate for VT and hence the expected number of GW-detected events, detailed analysis of gravitational wave searches in real data with synthetic sources can evaluate μ and hence the search sensitivity directly [4,5,66,67]. Such an approach will be particularly necessary when search selection biases (e.g., due to detector noise non-Gaussianity) cause the search sensitivity threshold to deviate away from the simple SNR threshold described here.

D. Examples of phenomenological population models

Motivated by the qualitative features of predictions produced by detailed binary formation calculations, several groups have proposed purely or weakly phenomenological models for the binary mass distribution [5,21,21,22,68]. Following [5,21], we adopt a pure truncated power law for the relative intrinsic probability $p(m_1, m_2)$ for the source-frame masses in m_1 and m_2 . Departing from previous work, we assume the probability density is nonzero only in a region $m_{\min} \leq m_2 \leq m_1 \leq m_{\max}$, and $m_1 + m_2 \leq M_{\max}$. Unless otherwise noted, we assume that M_{\max} is a property of the detector, not astrophysics, and following the conservative scenario described in [21] fix it at $200 M_{\odot}$. With these assumptions, our mass distribution model has parameters α_m , k_m , m_{\min} , m_{\max} and a functional form

$$p(m_1, m_2) = \frac{(m_2/m_1)^{k_m} m_1^{-\alpha_m}}{(m_1 - m_{\min})} \times C(\alpha_m, k_m, m_{\min}, m_{\max}, M_{\max}) \quad (7)$$

inside our mass limits and zero elsewhere, representing a truncated power law in m_1 with index $-\alpha_m$ and a simple power-law conditional distribution $p(m_2|m_1)$ in

secondary mass. The normalization constant C is defined so $\int_A dm_1 dm_2 p(m_1, m_2) dm_1 dm_2 = 1$. Unless otherwise noted, we will adopt $k_m = 0$ in this work. Because GW networks are much more sensitive to more massive BHs with $M \gtrsim 200 M_\odot$, this model and its fiducial choices (e.g., $\alpha_m \simeq 2$) produce a detected merger distribution $\propto Rp(m_1, m_2)VT$ which is roughly uniform over a wide range of masses, usually terminated by the specific cutoff choices m_{\max} , m_{\min} rather than by selection biases against low mass black holes or the rarity of massive BBHs. In the analysis described below, we leave M_{\max} fixed.

Motivated by binary neutron star observations as well as the desire to reproduce arbitrary substructure and features in the mass distribution, we will also examine Gaussian mass distributions in component mass m_i ,

$$p_G(m_1) = \mathcal{N}(\bar{m}, \sigma_m)(m_1), \quad (8)$$

which is characterized by its mean value \bar{m} and variance σ_m . In this work, we will typically explore the special case of $p(m_1, m_2) = p_G(m_1)p_G(m_2)$ and apply this distribution to the case of binary neutron stars, where the narrow width σ relative to the mean \bar{m} implies the distribution has effectively no support for undesirable regions (e.g., $m < 0$). Finally, for complete generality, we also discuss mixtures of mass distributions, including Gaussian mixture models as previously employed in [29]:

$$p(m_1, m_2 | \Lambda) = \sum_\alpha w_\alpha p_\alpha(m_1, m_2 | \Lambda_\alpha). \quad (9)$$

This latter approach allows complete generality and, with suitable smoothing priors on w , the ability to reproduce arbitrarily complicated mass distributions and circumvent systematic limitations due to our choice of model. In particular, these more generic models would allow us to reproduce features previously proposed in the literature, including overabundances at specific masses near the pair-instability supernova threshold [69–74].

For binary black hole spins, we adopt a simple flexible phenomenological model for each BH spin magnitude χ_i : a beta distribution,

$$p(\chi_i | \alpha_{\chi_i}, \beta_{\chi_i}) = \frac{\chi_i^{\alpha_{\chi_i}-1} (\chi_{\max} - \chi_i)^{\beta_{\chi_i}-1}}{B(\alpha_{\chi_i}, \beta_{\chi_i}) \chi_{\max}^{\beta_{\chi_i} + \alpha_{\chi_i}}}, \quad (10)$$

with unknown shape parameters α_{χ_i} and β_{χ_i} ($i = 1, 2$). This tractable two-parameter distribution allows us to fit to the observed mean and variance—all that the sparse sample of existing observations will allow. In this work, we for simplicity assume both black hole spins are drawn from the same distribution and $\chi_{\max} = 1$. Likewise, for simplicity we adopt the unphysical but easily described parametrization of the spin-orbit misalignment $\theta_i = \arccos \hat{\mathbf{L}} \cdot \hat{\mathbf{S}}_i$ proposed by Talbot and Thrane [32]: a unimodal distribution based

on a Gaussian in $\cos \theta$ that smoothly deforms into a uniform distribution in the limit of large σ_{χ_i} :

$$p(\cos \theta_i | \sigma_{\chi_i}) \propto \mathcal{N}(\cos \theta_i; 1, \sigma_{\chi_i}). \quad (11)$$

When using this model, we assume the polar angles ϕ_i of each spin vector relative to the orbital angular momentum direction $\hat{\mathbf{L}}$ are uniformly distributed between $0, 2\pi$. In this work, we assume BH spins are drawn from the same spin misalignment distribution $\sigma_{\chi_1} = \sigma_{\chi_2}$. In this approach, as in our parameter inference, all spins are assumed specified at a gravitational wave frequency $f_{\text{ref}} = 20$ Hz. No compelling reason exists that astrophysical formation processes should cause binaries of different masses and spins to be drawn from a single, universal misalignment distribution at an arbitrary reference frequency f_{ref} ; see, e.g., [16,75] for more detailed models. That said, this phenomenological approach is qualitatively consistent with the kinds of misalignments produced by binary SN natal kicks (e.g., $1 - \cos \theta_i \lesssim 0.1$ for BH natal kicks of order 50 km/s [40]), allowing us a simple way to characterize whether observations support or disfavor plausible amounts of spin-orbit misalignment.

E. Useful phenomenological parameters

Observations will constrain combinations of these phenomenological parameters which reflect clear physical features in the *observed* (selection-biased) distribution of binary black holes. We can better characterize what we learn from GW observations early on by adopting coordinates conforming to these features.

For example, we could have a mixture model [Eq. (9)] consisting only of elements with distinctive features, each characterizing a distinctive subpopulation of BHs. Such subpopulations might be BHs near the pair-instability supernova peak, binary neutron stars, and a population of binaries with a continuous mass spectrum formed through hierarchical growth in globular clusters (see, e.g., [21,34,76] and references therein). In such a scenario, observations quickly constrain each element, leveraging their distinctive features to identify the relative rates $\mathcal{R}w_\alpha$ and the subpopulations from each domain to constrain that region's parameters. For the first few tens of events, these observations will principally constrain the mean and variance of the *detection-weighted* subpopulation $p_\alpha(m_1, m_2)VT$. We therefore expect that the following coordinate system will produce roughly uncorrelated observables for a typical model: (a) the relative rates $\mathcal{R}w$ for different subpopulations; (b) the mean chirp mass $\overline{\mathcal{M}_{c,\alpha}}$, symmetric mass ratio $\bar{\eta}_\alpha$, effective spin $\bar{\chi}_{\text{eff},\alpha}$, and mean spin $\bar{\chi}$ in each subpopulation, based on our understanding of GW measurement errors; and (c) the respective widths $\Sigma_{\mathcal{M}_{c,\alpha}}, \Sigma_{\eta,\alpha}, \Sigma_{\chi_{\text{eff},\alpha}}, \Sigma_\chi$, where we adopt upcase to distinguish between these symbols and our

model hyperparameters. In Appendix C, we use order-of-magnitude arguments to explain how reliably each of these quantities can be measured.

In the context of our fiducial single-component model, we adopt a reference mass $m_1 = m_{\text{ref}} = 15 M_\odot$ and characterize the overall event rate not by its normalization, which depends on unobserved binaries with high and low masses, but by the event rate $\mathcal{R}p(m_{\text{ref}})$ of binaries whose primary m_1 has a mass comparable to GW151226 [77]. We identify other natural coordinates for the distribution of m_1 via its detection-weighted cumulative distribution $\mathcal{P}(< m_1)$:

$$\mathcal{P}(< x) = \frac{\int d\lambda V(\lambda)p(\lambda)\Theta(x - m_1(\lambda))}{\int d\lambda V(\lambda)p(\lambda)}. \quad (12)$$

The mass corresponding to the upper (lower) bound of the 90% symmetric detection-weighted probability on m_1 serves as a proxy for m_{max} (m_{min}) which is directly observable and thus a more natural coordinate.² In this work, we emphasize the upper bound m_* of the detection-weighted mass distribution:

$$\mathcal{P}(m_*) \equiv 0.95. \quad (13)$$

For BH spins, closed-form expressions for the appropriate mean values and variances are generally not available for arbitrary selection biases VT ; however, to the extent that VT depends only weakly on BH spin, our model for BH spins and misalignments [Eqs. (10), (11)] implies that

$$\bar{\chi} \simeq \frac{\alpha_\chi}{\alpha_\chi + \beta_\chi}, \quad (14a)$$

$$\Sigma_\chi^2 \simeq \frac{\alpha_\chi \beta_\chi}{(\alpha_\chi + \beta_\chi)^2 (\alpha_\chi + \beta_\chi + 1)}, \quad (14b)$$

$$\bar{\chi}_{\text{eff}} \simeq \bar{\chi} \overline{\cos \theta}, \quad (14c)$$

$$\overline{\cos \theta} \simeq \frac{\text{erf}(\sqrt{2}/\sigma) + 2\sigma(e^{-2/\sigma^2} - 1)/\sqrt{2\pi}}{\text{erf}(\sqrt{2}/\sigma)}, \quad (14d)$$

for our fiducial case where both BH spins are drawn from the same distributions; in these expressions, Σ_χ^2 refers to the variance of the one-dimensional χ distribution, while $\bar{\chi}$ refers to its mean.

F. Interpreting results: Posterior predictive distributions and revised priors

If we ask any question about compact binary properties x rather than model hyperparameters Λ , the only quantity that

²By contrast, Talbot and Thrane [22] introduce a model which depends on both a minimum mass m_{min} and a tapering mass scale δm , but only a linear combination of them is easily observable; see their Fig. 5.

appears in our posterior inferences $p(\Lambda|\{d\})$ informed by our observations $\{d\}$ is the posterior predictive distribution $p_{\text{ppd}}(x|\{d\})$:

$$p_{\text{ppd}}(x|\{d\}) = \int d\Lambda p(x|\Lambda)p(\Lambda|\{d_k\}). \quad (15)$$

The posterior predictive distribution (PPD) encodes our best estimates of the properties of any randomly selected future binary, based on observations to date and accounting for our initial prior knowledge about Λ . Unlike the model parameters themselves, which may be highly degenerate and lack physical meaning, the PPD provides an unambiguous estimate for how likely different binary parameters are, given our knowledge. Note that by design, the PPD is a probability distribution and, folding in all uncertainties, does not have an error estimate.

As events accumulate, we can use posterior constraints $p(\Lambda|\{d\}_k)$ on model hyperparameters Λ based on the first $k = 1 \dots N$ observations to provide a nuanced, observationally revised perspective on future measurements $k > N$. These prior insights can be particularly powerful when individual future measurements are only weakly informative about certain binary parameters such as the mass ratio or spin; see, e.g., [78,79] for examples.

To be concrete, our usual population inferences are performed using a single fiducial choice of reference prior $p_{\text{ref}}(x) = p(x|\Lambda_{\text{ref}})$: the posterior is $p(x|d_k, \Lambda_*) = p(d_k|x)p(x|\Lambda_*) / \int p(d_k|x)p(x|\Lambda_{\text{ref}})$. We exploit prior measurements via

$$p(x|d_k, \{d\}) = \frac{p(d_k|x) \int d\Lambda p(x|\Lambda)p(\Lambda|\{d_k\})}{\int dx p(d_k|x) \int d\Lambda p(x|\Lambda)p(\Lambda|\{d_k\})}. \quad (16)$$

In this expression, the numerator $\int d\Lambda p(x|\Lambda)p(\Lambda|\{d_k\})$ is the posterior predictive distribution described above.

III. CONTROLLED TESTS WITH SYNTHETIC POPULATIONS AND MEASUREMENTS

To demonstrate our method can infer population parameters, we perform several validation studies using toy models which mimic key features of real gravitational wave observations. These completely controlled illustrations also let us highlight what can be inferred and why about the mass and spin distribution, within the context of our approach. Finally, these examples allow us to demonstrate how population inference can strongly inform the interpretation of individual future GW observations.

A. BNS mass and (aligned) spin distribution

For each component of a binary neutron star (BNS), observations of galactic pulsars suggest that the component masses are drawn from a Gaussian distribution with mean

$1.33 M_{\odot}$ and standard deviation $0.09 M_{\odot}$ [80]. Observations of pulsars and theoretical models of pulsar spin-down suggest that if both NS are not recycled, then their dimensionless spins will be small [$\simeq \mathcal{O}(0.05)$]. Under the assumption that NS spins are parallel to their orbital angular momentum, we construct a synthetic population drawn from this phenomenological model; construct synthetic observations for each binary, recovering 13 synthetic sources based on a three-detector advanced LIGO/Virgo network using a threshold set by the second-most-sensitive detector's recovered amplitude; perform full GW inference on each source using RIFT [49]; and, with the resulting posterior distributions, use the techniques of Sec. II to infer the underlying NS mass and spin distribution. In our reconstruction, we assume both components of a NS binary are independently drawn from a Gaussian distribution with unknown mean and variance, and with spins $\chi_{i,z}$ drawn from a beta distribution with unknown mean and variance, such that $|\chi_{i,z}| \leq 0.05$.

Figure 2 shows the synthetic measurements used as inputs in our calculation. These synthetic measurements incorporate significant uncertainty in each source's redshift, which contributes to the overall uncertainty in each binary's chirp mass. For each neutron star in our synthetic population, we use the APR4 equation of state to calculate each neutron star's tidal deformability $\lambda_i = \lambda(m|\text{APR4})$. We generate and recover our synthetic sources with IMRPHENOMD_NRTIDAL [81]. Figure 3 compares our recovered NS mass and spin distribution. When inferring source parameters, our waveform model and parameter inferences include the effects of NS tides, treating each NS tidal deformability λ_i as a free parameter. Despite considerable uncertainties in each measurement, each BNS observation constrains that binary's chirp mass reasonably well, to an accuracy $\sigma_{\mathcal{M}_c} \simeq 0.05 M_{\odot}$, dominated by uncertainty in source redshift. Because GW measurements are only weakly informative about the mass ratio, these measurements each constrain the total mass to be $m_1 + m_2 \simeq 2^{6/5} \mathcal{M}_c$ to an accuracy $\sigma_{\mathcal{M}_c} 2^{6/5}$; averaging all such observations, we can deduce the mean NS mass \bar{m} . With $n = 13$ such measurements, we expect to constrain the mean mass of the population to a 1 standard deviation accuracy $\sqrt{\sigma_{\mathcal{M}_c}^2 2^{12/5} / 4 + \sigma^2} / \sqrt{n} \simeq 0.027 M_{\odot}$, which compares favorably to $0.02 M_{\odot}$, the standard deviation of our Bayesian estimate for \bar{m} . (A similar analysis shows that we constrain the NS population standard deviation σ_m almost entirely through these one-dimensional chirp mass constraints.) Because GW measurements have a smaller statistical uncertainty than the astrophysical population width in total mass, the accuracy to which we constrain the mean NS mass is dominated by a simple frequentist error estimate (σ/\sqrt{n}), allowing us to reliably project the information we will extract about NS masses from future GW observations.

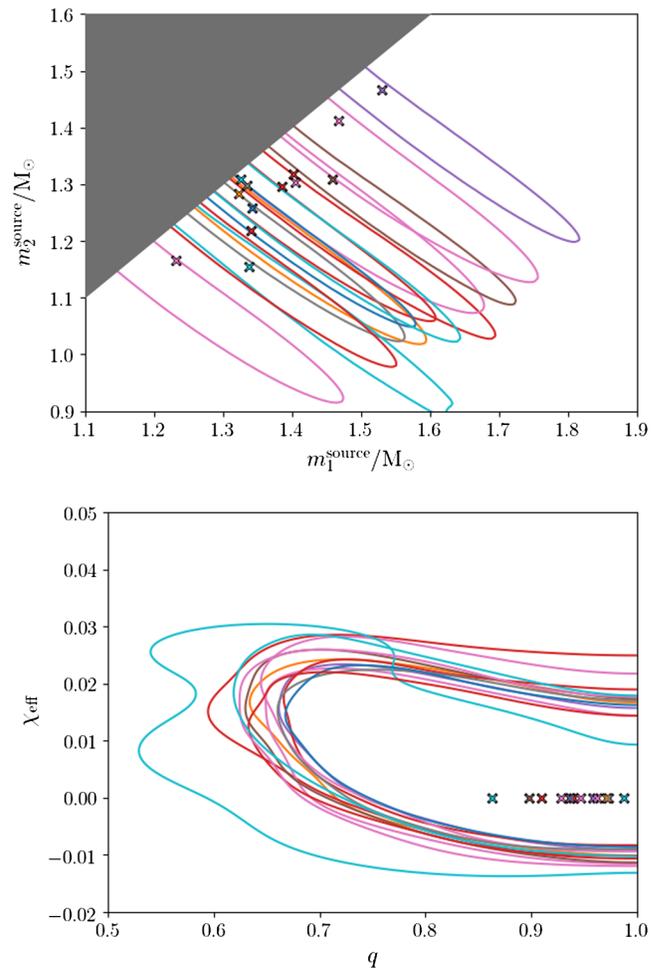


FIG. 2. Source information for our synthetic BNS population: For each synthetic signal used in the BNS population reconstruction calculation described in Sec. III A, these two panels show the true injected source-frame parameters (as crosses) and posterior distributions (contours of their 95% highest posterior density regions). Each color corresponds to a different source. Source parameters have been inferred using full Bayesian parameter inference via RIFT, as described in the text.

The measurement accuracy for GW measurements of BNS has been long known [54], and their implications for astrophysics (e.g., mass and BNS spin distributions) have been immediately apparent; see, e.g., [82–84] and references therein. We provide the first end-to-end demonstration of how well binary NS population parameters can be measured, using a detailed waveform model at a level where waveform systematics should not dramatically impact the mass, spin, or tidal parameter inferences being performed. By contrast, many previous studies focusing on NS tidal deformation have demonstrated that waveform systematics could bias inferences [85–87], if not controlled. Only recently have systematic errors between waveform models diminished enough to enable consistent inference; see, e.g., [88].

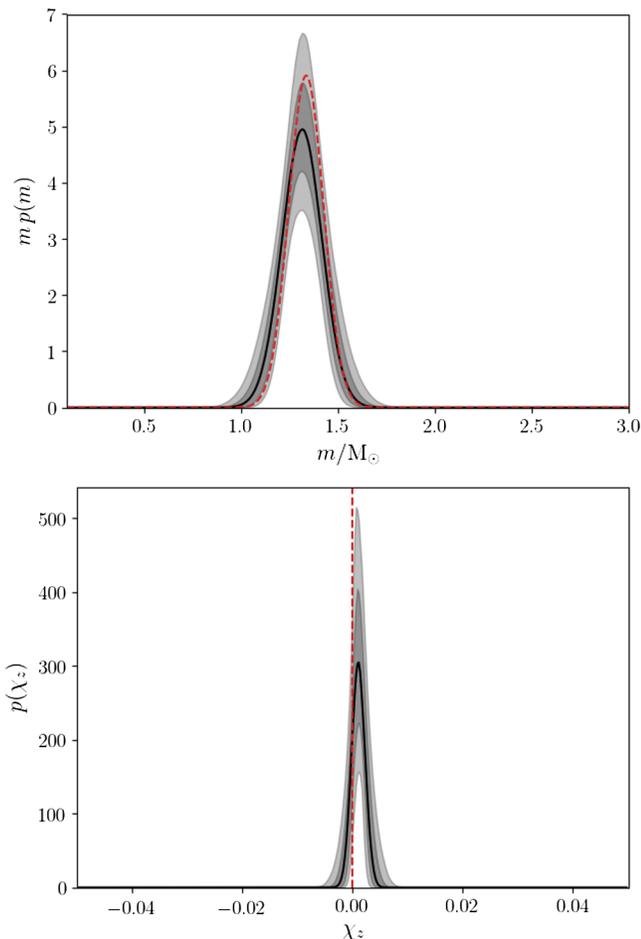


FIG. 3. Recovered properties of NS mass and spin distribution: For the synthetic population of BNS sources illustrated in Fig. 2, this figure shows the recovered mass distribution (top figure) and spin distribution parameters (bottom figure) derived using the Gaussian mass and β -distribution spin model described in the text. The solid line indicates the median distribution; the shaded regions indicate the 68% and 95% credible intervals. Red dashed lines denote the true underlying distribution. In the case of spin, note that the truth is a delta function at zero, so it would require an infinite number of detections to fall within the constraints on this plot.

Reliable population inference allows us to draw informed conclusions about future measurements, using previous observations as prior input. Particularly for cases like NS binaries where individual measurements can be weakly informative and produce highly correlated constraints on NS parameters, these prior inputs enable much sharper constraints on astrophysical parameters. As a concrete example, Fig. 4 shows inferences about one parameter ($\tilde{\Lambda} = \frac{16}{13} [(m_1 + 12m_2)m_1^4\lambda_1 + (m_2 + 12m_1)m_2^4\lambda_2] / (m_1 + m_2)^5$) of one of our synthetic NS binaries, where the inferences are performed in isolation (blue line) and using information obtained from all other NS observations in our sample about NS masses and spins (but not tides $\tilde{\Lambda}$, which are presumed arbitrary and spin). Because our other

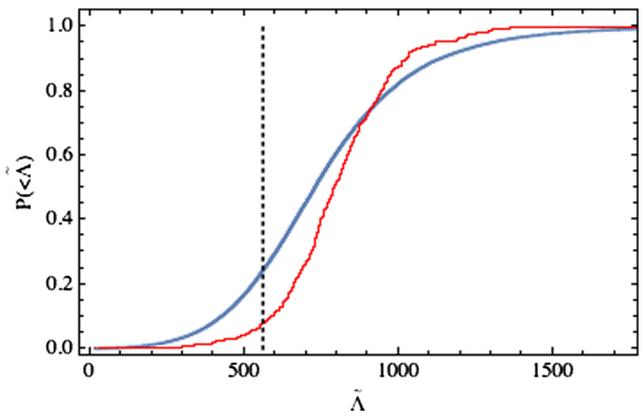


FIG. 4. Population measurement enables sharper constraints on NS tides: Cumulative posterior distribution of $\tilde{\Lambda}$ for one of the synthetic sources in our BNS population model. Blue curve shows a single-event analysis, not exploiting information about the mass and spin distribution from other events; red curve shows an analysis based on Eq. (16) that employs our best estimate for the underlying mass and spin distribution, as constrained from the population of events in our BNS synthetic sample.

measurements have allowed us to strongly constrain the NS population’s mass and spin distribution, we can exploit correlations between our inferences about these parameters and the NS tidal deformability to more tightly constrain this parameter. In this way, even though only the strongest few GW measurements will provide most of the information about NS tides and the nuclear EOS, by exploiting population measurements we expect to more efficiently draw conclusions using all available information about the NS population.

B. BBH mass and (precessing) spin distribution

To assess our ability to simultaneously constrain both the mass and spin distribution of binary black holes using GW observations, we constructed a synthetic population drawn from our fiducial BBH population model, with parameters as described in Table I. Following the procedure described in Appendix B, we drew freely from this population, then selected a subsample based on their relative probability of detection, producing 25 events based on 300 days of synthetic observation at O1 sensitivity. For both the synthetic population and sensitivity model, we approximate VT by neglecting any effects of spin, as a self-consistent leading-order approximation. For each event, we generated 1000 fair draws from a synthetic posterior distribution, using the procedure described in Appendix A. These synthetic or “mock” posterior distributions mimic the effects of full GW parameter inference, but by construction only explicitly constrain the binary chirp mass, mass ratio, and effective spin χ_{eff} of each event. Figure 5 shows the specific source population and synthetic posteriors used in this analysis. Using these synthetic posterior distributions, we apply the population inference procedure described in

TABLE I. Synthetic BBH population model: This table shows the parameters of the population model family we adopt to generate and recover a synthetic binary black hole population as described in Sec. III B. The population is characterized by an overall BBH merger rate \mathcal{R} ; a power-law slope α_m for the primary mass, between m_{\min} and m_{\max} ; a beta distribution for spin magnitude, characterized by the two parameters α, β [Eq. (10)]; and a characteristic misalignment σ_χ for the angle between BH spins and the orbital angular momentum at our reference frequency [Eq. (11)]. This analysis also fixes the maximum allowed total mass M_{\max} (i.e., $m_1 + m_2 \leq M_{\max}$) to $200 M_\odot$. In this model, both black hole spins are assumed drawn independently from the same distribution. The second row shows the values of these parameters used to generate our synthetic population. The third row shows the range of parameter space we explore when attempting to reproduce our data. The fourth row shows the prior distribution adopted for each parameter, all assumed *a priori* independent; in this row, “log uniform” implies the prior distribution for any variable x is uniform as a function of $\log x$ [i.e., $p(x) \propto 1/x$]. Note that for simplicity we have assumed the minimum mass is known.

Quantity	\mathcal{R} [Gpc ⁻³ yr]	α_m	m_{\min} [M_\odot]	m_{\max} [M_\odot]	α_χ	β_χ	σ_χ
Synthetic population	100	0.8	5	40	1.1	5.5	0.4
Prior range	[10 ⁻¹ , 10 ⁶]	[-5, 5]	[5, 5]	[30, 195]	[10 ⁻⁴ , 10 ⁴]	[10 ⁻⁴ , 10 ⁴]	[10 ⁻² , 10 ²]
Prior distribution	Log-uniform	Uniform	Uniform	Uniform	Log-uniform	Log-uniform	Log-uniform

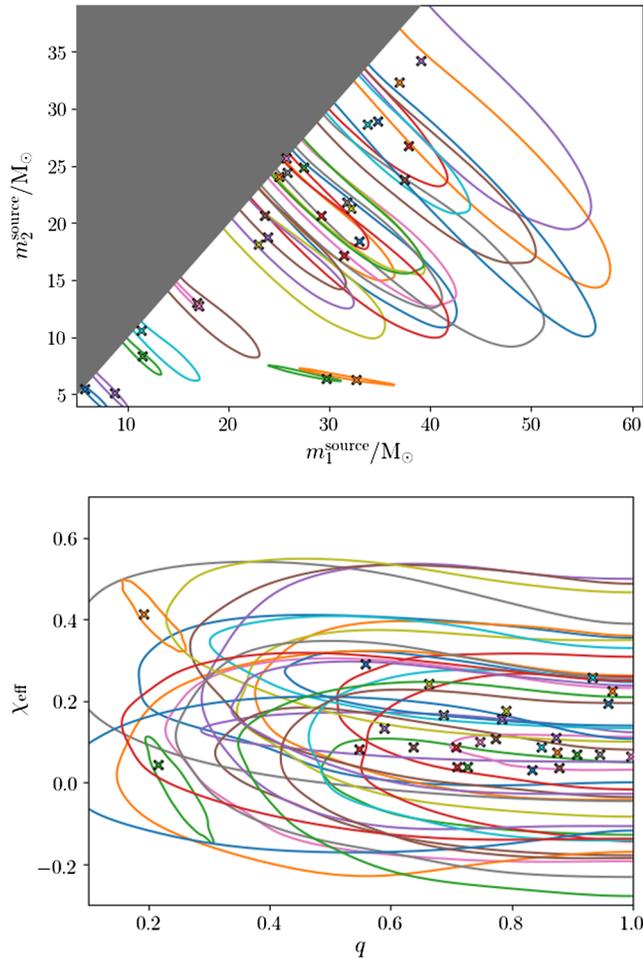


FIG. 5. Source information for our synthetic BBH population: For each synthetic signal used in the BH population reconstruction calculation described in Sec. III B, these two panels show the true injected source-frame parameters (as crosses) and posterior distributions (contours of their 95% highest posterior density regions). Each color corresponds to a different source.

Sec. II to produce our best estimates for the population parameters responsible for our synthetic observations. As summarized in Table I, our model has parameters

$$\Lambda \equiv (\mathcal{R}, \alpha_m, m_{\min}, m_{\max}, \alpha_\chi, \beta_\chi, \sigma_\chi). \quad (17)$$

To be consistent with the priors adopted in other work [5], we express our results after reweighting to correspond to a Jeffries prior on the rate [$\pi(\mathcal{R}) \propto \mathcal{R}^{-1/2}$]. Even with only 25 events drawn from a preferentially low-spin population, our calculations show that GW measurements should strongly constrain the mass and spin distribution of binary black holes.

Figure 6 shows how well we can determine the merger rate versus binary masses, such as the primary mass. Notably and in good agreement with previous work, we find we can strongly constrain the maximum detectable mass in the population [21,32]. Following the discussion Sec. II E, however, we emphasize that while the maximum *detectable* mass—demarcated by a sharp cutoff in the observed population—is well constrained, the *parameters* $\mathcal{R}, m_{\max}, \alpha_m$ have a degeneracy: as shown in Fig. 6, a population with extremely few but very massive BHs is hard to rule out, enabling larger m_{\max} to be consistent with our synthetic observations. Additionally and for the first time, we demonstrate how to self-consistently compute both the overall event rate distribution, including Poisson error, while simultaneously constraining the mass distribution. Previous investigations have used specially devised calculations which marginalize over the event rate distribution, producing results that (for a suitable Jeffries prior) are consistent with our results for the marginal mass distribution. As demonstrated in Fig. 6, to produce a self-consistent rate distribution, due to strong correlations between the event rate and mass distribution, we must simultaneously measure the mass-dependent merger rate in the local universe. Because the correlation between the event rate and mass distribution arises through the expected number of events, we can provide a simple analytic model

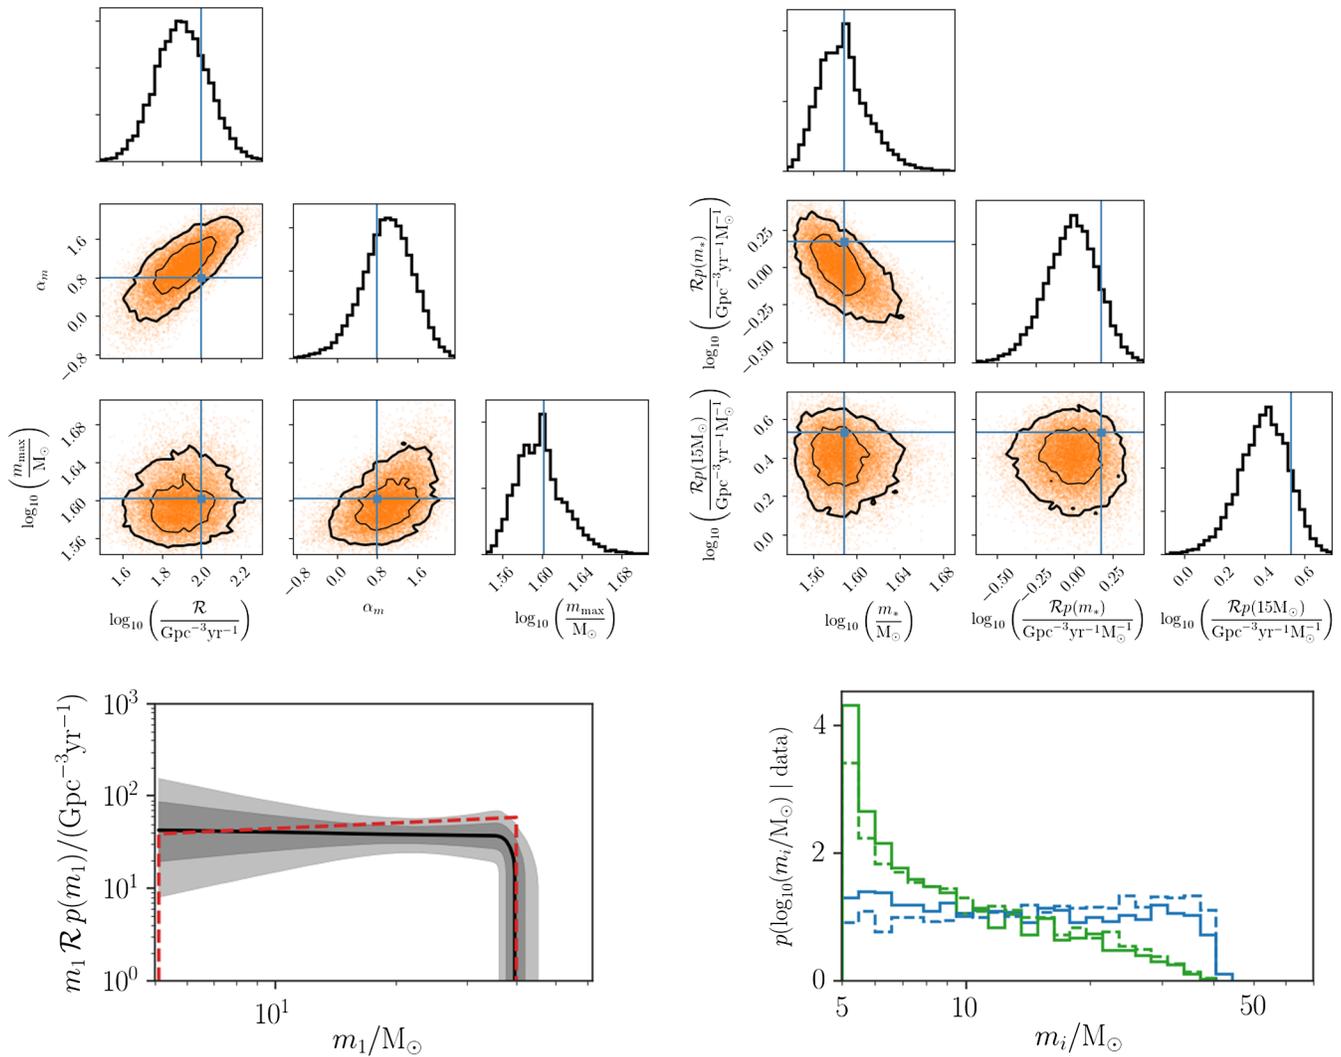


FIG. 6. Inferred merger rate versus mass: This figure shows how our estimated merger rate versus mass compares with the known distribution used to generate our synthetic source population. For a more thorough statistical test, see the P - P plots in Appendix D. *Top left*: This group of figures represents the one- and two-dimensional marginal posterior distributions for \mathcal{R} , α_m , and m_{\max} , with the true values overlaid as blue crosshairs. *Top right*: This group of figures represents the one- and two-dimensional marginal posterior distributions for m_1^* , $\mathcal{R}p(m_1^*)$, and $\mathcal{R}p(15 M_\odot)$, with the true values overlaid as blue crosshairs. *Bottom left*: In this figure, the red dashed line shows the characteristic merger rate associated with a given mass scale $[m_1 \mathcal{R}p(m_1)]$ versus primary mass m_1 . The black line shows the median inferred value, and the two gray shaded regions show the symmetric 68% and 95% credible regions. *Bottom right*: The solid lines in this figure show our posterior predictive distribution $p(m_i|D)$: the best estimate for the probability of a future event being detected having masses m_i . In this figure, blue and green correspond to the primary and secondary masses. For comparison, the dotted lines show the true astrophysical distribution.

for the correlation between the mass distribution and event rate, as described in Appendix C.

With 25 events, our population model has enough information to produce strong constraints on the underlying phenomenological distributions, even for parameters such as spin which are weakly constrained by individual measurements. Figure 7 illustrates how informative these constraints can be about the spin distribution. This figure compares the true marginal distribution of q, χ_{eff} for the BH-BH population to our best (posterior predictive) estimate of that distribution. Even with only a few tens of detections,

the estimate traces the general structure of the true distribution. In particular, we can clearly and unambiguously identify that a bias in the χ_{eff} distribution toward positive values suggests an underlying tendency toward alignment. Of course, our synthetic observations were intentionally drawn from the model family we use to fit it; in general, the underlying astrophysical distribution may have a form outside the model family we adopt, introducing small biases into our interpretation. Nonetheless, our analysis substantially generalizes previous proof-of-concept demonstrations on how well BH measurements can measure BH spin distributions,

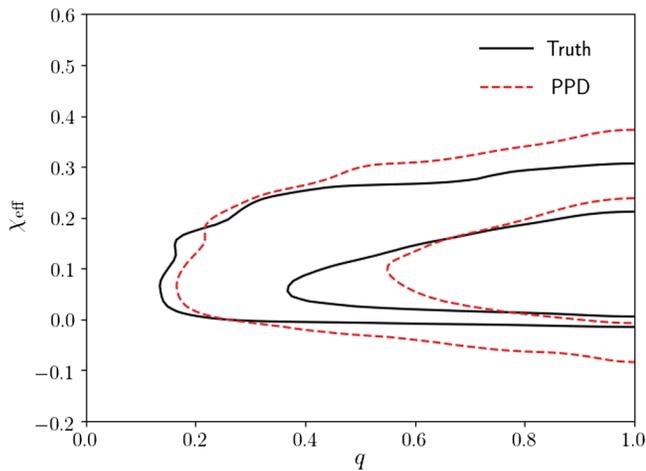


FIG. 7. Recovering the true mass ratio and χ_{eff} distribution: A comparison between the underlying truth (black solid contours) and the inferred posterior predictive (red dashed contours) for the q, χ_{eff} marginal distribution. The inner (outer) contour for each denotes the 50% (90%) highest probability density credible region.

not being limited to a single spin magnitude, a discrete and restrictive family of orientation distributions, or similar strong prior adopted in previous investigations [36,39].

Even with only 25 events, we strongly constrain the BH spin distribution, in both magnitude and orientation (Fig. 8). As described in Appendix C in greater quantitative detail, these two constraints are easily understood. For this synthetic analysis, the upper limit on spin follows from the χ_{eff} distribution of recovered sources. Since our synthetic observations included no events with large χ_{eff} , we can be confident BH spins are not extremely large, since by chance we ought to have found one large value of χ_{eff} out of 25, even allowing for uncertainty in how they are oriented. Similarly, because our synthetic population is preferentially aligned ($\sigma = 0.4$), the recovered population shown in Fig. 2 has a χ_{eff} distribution biased toward positive values. Using Eq. (14) for $\bar{\chi}_{\text{eff}}$, the bias in χ_{eff} inevitably implies $\cos \theta$ is preferentially positive and, as described in Appendix C, allows us to limit σ .

In this analysis, we employ conservative synthetic posteriors which assume only the chirp mass, mass ratio, and effective spin can be constrained with GW measurements. Precessing, coalescing binaries can produce a rich symphony of gravitational waves just prior to and during merger, reflecting complex binary dynamics and strong-field multimodal radiation. Given the high expected event rate in ongoing gravitational wave surveys, we expect that future observations will provide clear examples of precessional dynamics, if nature produces them, and that these measurements will allow us to much more sharply constrain the BH spin distribution. However, for massive BH binaries, model systematics complicate attempts to measure BH parameters, including spin. We will conduct full

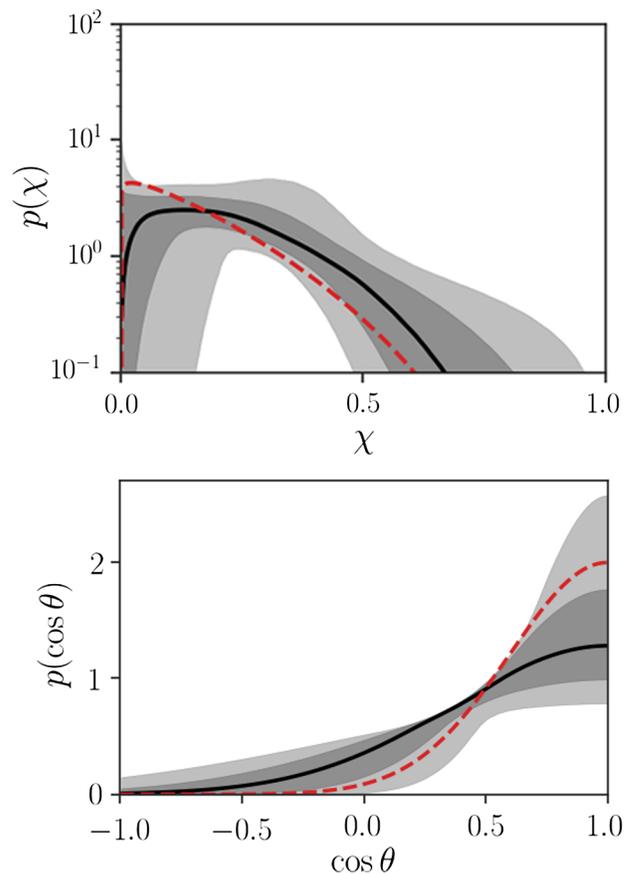


FIG. 8. Inferred spin distribution derived from synthetic BBH observations: The top panel shows our inferences about the total BH spin; the bottom panel shows our inferences about BH spin-orbit misalignment. In both panels, the red dashed lines show the underlying distribution, while the black solid lines and shaded regions show the median recovered parameter distribution. To a first approximation, the constraints on spin magnitude and misalignment are as needed for the population model to reproduce the mass and χ_{eff} distribution of the underlying population as shown in Fig. 7.

end-to-end calculations with synthetic data and state of the art models in future work.

IV. ANALYSIS OF REPORTED OBSERVATIONAL RESULTS

To date, five confident binary black hole mergers have been reported: GW150914 [89], GW151226 [77], GW170104 [90], GW170608 [91], and GW170814 [92]—the latter discovered jointly with the Advanced Virgo instrument [3]. Additionally, an astrophysically plausible candidate BBH signal has been reported (LVT151012) [5]. In this section, we describe inferences about the binary black hole population based on reported events, deduced from these reported observations and a simplified model for the network’s search sensitivity. For O1 events, most notably for GW151226, we use full posterior inferences

TABLE II. Inferences about astrophysical binary BH model parameters: This table provides 90% credible intervals for the underlying parameters of our fiducial BBH population model, applied to O1 and reported O2 observations as described in the text. Parameters with clear unimodal structure are represented by their median and the widths of their 90% symmetric probability confidence interval, whereas we only report the 90% upper and lower limits for more poorly constrained parameters. For the spin magnitude distribution, rather than show the (highly correlated) credible intervals for the underlying sampling variables α_χ , β_χ , we instead show credible intervals for the mean value of χ and the standard deviation of χ . We also show the posterior predictive range of spin magnitudes χ and effective spins χ_{eff} . We apply an asterisk (O2*) to all O2 results, to highlight the nonfinal sample, simplified sensitivity model VT, and mocked-up posteriors used in this proof-of-concept analysis.

	\mathcal{R} [$\text{Gpc}^{-3} \text{yr}^{-1}$]	α_m	m_{max} [M_\odot]	$\mathbb{E}[\chi]$	$\text{Std}[\chi]$	$\log_{10} \sigma_\chi$	χ_{eff}	χ
O1	122_{-96}^{+291}	$2.8_{-2.5}^{+1.4}$	70_{-30}^{+110}	$0.28_{-0.15}^{+0.31}$	$0.02_{-0.02}^{+0.25}$	0.1–9.5	$0.00_{-0.24}^{+0.24}$	0.03–0.68
O2*	...	$1.9_{-2.0}^{+1.5}$	39_{-6}^{+98}	$0.24_{-0.12}^{+0.21}$	$0.01_{-0.01}^{+0.19}$	0.3–9.4	$0.00_{-0.19}^{+0.19}$	0.04–0.49

derived from GW data, provided by the LIGO Scientific Collaboration. For O2 events, in lieu of full posterior inferences, we use the procedure described in Appendix A to generate synthetic posterior distributions which closely resemble the reported parameter estimates for mass and χ_{eff} . For simplicity as well as to enable a concrete illustration of our method using real data, we will produce estimates under the (unwarranted) assumption that reported O2 results available to date represent a comprehensive and fair sample of binary black holes seen during LIGO's O2 observing run. In these estimates, we assume O1 and O2 share a common sensitive volume V as estimated in Sec. II C, with observing duration $T_{\text{O1}} = 48.6$ days [5] and $T_{\text{O2}} = 117$ days [93]. Keeping in mind model systematics such as the omission of a salient feature in the mass distribution can demonstrably strongly bias recovered model parameters [21,22], as well as sample incompleteness for our O2-scale analysis, in Table II we provide our inferences about the O1 and O2 population within the context of the fiducial BBH population model described in Sec. III B. For O2 in particular, we emphasize the simplified VT and nonfinal sample used in that analysis, which is provided solely for illustration and to connect to previously published investigations about O2-scale events [16,21,33]; applying our methods to final O2 results with real samples and carefully calibrated VT could produce substantially different astrophysical conclusions.

Figure 9 shows our best estimates for the merger rate of BH-BH binaries of different masses, inferred within the context of the model described in Table I and demonstrated on synthetic data in Sec. III B. Naturally, we estimate an overall BH-BH merger rate and mass distribution consistent with previously reported results [5]. Using a Jeffries' prior for the merger rate, we find $\mathcal{R} = 122_{-96}^{+291} \text{ Gpc}^{-3} \text{ yr}^{-1}$ based on O1. For O2, we find uncertainty in the event rate is reduced by roughly a factor of 2, both through reduced Poisson error (e.g., six instead of three events) and through sharper constraints on the mass distribution (e.g., reducing prospects for a large maximum mass). Our result for O1 is more conservative (wider) than the power-law result reported previously in Abbott *et al.* [5], $97_{-67}^{+135} \text{ Gpc}^{-3} \text{ yr}^{-1}$, because we employ a more flexible

model and therefore incorporate more model systematics, notably including the correlation between event rate and mass spectrum and also the impact of the upper mass cutoff. Conversely, if we employ consistent assumptions, we arrive at the same answers previously reported for O1 [5]. As we adopt a merger rate model that reduces to previously investigated power laws, by design we reproduce the analysis reported in [21]: the events reported during O2 suggest the absence of very massive BHs in the observable population.³ For this reason our inferences about the mass spectrum exponent α_m are considerably wider than prior work which does not take a possible upper mass cutoff into account. Even with the small sample publicly reported so far, our analysis corroborates the analysis in [21] that O2-scale GW measurements could be weakly informative about the maximum mass of coalescing BHs.

As demonstrated in several previous investigations [16,33], we know that BHs in merging binaries likely have low typical spin. For example, based on the distribution of χ_{eff} , Farr *et al.* [33] argued that several members of a discrete array of candidate spin orientations (aligned or isotropic) and magnitude distributions are inconsistent with observations to date, and that BH spins were likely randomly oriented or small. Later, Wysocki and collaborators [16] demonstrated that, if binary black holes arose from isolated binaries whose spins were weakly misaligned by SN natal kicks, then only relatively small BH natal spins were consistent with observations available at the time. As shown in Fig. 10, with more events available to our analysis, and using much more flexible models, we can draw sharper and more generic conclusions about the BH spin distribution, even using only six reported events. First and foremost, exactly as seen with synthetic data, the absence of large χ_{eff} allows us to with increasing confidence bound above the fraction of BHs in merging binaries that have large spin. Too, because collectively the observed population distribution of

³While our assumptions about the mass distribution model have modestly changed relative to Fishbach *et al.* [21], we reproduce their results when adopting the same inputs and mass model.

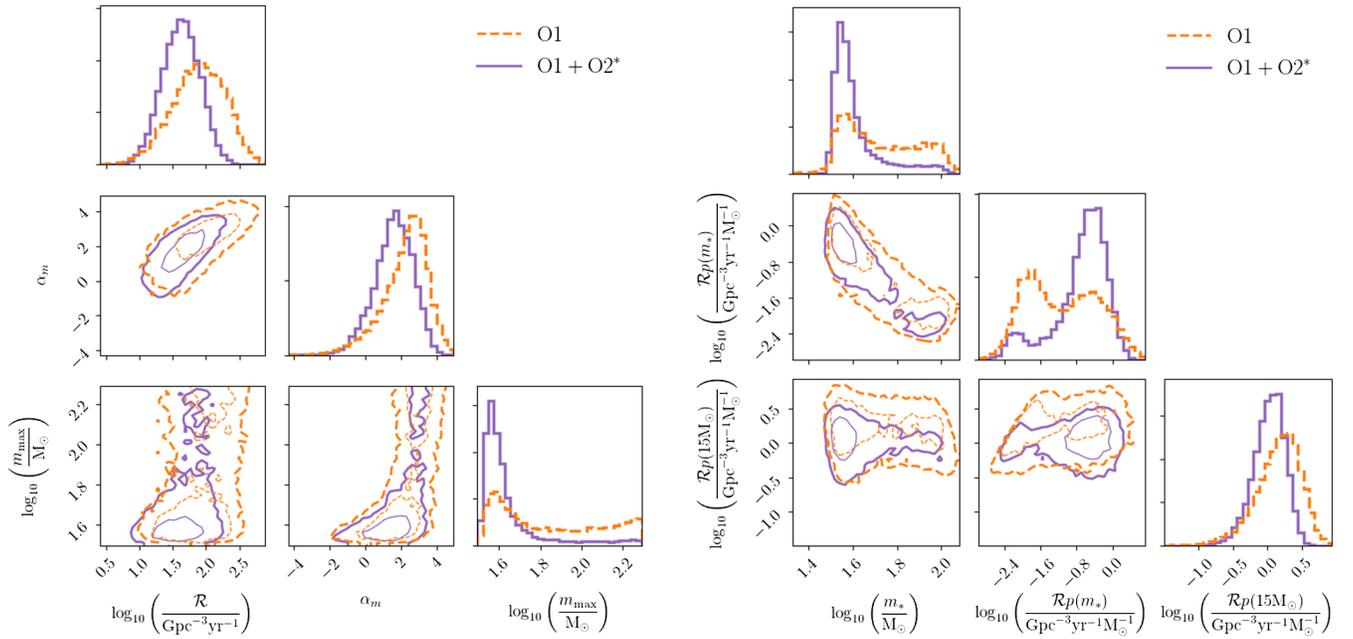


FIG. 9. Inferences about astrophysical binary BH mass distribution: Inferences about the merger rate versus mass of coalescing BH-BH binaries, using only O1 observations (dashed orange) and using O1 and reported O2 observations (solid purple), for simplicity assuming the latter represents a comprehensive and fair sample. We apply an asterisk (O2*) to all O2 results, to highlight the nonfinal sample, simplified sensitivity model *VT*, and mocked-up posteriors used in this proof-of-concept analysis. The panels in this figure follow the format of Fig. 6 for representing one- and two-dimensional marginal posterior distributions.

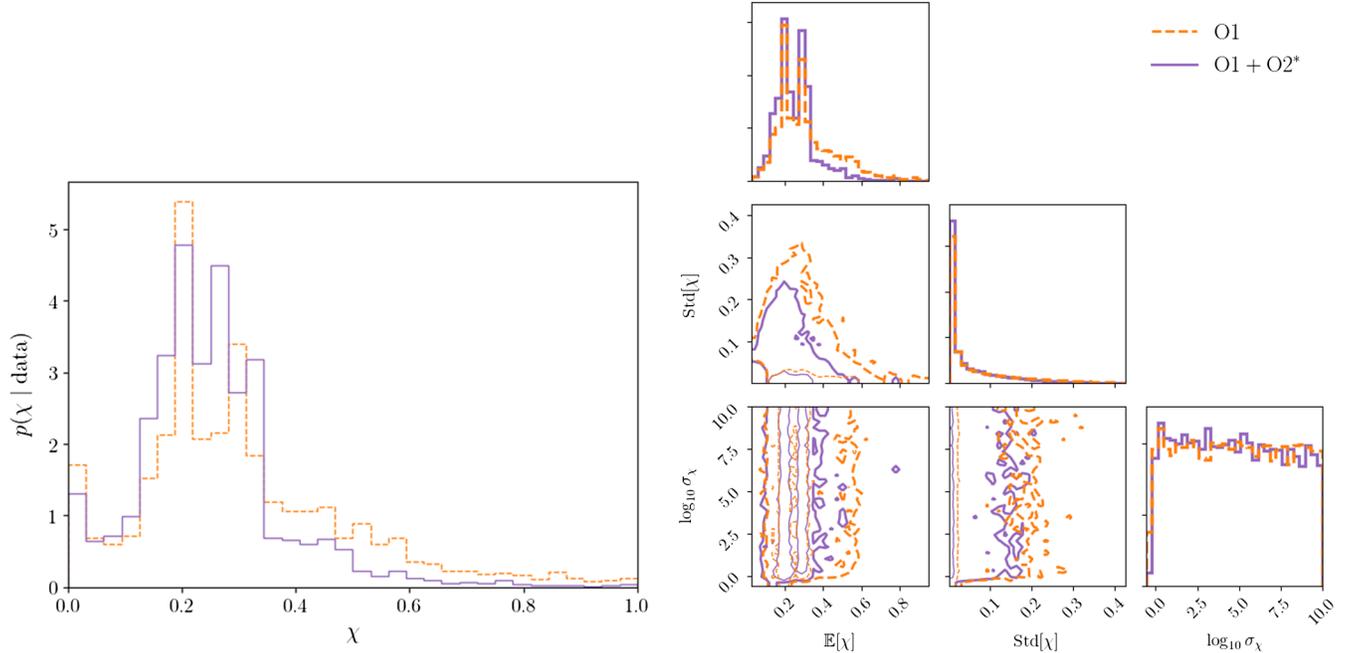


FIG. 10. Inferences about astrophysical binary BH spin distribution. *Left*: Our best estimates for the binary BH spin magnitude distribution (PPD) based on O1 (dashed orange) and O2 (solid purple) observations. We apply an asterisk (O2*) to all O2 results, to highlight the nonfinal sample, simplified sensitivity model *VT*, and mocked-up posteriors used in this proof-of-concept analysis. Due to the low characteristic spin and within the context of the information used in this analysis, these observations remain uninformative about BH spin-orbit orientations. *Right*: Our best estimates for the binary BH spin distribution, as expressed using our model hyperparameters, for O1 and O2.

χ_{eff} remains nearly symmetrically distributed around zero, we can with increasing confidence bound the fraction of binaries that are preferentially aligned and with modest spin. With at least one BH known to have spin (GW151226) and for simplicity assuming the BH spin and mass distribution are uncorrelated, we are led to weakly disfavor scenarios where BHs are preferentially aligned (i.e., small σ is disfavored). We emphasize, however, that this conclusion is driven by the absence of strong support for *any* spin in all but one binary (GW151226). We would arrive at the same nominal conclusion for a comparable number of random draws from a binary population model with perfectly aligned binaries with small BH spins. Future and more informative observations of BH binaries could significantly alter this conclusion.

V. DISCUSSION

In this work, we present concrete examples for how well just a handful of GW measurements can improve our phenomenology of the BH mass and spin distribution. Our examples include real observational data from LIGO's O1 and (an incomplete sample from) O2 observing run, suggesting current observations could be on the cusp of constraining BH spins and maximum masses. We provide simple estimates to understand how well these parameters have been constrained, allowing the reader to extrapolate to larger sample sizes. For example, in the absence of positive support for spin, the upper limit on BH spin will decrease rapidly, allowing us to place strong upper limits for (or enable discovery of) BH natal spin.

Because each empirical marginal distribution possesses an infinite number of degrees of freedom (d.o.f.), any phenomenological parametrization such as our own can quickly be exhausted by the data [23], particularly when the population must reproduce multiple observational features. In the short run, therefore, we anticipate a fully generic and regularized infinite-dimensional approach will soon be required to adequately reproduce the thousands of events that even the current generation of instruments will discover. A fully generic approach, however, can easily be misled, not least because GW measurements are subject to many subtle strong-field systematics due to model incompleteness. For example, a waveform approximation widely used for rapid parameter inference of binary black holes (IMRPv2 [94]) omits astrophysically critical d.o.f.—the calculation allows for only one precessing spin instead of the two necessary to fully describe the dynamics—and demonstrably has systematic errors large enough to shift posterior distributions for O3-scale events by an appreciable fraction of their statistically expected extent [49,79]. To illustrate the pernicious impact of these systematic biases, we can consider a simple order-of-magnitude estimate: a single quantity, with intrinsic Gaussian distribution of mean μ and width σ , being observed multiple times by an apparatus with a (Gaussian, random) measurement error Δx and bias δx . The bias will be

important when it influences our best estimate of the average (i.e., when $\delta x \gtrsim \sqrt{\sigma^2 + \Delta x^2}/\sqrt{N}$). Applying this order-of-magnitude approach to GW measurements, we expect that after only a few tens of binary mergers, these modeling systematics will progressively contaminate the interpretation of coalescing binaries, as posterior biases in each event become reflected in biases in the inferred population distribution. Waveform systematics will be even more important because BH spins appear to be small: greater accuracy is needed to separate the secular effects of spin. In this work, when carrying out a full parameter inference, we use the newly developed RIFT parameter inference engine [49] to produce posteriors. We will discuss the impact of waveform systematics on BH spin misalignment measurements in future work.

VI. CONCLUSIONS

We have introduced a flexible, ready-to-use, and self-consistent parametric method to estimate the compact binary merger rate as a function of binary parameters, specifically emphasizing mass and spin. Unlike prior work, our procedure self-consistently estimates the merger rate and binary parameter distribution, accounting for statistical sampling error, measurement error, and selection bias. Using this procedure, we show by example that only a handful of NS-NS and BH-BH measurements can enable strong constraints on their respective populations via GW observations alone. Even in the astrophysically likely scenario of small BH spin, we emphasize that just a few measurements will enable sharp constraints on the BH spin distribution. Interpreting current observations, we show that GW measurements are already beginning to place astrophysically interesting constraints on the spin of BHs. We reproduce prior results about the lack of reported BHs at high mass and its implications for the BH mass spectrum. Finally, particularly in our Appendix, we explain how to extrapolate toward the measurement prospects available in the very near future.

The procedure described here assumes all sources have been unambiguously resolved from observational data, omitting any treatment of source significance aside from a naive selection bias. Farr *et al.* [95] demonstrated and popularized an approach to self-consistently perform the detection and population inference process, estimating the foreground and background distributions simultaneously; see also [96–98]. Recently, Gaebel and collaborators [99] developed a concrete procedure to apply this technique to gravitational wave observations. Owing to many deep similarities between our strategies, we anticipate we will shortly incorporate this technique in our own analysis.

The approach described here also employs several strong assumptions about the (lack of) correlations between model parameters. For example, our fiducial BH model assumes the mass-dependent BH merger rate is independent of redshift; that BH masses and spins are completely independent; and that BH spin misalignment and spin

magnitudes are likewise uncorrelated. We will explore more physically motivated correlations in future work.

In the long run, phenomenology is only as sound as the underlying parametrization. Previous analyses have repeatedly shown that adopting an overly restrictive model will produce biased results, as demonstrated by Fishbach *et al.* (with the maximum mass) [21] and Talbot *et al.* [22] (with the shape of the maximum mass cutoff). With sufficient data, a suitably regularized infinite-dimensional parametrization will make unintended systematic biases less frequent. Mature methods for infinite-dimensional or non-parametric inference exist [100–102], beginning with simple infinite-dimensional parametrizations plus smoothing priors or with Gaussian processes [103]. Early investigations have applied nonparametric methods to GW population estimates [29,30]. However, because the GW signal is so rich, many parameters can be measured for each event, several of which are believed to be correlated in most astrophysical formation scenarios. These correlations should be more sharply identified with strong theoretical priors for the immediate future.

Finally, several technical improvements can make this approach faster and more robust. For example, we can perform inference on all events simultaneously, using direct estimates of the likelihood $\ell(\lambda)$ naturally reported by RIFT, to ensure any population inferences are not limited by the compact support of fiducial priors. Using accelerated general-purpose inference engines, we expect to dramatically accelerate the speed with which our population inferences are provided, with a long-term goal of enabling low-latency population-informed identification and classification of candidate sources.

ACKNOWLEDGMENTS

The authors appreciate the opportunities to talk about this work during its development with Maya Fishbach, Tom Dent, Jonah Kanner, Will Farr, Colm Talbot, Eric Thrane, and Salvatore Vitale. R. O. S., J. L., and D. W. gratefully acknowledge NSF Grant No. PHY-1707965. D. W. also acknowledges support from the Rochester Institute of Technology through the Frontiers in Gravitational Wave Astrophysics (FGWA) Signature Interdisciplinary Research Areas (SIRA) initiative. The authors thank the LIGO Scientific Collaboration for access to the data and gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, and the Max-Planck-Society (MPS) for support of the construction of Advanced LIGO. Additional support for Advanced LIGO was provided by the Australian Research Council. For their use in building the PopModels package, we would like to acknowledge Numpy and Scipy [104], Emcee [57], Matplotlib [105], AstroPy [106,107], and h5py [108].

APPENDIX A: MOCK POSTERIOR POPULATIONS PRECESSING BINARIES: ALIGNED FISHER MATRIX APPROACH

We test our code using synthetic or “mock” posterior distributions for binary black hole parameters, designed to mimic the results of full end-to-end Bayesian inference on synthetic data. For the mock BBH posterior distributions constructed in this work, we adopt a very simple approximation, motivated by decades of experience suggesting that for short BBH signals the likelihood for gravitational wave signals is nearly Gaussian in three coordinates $(\mathcal{M}_c, \eta, \chi_{\text{eff}})$ and does not strongly constrain any other d.o.f. Specifically, if λ_0 are the true binary parameters and ρ is the true network signal amplitude; if $\Gamma_{ab} = \langle \partial_a h | \partial_b h \rangle$ is the Fisher matrix for the binary parameters λ , evaluated at $\lambda = \lambda_0$ and for a signal amplitude ρ using a fiducial detector power spectrum; and if $p(\lambda)$ is the prior distribution on λ , then we approximate the posterior distribution by a distribution proportional to

$$e^{-\Gamma_{ab}(\lambda-\lambda_*)_a(\lambda-\lambda_*)_b} p_{\text{ref}}(\lambda), \quad (\text{A1})$$

where λ_* is a fixed random realization from a normal distribution with mean λ_0 and covariance matrix Γ^{-1} . We generate samples from this distribution via Monte Carlo techniques. We evaluate the approximate Fisher matrix Γ using the effective Fisher technique [55,82,109], applied to a nonprecessing binary waveform model assigned the same values of $\mathcal{M}_c, \eta, \chi_{\text{eff}}$ (i.e., via $\chi_{1,z} = \chi_{2,z} = \chi_{\text{eff}}$).

This approximate posterior distribution has several distinct advantages. First and foremost, it captures in Γ_{ab} the strong, parameter-dependent, and well-understood correlations between the variables that most significantly impact the GW inspiral signal, while simultaneously populating all intrinsic binary parameters. For example, it captures the shape of the posterior distribution in mass ratio and spin while correctly accounting for parameter boundary effects, as described in [110]. Second, it accounts via λ_* for the effect of random noise realizations, which impact the best-fitting parameters associated with each set of synthetic data. By including an explicit prior $p_{\text{ref}}(\lambda)$, it allows us to carefully adopt fiducial prior assumptions, which have a substantial impact on inferred binary masses and spins.

A ready-to-use implementation of this algorithm is available.⁴

For simplicity, in this implementation, no cosmological effects are applied. If used unaltered, this approximate posterior applies either if cosmological redshift effects are small compared to the width of the distribution in mass (i.e., bias is small compared to the statistical uncertainty) or if these ambiguity distributions are used to approximate the source-frame ambiguity function. Cosmological effects

⁴See <https://git.ligo.org/daniel.wysocki/synthetic-PE-posteriors>.

dominate the accuracy to which a binary neutron star's chirp mass can be measured; to be used in such a scenario, this approximation must be refined to reflect the significant impact of the sources' unknown redshift.

APPENDIX B: MOCK POPULATIONS

To generate a synthetic population of events, we employ the following procedure. Using O1 sensitivity, and a detection criterion of $\rho > 8$ in a single interferometer, we used our estimate of V and a fiducial observation time T to compute the expected number of events μ . Using the Poisson distribution, we select a total number of events N to observe. We assumed each detected binary had a network SNR drawn from a power law $p(\rho_{\text{network}}) \propto \rho_{\text{network}}^{-4}$, with a lower cutoff of 12 (roughly corresponds to 8 in two detectors).

APPENDIX C: OVERVIEW OF KEY PHENOMENOLOGICAL CONSTRAINTS

1. How well can we measure distribution hyperparameters?

Classical frequentist statistical methods provide a quick way to assess how rapidly observations will constrain model hyperparameters. For example, the sample mean of maximum likelihood estimators converges rapidly to the true mean, and (to a first approximation) the sample variance is approximately χ^2 distributed. Thus, by adopting the mean and variance of our underlying distributions as coordinates on the space Λ of hyperparameters, we can estimate how efficiently observations will constrain them. For example, if we account for measurement error, we can measure the mean spin to an accuracy $\sqrt{V(\chi) + \sigma_\chi^2}/\sqrt{N}$ where $V(\chi)$ is the variance of the spin magnitude distribution and σ_χ is the typical spin measurement accuracy for the mass range of interest [typically $\mathcal{O}(0.3)$]. Because of sharp cutoffs, the maximum and minimum masses have a qualitatively different behavior; see, e.g., [111]. Both the maximum and minimum masses are best estimated using the most extreme individual event, with an accuracy converging as $1/N$. In our context—the power-law mass distribution—the accuracy with which these maximum masses can be determined scales directly with the number of events in a given region. We therefore expect the maximum mass can be determined to an accuracy of order m_{max}/N ; the appropriate scale factor can be calibrated to detailed analyses of the kind performed in Sec. III. Similarly, as described below in Appendix C 2, we can use the observed range of χ_{eff} to constrain spin magnitudes and misalignments.

While providing a useful order-of-magnitude estimate into how well we can measure distribution parameters, the simple estimates above become cumbersome when trying to capture correlations between our phenomenological parameters, notably the event rate and mass distribution.

Following [23], we assess how well we can distinguish model hyperparameters from the (expected) log-likelihood as a function of model hyperparameters Λ of

$$\langle \ln \mathcal{L} \rangle = -\mu_* + \mu_* \left\langle \ln \int d\lambda p(d|\lambda) R p(\lambda|\Lambda) \right\rangle_*, \quad (\text{C1})$$

where the expectation is performed relative to some reference model characterized by parameters Λ_* such that $p_*(\lambda) \equiv p(\lambda|\Lambda_*)$ and $\mu_* = \mu(\Lambda_*)$. Rather than work in full generality, we perform a Taylor series expansion of the likelihood around the local maximum, characterizing the second order term by its inverse covariance or Fisher matrix Γ_{ab} ,

$$\langle \ln \mathcal{L} \rangle \simeq \ln \mathcal{L}_* - \frac{1}{2} \Gamma_{\alpha\beta} (\Lambda - \Lambda_*)_\alpha (\Lambda - \Lambda_*)_\beta. \quad (\text{C2})$$

If γ_k are eigenvalues of Γ , then hyperparameters can be measured to an accuracy $1/\sqrt{\gamma_k}$, which scales as $1/\sqrt{N}$ for N the number of observed events.

We first illustrate this technique in the idealized case of zero measurement error, following previous work [23] which characterized differences between two distributions q, p using the Kullback–Leibler (KL) divergence $D_{\text{KL}}(p|q) \equiv \int p(x) \ln[p(x)/q(x)] dx$. The marginalized log likelihood only depends on model hyperparameters Λ through the KL divergence between our proposed model μ, p (which depends on Λ) and the reference model μ_*, p_* (which does not):

$$\langle \ln \mathcal{L} \rangle = -D_{\text{KL}}(\mu_*|\mu) - \mu_* D_{\text{KL}}(p_*|p) + \text{const}. \quad (\text{C3})$$

As a result, the Fisher matrix has two model-dependent terms, each reflecting second derivatives of D_{KL} with respect to model parameters:

$$\Gamma_{\alpha\beta}^{(\text{zero})} = \Gamma_{\alpha\beta}^{(\mu)} + \mu_* \Gamma_{\alpha\beta}^{(p)}, \quad (\text{C4})$$

where the first term arises from differences in the observed number; where the second term reflects differences in shape; and where we use the fact that D_{KL} has a local minimum (of 0) when the two distributions are equal to eliminate cross terms. Thus, we can evaluate the Fisher matrix simply by computing KL divergences and carrying out the necessary derivatives. For example, for the mass power-law model with fixed mass range, $p(m|\alpha) = C(\alpha, m_+, m_-) m_1^{-\alpha} / (m_1 - m_-)$, the KL divergence $D_{\text{KL}}(p_*, p)$ becomes

$$D_{\text{KL}}(\alpha_*|\alpha) \equiv \int p(x|\alpha_*) \ln p(x|\alpha_*) / p(x|\alpha) \quad (\text{C5})$$

$$= (\alpha - \alpha_*) \langle \ln x \rangle_{\alpha_*} + \ln C(\alpha_*) / C(\alpha), \quad (\text{C6})$$

where the conditional average is $\langle f \rangle_\alpha \equiv \int dx f(x) p(x|\alpha)$. In this expression, only the last term $-\ln C(\alpha)$ does not cancel in $\partial_\alpha^2 D_{\text{KL}}(\alpha_*|\alpha)$.

Again using the same concrete power-law example, we next use this technique to show how, because μ [Eq. (6)] and the mass distribution can be independently constrained, the “overall event rate” \mathcal{R} and the mass distribution are correlated. Representing $\mu = e^X$, the second derivative of $D_{\text{KL}}(\mu_*|\mu)$ becomes [23]

$$D_{\text{KL}} \simeq \frac{1}{2} \mu_* (\partial_a X)(\partial_b X) (\Lambda - \Lambda_*)_a (\Lambda - \Lambda_*)_b. \quad (\text{C7})$$

For the power-law model described above, the only two derivatives needed are $\partial_{\ln R} X = 1$ and $\partial_\alpha X = \partial_\alpha \ln(VT)_\alpha$, the latter of which can be well approximated by -1 . This term introduces correlations between the rate variable ($\ln \mathcal{R}$) and shape (α). Conversely, using coordinates μ and α to characterize the observed population, by construction our inferred posterior distribution on the total number and mass distribution are uncorrelated.

Roughly speaking, the effects of measurement error add in quadrature in the Fisher matrix:

$$\Gamma = \Gamma^{(\text{zero})} + \Gamma^{(\text{measure})}. \quad (\text{C8})$$

We can therefore refine the estimates provided above to incorporate simple estimates of GW measurement errors and their correlations. For the simple power-law estimate described above, however, these measurement errors are relatively small compared to the range of the distribution, unless α is very large.

In the above order-of-magnitude discussion, we have not accounted for parameter-dependent selection bias. To a good first approximation, GW selection bias enters only through the masses, roughly as the (chirp) mass to a power. We can therefore treat the observed population as a (different) power law, which observations constrain to an accuracy loosely characterized by the analysis above.

Therefore, for the power-law mass distribution, we expect the posterior distribution of (log) rate and power-law exponent will be correlated and follow a Gaussian distribution characterized by the inverse covariance

$$\Gamma \simeq \mu \begin{bmatrix} 1 & -1 \\ -1 & 1 + 2\partial_\alpha^2 \ln C(\alpha) \end{bmatrix} \quad (\text{C9})$$

relative to the coordinates $(\ln \mathcal{R}, \alpha)$, if we adopt a uniform prior on α and $\ln R$. This expression captures the correlations between the rate and mass ratio seen in our inferences, when only varying the total event rate and mass ratio.

2. Semianalytic model for constraints on the spin magnitude and misalignment distribution

In this paper, for the purposes of illustration and as a leading-order approximation suitable for the BH-BH binaries reported to date, we adopt three simplifying approximations: that the sensitive volume depends weakly on spin; that GW measurements will only constrain χ_{eff} ; and that the

underlying mass and spin distributions of BH-BH binaries are uncorrelated. In this framework of approximations, only χ_{eff} measurements and hence the underlying χ_{eff} distribution of the population determines how well we can distinguish between population models via spin measurements. Within this framework, we can simply and largely analytically estimate how much information we gain about the BH spin distribution from repeated measurements.

In our synthetic model (and nature) where BH spins appear to be small, the first few measurements will principally inform our upper limit on the BH spin distribution, via the absence of observations consistent with large χ_{eff} . For example, in our synthetic model, the 90% upper limit expected in 25 events is $\chi_{\text{eff}} < 0.31$; for our inferred posterior predictive distribution based on all published events, it is 0.19. In Fig. 11, we use a simple toy model to illustrate how upper limits loosely inform our

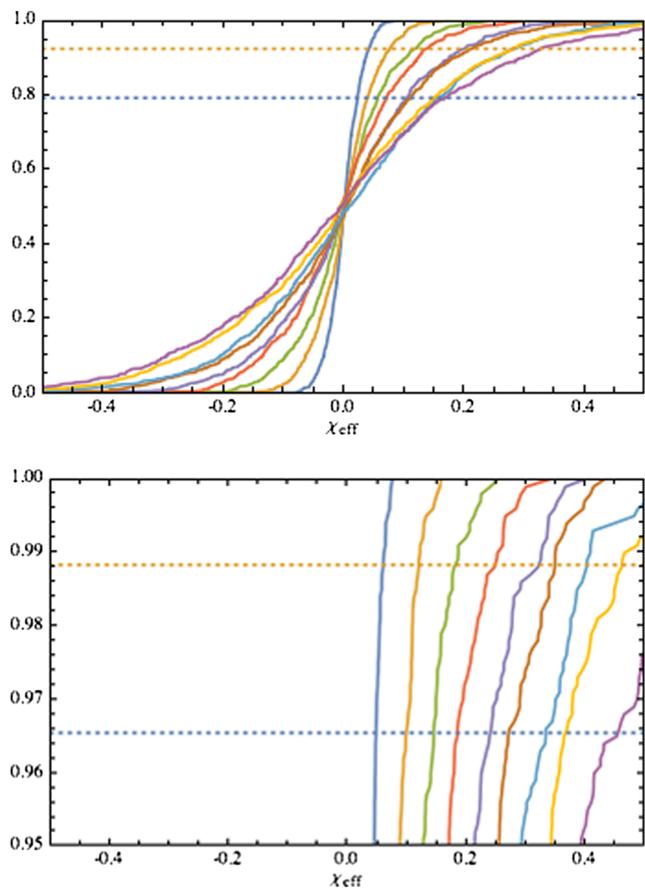


FIG. 11. Why χ_{eff} measurements constrain the maximum spin: cumulative distribution function for χ_{eff} for toy models with isotropic spins and uniform spin magnitude distributions limited by 0.1, 0.2, 0.3, In the top panel, the vertical lines, corresponding to $0.5^{1/3}$ and $0.5^{1/25}$, indicate the locus of points in each cumulative distribution function we can begin to constrain with the absence of events above X with 3 and 25 events, respectively. In the bottom panel, the lines have been changed to $0.9^{1/3}$ and $0.9^{1/25}$, respectively.

estimates of the BH spin distribution. In this model, we assume each BH in a binary has a random spin magnitude drawn from a uniform distribution between 0 and χ_{\max} , randomly (isotropically) oriented, for binaries with a random mass ratio uniformly drawn between 0.1 and 1. This figure shows the cumulative distribution of χ_{eff} implied by these assumptions, for different choices of χ_{\max} . (These cumulative distributions are well approximated by analytic expressions for the cumulative distribution of $\chi_{1,z}$ and χ_{eff} under these assumptions; see [49] for concrete expressions.) For comparison, the vertical shaded regions show the largest values of χ_{eff} which have significant support in our synthetic sample ($\chi_{\text{eff}} \lesssim 0.5$), consistent with the largest plausible spins reported for O1 and O2 events. The lack of support for large χ_{eff} in any observation to date strongly suggests that BH spins cannot be large. Conversely, an observation of a binary with χ_{eff} bounded below by ϵ (e.g., GW151226) implies that a significant fraction of BH spins must be greater than of order ϵ .

We emphasize that we provide these estimates (and perform our calculation within these underlying approximations) to produce a *conservative, well-understood benchmark* for how well the BH spin distribution can be constrained with present and future GW measurements. Real GW measurements, particularly of low-mass or closer and therefore higher-amplitude BH-BH mergers, will provide additional direct constraints on the other spin d.o.f.

APPENDIX D: END-TO-END TESTS OF POPULATION HYPERPARAMETER RECOVERY: P - P PLOTS

A standard technique to test Bayesian parameter inference codes is a probability-probability or P - P plot. We employ this test both on our population inference engine and on the procedure for making synthetic observations. For our population inference code, we generate $k = 1 \dots 1000$ synthetic BBH populations, each a fair draw from a set of population hyperparameters controlling the rate, mass, and spin distribution. For each synthetic population, we generate one random observing run with O1 LIGO sensitivity and $T = 300$ days coincident observing time, by computing the expected number of detections μ [Eq. (6)] and taking one random Poisson draw $p(n_k) \propto e^{-\mu} \mu^{n_k} / n_k!$. We take n_k detection-weighted binaries, generating parameter estimates according to the procedure in Appendix B. We then apply our population parameter inference code to generate posterior distributions on the population hyperparameters Λ , and from that one-dimensional marginal cumulative distributions $\hat{P}_{k,i}(\Lambda_i)$, for each parameter Λ_i . It should be noted here that we used as our prior the same distribution that these population hyperparameters were drawn from, as anything else would produce biases. Using the true hyperparameter values $\Lambda_{k,i}^*$,

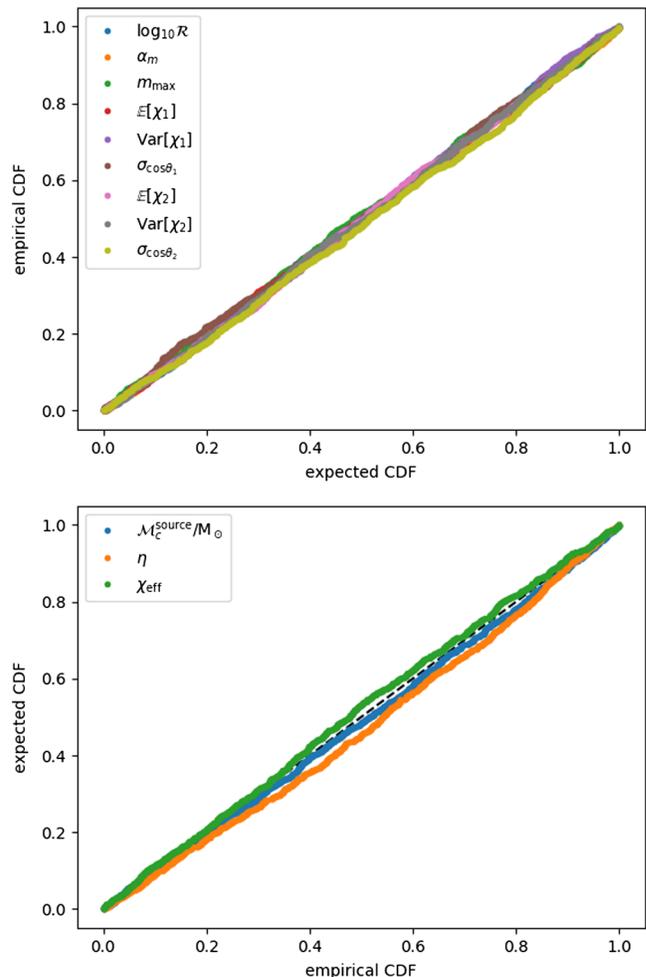


FIG. 12. P - P plots for hyperparameter recovery. Top (bottom) panel shows the P - P plot for population (single synthetic event) inferences.

we generate a single number for each hyperparameter $\hat{P}_{k,i}(\Lambda_{k,i}^*)$. A P - P plot is the cumulative distribution of these $\hat{P}_{k,i}(\Lambda_{k,i}^*)$. If the code is behaving correctly, these should be uniformly distributed from 0 to 1: the plot should be diagonal. The top panel of Fig. 12 shows the P - P plots for each of our model hyperparameters.

In addition to our population inference code, we made P - P plots for our synthetic parameter estimation code, described in Appendix B, as our population inference tests make use of it. Here we generated $k = 1 \dots 1000$ synthetic BBH signals, drawing true values from the prior we used for measuring the posteriors. We repeated the same process just described, making posterior distributions on the intrinsic parameters λ , and evaluating the marginal cumulative distribution functions at the true values λ_k^* . P - P plots for some representations of the intrinsic parameters are shown in the bottom panel of Fig. 12.

- [1] B. Abbott *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
- [2] T. Accadia *et al.*, *J. Instrum.* **7**, P03012 (2012).
- [3] F. Acernese *et al.* (VIRGO Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
- [4] B. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **833**, L1 (2016).
- [5] B. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **6**, 041015 (2016).
- [6] LIGO Scientific and Virgo Collaborations, *Astrophys. J. Lett.* **818**, L22 (2016).
- [7] I. Mandel and S. E. de Mink, *Mon. Not. R. Astron. Soc.* **458**, 2634 (2016).
- [8] P. Marchant, N. Langer, P. Podsiadlowski, T. Tauris, and T. Moriya, *Astron. Astrophys.* **588**, A50 (2016).
- [9] C. L. Rodriguez, C.-J. Haster, S. Chatterjee, V. Kalogera, and F. A. Rasio, *Astrophys. J. Lett.* **824**, L8 (2016).
- [10] S. Bird, I. Cholis, J. B. Muñoz, Y. Ali-Haïmoud, M. Kamionkowski, E. D. Kovetz, A. Raccanelli, and A. G. Riess, *Phys. Rev. Lett.* **116**, 201301 (2016).
- [11] D. Kushnir, M. Zaldarriaga, J. A. Kollmeier, and R. Waldman, *Mon. Not. R. Astron. Soc.* **462**, 844 (2016).
- [12] A. Lamberts, S. Garrison-Kimmel, D. R. Clausen, and P. F. Hopkins, *Mon. Not. R. Astron. Soc.* **463**, L31 (2016).
- [13] I. Dvorkin, E. Vangioni, J. Silk, J.-P. Uzan, and K. A. Olive, *Mon. Not. R. Astron. Soc.* **461**, 3877 (2016).
- [14] B. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **116**, 131102 (2016).
- [15] L. Barack, V. Cardoso, S. Nissanke, T. P. Sotiriou, A. Askar, C. Belczynski, G. Bertone, E. Bon *et al.*, *Classical Quantum Gravity* **36**, 143001 (2019).
- [16] D. Wysocki, D. Gerosa, R. O’Shaughnessy, K. Belczynski, W. Gladysz, E. Berti, M. Kesden, and D. E. Holz, *Phys. Rev. D* **97**, 043014 (2018).
- [17] I. Mandel and R. O’Shaughnessy, *Classical Quantum Gravity* **27**, 114007 (2010).
- [18] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, *Astrophys. J. Lett.* **832**, L2 (2016).
- [19] K. Breivik, C. L. Rodriguez, S. L. Larson, V. Kalogera, and F. A. Rasio, *Astrophys. J. Lett.* **830**, L18 (2016).
- [20] A. Nishizawa, E. Berti, A. Klein, and A. Sesana, *Phys. Rev. D* **94**, 064020 (2016).
- [21] M. Fishbach and D. E. Holz, *Astrophys. J. Lett.* **851**, L25 (2017).
- [22] C. Talbot and E. Thrane, *Astrophys. J.* **856**, 173 (2018).
- [23] R. O’Shaughnessy, *Phys. Rev. D* **88**, 084061 (2013).
- [24] S. Stevenson, F. Ohme, and S. Fairhurst, *Astrophys. J.* **810**, 58 (2015).
- [25] K. Belczynski, D. E. Holz, T. Bulik, and R. O’Shaughnessy, *Nature* **534**, 512 (2016).
- [26] M. Zevin, C. Pankow, C. L. Rodriguez, L. Sampson, E. Chase, V. Kalogera, and F. A. Rasio, *Astrophys. J.* **846**, 82 (2017).
- [27] J. W. Barrett, S. M. Gaebel, C. J. Neijssel, A. Vigna-Gómez, S. Stevenson, C. P. L. Berry, W. M. Farr, and I. Mandel, *Mon. Not. R. Astron. Soc.* **477**, 4685 (2018).
- [28] A. Miyamoto, T. Kinugawa, T. Nakamura, and N. Kanda, *Phys. Rev. D* **96**, 064025 (2017).
- [29] D. Wysocki, arXiv:1712.02643.
- [30] I. Mandel, W. M. Farr, A. Colonna, S. Stevenson, P. Tiño, and J. Veitch, *Mon. Not. R. Astron. Soc.* **465**, 3254 (2017).
- [31] E. D. Kovetz, I. Cholis, P. C. Breysse, and M. Kamionkowski, *Phys. Rev. D* **95**, 103010 (2017).
- [32] C. Talbot and E. Thrane, *Phys. Rev. D* **96**, 023012 (2017).
- [33] W. M. Farr, S. Stevenson, M. C. Miller, I. Mandel, B. Farr, and A. Vecchio, *Nature (London)* **548**, 426 (2017).
- [34] D. Gerosa and E. Berti, *Phys. Rev. D* **95**, 124046 (2017).
- [35] M. Fishbach, D. E. Holz, and B. Farr, *Astrophys. J. Lett.* **840**, L24 (2017).
- [36] S. Stevenson, C. P. L. Berry, and I. Mandel, *Mon. Not. R. Astron. Soc.* **471**, 2801 (2017).
- [37] S. Vitale, R. Lynch, V. Raymond, R. Sturani, J. Veitch, and P. Graff, *Phys. Rev. D* **95**, 064053 (2017).
- [38] B. Farr, D. E. Holz, and W. M. Farr, *Astrophys. J. Lett.* **854**, L9 (2018).
- [39] S. Vitale, R. Lynch, R. Sturani, and P. Graff, *Classical Quantum Gravity* **34**, 03LT01 (2017).
- [40] R. O’Shaughnessy, D. Gerosa, and D. Wysocki, *Phys. Rev. Lett.* **119**, 011101 (2017).
- [41] V. Kalogera, *Astrophys. J.* **541**, 319 (2000).
- [42] K. Belczynski, M. Dominik, T. Bulik, R. O’Shaughnessy, C. L. Fryer, and D. E. Holz, *Astrophys. J. Lett.* **715**, L138 (2010).
- [43] R. O’Shaughnessy, R. Kopparapu, and K. Belczynski, *Classical Quantum Gravity* **29**, 145011 (2012).
- [44] M. Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, T. Bulik, I. Mandel, and R. O’Shaughnessy, *Astrophys. J.* **759**, 52 (2012).
- [45] D. Wysocki, D. Gerosa, R. O’Shaughnessy, K. Belczynski, W. Gladysz, E. Berti, M. Kesden, and D. E. Holz, *Phys. Rev. D* **97**, 043014 (2018).
- [46] T. Damour, *Phys. Rev. D* **64**, 124013 (2001).
- [47] É. Racine, *Phys. Rev. D* **78**, 044021 (2008).
- [48] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brüggemann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaría, and J. Seiler, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [49] J. Lange, R. O’Shaughnessy, and M. Rizzo, arXiv:1805.10457.
- [50] LIGO Scientific and Virgo Collaborations, <https://dcc.ligo.org/LIGO-P1500218>.
- [51] J. Veitch, V. Raymond, B. Farr, W. M. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
- [52] B. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **94**, 064035 (2016).
- [53] C. Cutler and É. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [54] E. Poisson and C. M. Will, *Phys. Rev. D* **52**, 848 (1995).
- [55] H. Cho, E. Ochsner, R. O’Shaughnessy, C. Kim, and C. Lee, *Phys. Rev. D* **87**, 024004 (2013).
- [56] J. Goodman and J. Weare, *Commun. Appl. Math. Comput. Sci.* **5**, 65 (2010).
- [57] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [58] J. Abadie *et al.* (LIGO Scientific and Virgo Collaborations), *Classical Quantum Gravity* **27**, 173001 (2010).
- [59] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993).

- [60] R. O'Shaughnessy, V. Kalogera, and K. Belczynski, *Astrophys. J.* **716**, 615 (2010).
- [61] M. Dominik, E. Berti, R. O'Shaughnessy, I. Mandel, K. Belczynski, C. Fryer, D.E. Holz, T. Bulik, and F. Pannarale, *Astrophys. J.* **806**, 263 (2015).
- [62] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016).
- [63] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044007 (2016).
- [64] D. A. Brown, A. Lundgren, and R. O'Shaughnessy, *Phys. Rev. D* **86**, 064020 (2012).
- [65] R. O'Shaughnessy, B. Vaishnav, J. Healy, and D. Shoemaker, *Phys. Rev. D* **82**, 104006 (2010).
- [66] R. Biswas, P. R. Brady, J. D. E. Creighton, and S. Fairhurst, *Classical Quantum Gravity* **26**, 175009 (2009).
- [67] V. Tiwari, *Classical Quantum Gravity* **35**, 145009 (2018).
- [68] P. Christian, P. Mocz, and A. Loeb, *Astrophys. J. Lett.* **858**, L8 (2018).
- [69] G. S. Fraley, *Astrophys. Space Sci.* **2**, 96 (1968).
- [70] C. L. Fryer, S. E. Woosley, and A. Heger, *Astrophys. J.* **550**, 372 (2001).
- [71] S. E. Woosley, A. Heger, and T. A. Weaver, *Rev. Mod. Phys.* **74**, 1015 (2002).
- [72] S. E. Woosley, S. Blinnikov, and A. Heger, *Nature (London)* **450**, 390 (2007).
- [73] D. Kasen, S. E. Woosley, and A. Heger, *Astrophys. J.* **734**, 102 (2011).
- [74] K. Belczynski, A. Heger, W. Gladysz, A. J. Ruitter, S. Woosley, G. Wiktorowicz, H.-Y. Chen, T. Bulik, R. O'Shaughnessy, D. E. Holz, C. L. Fryer, and E. Berti, *Astron. Astrophys.* **594**, A97 (2016).
- [75] C. L. Rodriguez, P. Amaro-Seoane, S. Chatterjee, and F. A. Rasio, *Phys. Rev. Lett.* **120**, 151101 (2018).
- [76] M. C. Miller and D. P. Hamilton, *Mon. Not. R. Astron. Soc.* **330**, 232 (2002).
- [77] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.*, *Phys. Rev. Lett.* **116**, 241103 (2016).
- [78] S. Vitale, D. Gerosa, C.-J. Haster, K. Chatziioannou, and A. Zimmerman, *Phys. Rev. Lett.* **119**, 251103 (2017).
- [79] A. R. Williamson, J. Lange, R. O'Shaughnessy, J. A. Clark, P. Kumar, J. Calderón Bustillo, and J. Veitch, *Phys. Rev. D* **96**, 124041 (2017).
- [80] F. Özel and P. Freire, *Annu. Rev. Astron. Astrophys.* **54**, 401 (2016).
- [81] T. Dietrich, S. Khan, R. Dudi, S. J. Kapadia, P. Kumar, A. Nagar, F. Ohme, F. Pannarale, A. Samajdar, S. Bernuzzi, G. Carullo, W. Del Pozzo, M. Haney, C. Markakis, M. Pürrer, G. Riemenschneider, Y. E. Setyawati, K. W. Tsang, and C. Van Den Broeck, *Phys. Rev. D* **99**, 024029 (2019).
- [82] R. O'Shaughnessy, B. Farr, E. Ochsner, H.-S. Cho, C. Kim, and C.-H. Lee, *Phys. Rev. D* **89**, 064048 (2014).
- [83] M. Hannam, D. A. Brown, S. Fairhurst, C. L. Fryer, and I. W. Harry, *Astrophys. J. Lett.* **766**, L14 (2013).
- [84] X. Zhu, E. Thrane, S. Osłowski, Y. Levin, and P. D. Lasky, *Phys. Rev. D* **98**, 043002 (2018).
- [85] L. Wade, J. D. E. Creighton, E. Ochsner, B. D. Lackey, B. F. Farr, T. B. Littenberg, and V. Raymond, *Phys. Rev. D* **89**, 103012 (2014).
- [86] M. Favata, *Phys. Rev. Lett.* **112**, 101101 (2014).
- [87] B. D. Lackey and L. Wade, *Phys. Rev. D* **91**, 043002 (2015).
- [88] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 011001 (2019).
- [89] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari *et al.*, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [90] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya *et al.*, *Phys. Rev. Lett.* **118**, 221101 (2017).
- [91] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **851**, L35 (2017).
- [92] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **119**, 141101 (2017).
- [93] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **119**, 141101 (2017).
- [94] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [95] W. M. Farr, J. R. Gair, I. Mandel, and C. Cutler, *Phys. Rev. D* **91**, 023005 (2015).
- [96] T. J. Loredo, *AIP Conf. Proc.*, **735**, 195 (2004).
- [97] J. Buchner, A. Georgakakis, K. Nandra, M. Brightman, M.-L. Menzel, Z. Liu, L.-T. Hsu, M. Salvato, C. Rangel, J. Aird, A. Merloni, and N. Ross, *Astrophys. J.* **802**, 89 (2015).
- [98] C. Messenger and J. Veitch, *New J. Phys.* **15**, 053027 (2013).
- [99] S. M. Gaebel, J. Veitch, T. Dent, and W. M. Farr, *Mon. Not. R. Astron. Soc.* **484**, 4008 (2019).
- [100] A. Gelman and J. Carlin, *Bayesian Data Analysis*, 3rd ed Chapman & Hall/CRC Texts in Statistical Science (Taylor & Francis, 2013).
- [101] P. Orbanz and Y. Teh, *Bayesian nonparametric models*, Encyclopedia of Machine Learning (Springer, 2010) pp. 81-89.
- [102] S. Ghosal and A. van der Vaart, *Fundamentals of Nonparametric Bayesian Inference*, Cambridge Series in Statistical and Probabilistic Mathematics (Cambridge University Press, 2017).
- [103] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, 2006).
- [104] E. Jones, T. Oliphant, P. Peterson *et al.*, Available at <https://www.scipy.org/>.
- [105] J. D. Hunter, *Computing in Science & Engineering* (IEEE, New York, 2007).
- [106] T. P. Robitaille, E. J. Tollerud, P. Greeneld, M. Droettboom, E. Bray, T. Aldcroft, M. Davis, A. Ginsburg, A. M. Price-Whelan *et al.* (Astropy Collaboration), *Astron. Astrophys.* **558**, A33 (2013).

- [107] A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, D. L. Shupe, M. W. Craig, N. Dencheva, A. Ginsburg *et al.*, *Astron. J.* **156**, 18 (2018).
- [108] A. Collete, *Python and HDF5* (O'Reilly Media, Inc., 2013).
- [109] H.-S. Cho and C.-H. Lee, *Classical Quantum Gravity* **31**, 235009 (2014).
- [110] K. K. Y. Ng, S. Vitale, A. Zimmerman, K. Chatziioannou, D. Gerosa, and C.-J. Haster, *Phys. Rev. D* **98**, 083007 (2018).
- [111] S.-i. Amari and H. Nagaoka, *Methods of Information Geometry* (American Mathematical Society, Providence, 2007), Vol. 191.