

Cover Your Basis: Comprehensive Data-driven Characterization of the Binary Black Hole Population

Bruce Edelman¹ , Ben Farr¹ , and Zoheyr Doctor² 

¹ Institute for Fundamental Science, Department of Physics, University of Oregon, Eugene, OR 97403, USA; bedelman@uoregon.edu

² Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Department of Physics and Astronomy, Northwestern University, Evanston, IL 60201, USA

Received 2022 October 24; revised 2022 December 10; accepted 2023 January 2; published 2023 March 22

Abstract

We introduce the first complete nonparametric model for the astrophysical distribution of the binary black hole (BBH) population. Constructed from basis splines, we use these models to conduct the most comprehensive data-driven investigation of the BBH population to date, simultaneously fitting nonparametric models for the BBH mass ratio, spin magnitude and misalignment, and redshift distributions. With GWTC-3, we report the same features previously recovered with similarly flexible models of the mass distribution, most notably the peaks in merger rates at primary masses of $\sim 10M_{\odot}$ and $\sim 35M_{\odot}$. Our model reports a suppressed merger rate at low primary masses and a mass-ratio distribution consistent with a power law. We infer a distribution for primary spin misalignments that peaks away from alignment, supporting conclusions of recent work. We find broad agreement with the previous inferences of the spin magnitude distribution: the majority of BBH spins are small ($a < 0.5$), the distribution peaks at $a \sim 0.2$, and there is mild support for a nonspinning subpopulation, which may be resolved with larger catalogs. With a modulated power law describing the BBH merger rate's evolution in redshift, we see hints of the rate evolution either flattening or decreasing at $z \sim 0.2\text{--}0.5$, but the full distribution remains entirely consistent with a monotonically increasing power law. We conclude with a discussion of the astrophysical context of our new findings and how nonparametric methods in gravitational-wave population inference are uniquely poised to complement to the parametric approach as we enter the data-rich era of gravitational-wave astronomy.

Unified Astronomy Thesaurus concepts: [Gravitational-wave astronomy \(675\)](#); [Gravitational waves \(678\)](#); [Black holes \(162\)](#); [Compact objects \(288\)](#); [High energy astrophysics \(739\)](#)

1. Introduction

Observations of gravitational waves (GWs) from compact binary mergers are becoming a regular occurrence, producing a catalog of events that recently surpassed 90 such detections (Abbott et al. 2019a, 2021a; The LIGO Scientific Collaboration et al. 2021a). As the catalog continues to grow, so does our understanding of the underlying astrophysical population of compact binaries (Abbott et al. 2019b, 2021b; The LIGO Scientific Collaboration et al. 2021b). Following numerous improvements to the detectors since the last observing run, the anticipated sensitivities for the upcoming fourth observing run of the LIGO–Virgo–KAGRA (LVK) Collaboration suggest detection rates as high as once per day (LIGO Scientific Collaboration et al. 2015; Acernese et al. 2015; Abbott et al. 2020a; Akutsu et al. 2021). With the formation history of these dense objects encoded in the details of their distribution (Rodriguez et al. 2016; Vitale et al. 2017b; Farr et al. 2017; Zevin et al. 2017; Farr et al. 2018), the likely doubling in size of the catalog with the next observing run could provide another leap in our understanding of compact binary astrophysics. Beyond formation physics, population-level inference of the compact binary catalog has also been shown to provide novel measurements of cosmological parameters (Farr et al. 2019; The LIGO Scientific Collaboration et al. 2021c; Ezquiaga & Holz 2022), constrain modified gravitational-wave propagation (Finke et al. 2022; Mancarella et al. 2022;

Okounkova et al. 2022), constrain a running Planck mass (Lagos et al. 2019), search for evidence of ultralight bosons through superradiance (Ng et al. 2021a, 2021b), constrain stellar nuclear reaction rates (Farmer et al. 2019, 2020), look for primordial black holes (Ng et al. 2021c, 2022), and to constrain physics of neutron stars (Landry & Read 2021; Golomb & Talbot 2022a). Through a better understanding of the mass, spin, and redshift distributions of compact binaries that will come with the increased catalog size, one can probe a wide range of different physical phenomena with even greater fidelity.

The binary black hole (BBH) mass distribution was first found to have structure beyond a smooth power law with simpler parametric models, exhibiting a possible high-mass truncation and either a break or a peak at $m_1 \sim 35\text{--}40 M_{\odot}$ (Fishbach & Holz 2017; Talbot & Thrane 2018; Abbott et al. 2019b, 2021b). Starting with the moderately sized catalog, GWTC-2, more flexible models found signs of additional structure (Tiwari & Fairhurst 2021; Edelman et al. 2022). The evidence supporting these features, such as the peak at $m_1 \sim 10 M_{\odot}$, has only grown after analyzing the latest catalog, GWTC-3, with the same models (The LIGO Scientific Collaboration et al. 2021b; Tiwari 2022). While this shows the usefulness of data-driven methods with the current relatively small catalog size, they will become more powerful with more observations. The canonical approach to constructing population models has been to use simple parametric descriptions (e.g., power laws, beta distributions) that aim to describe the data in the simplest way, employ astrophysically motivated priors where appropriate, then sequentially add complexities (e.g., Gaussian peaks) as the data demand. This simple approach was necessary when data were



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

scarce, but as we move into the data-rich catalog era, this approach is already failing to scale. More flexible and scalable methods, such as the nonparametric modeling techniques presented in this article, will be necessary to continue to extract the full information contained in the compact binary catalog. In contrast to parametric models, flexible and nonparametric models are data-driven and contribute little bias to functional form. They additionally are particularly useful to search for unexpected features in the data, providing meaningful insight into features that parametric models may fail to capture.

While we eventually hope to uncover hints of binary formation mechanisms in the mass spectrum of BBHs, the distribution of spin properties have been of particular interest. The measurements of spin properties of individual binaries often have large uncertainties, but the theorized formation channels are expected to produce distinctly different spin distributions (Rodriguez et al. 2016; Farr et al. 2017; Zevin et al. 2017; Farr et al. 2018; Gerosa et al. 2018). Isolated (or field) formation scenarios predict component spins that are preferentially aligned with the binary’s orbital angular momentum, although some small misalignment can occur depending on the nature of the supernova kicks as each star collapses to a compact object (Bavera et al. 2020, 2021; Zevin & Bavera 2022). Alternatively, dynamical formation in dense environments where many-body interactions between compact objects can result in binary formation and hardening (shrinking of binary orbits) should produce binaries with components’ spins distributed isotropically (Rodriguez et al. 2016, 2019). BBH spins have also been of controversial interest recently, with different parametric approaches to modeling the spin distribution coming to different conclusions. Studies have disagreed on the possible existence of a significant zero-spinning subpopulation, as well as the presence of significant spin misalignment (i.e., $\cos \theta_i < 0.0$) (The LIGO Scientific Collaboration et al. 2021b; Galauadage et al. 2021; Roulet et al. 2021; Callister et al. 2022; Tong et al. 2022). Another study recently showed that inferences of spin misalignment (or tilts) are sensitive to modeling choices and may not peak at perfectly aligned spins, as is often assumed (Vitale et al. 2022). While enlightening, these recent efforts to improve BBH spin models continue to build sequentially on previous parametric descriptions (Galauadage et al. 2021; Callister et al. 2022; Vitale et al. 2022). To ensure we are extracting the full detail the catalog has to offer, we extend our previous nonparametric modeling techniques to include spin magnitudes and tilts, as well as the binary mass ratio and redshift. The work of Golomb & Talbot (2022b) was released concurrently with this work (based on our previous work, Edelman et al. 2022), and they find similar conclusions on the spin distribution when applying similar flexible models constructed with cubic splines. The work presented in this article, however, does not need to analyze a suite of different model configurations, and it includes flexible nonparametric models for each of the mass, spin, and redshift distributions, rather than spin alone.

Polynomial splines have been applied with success across different areas of gravitational-wave astronomy. They have been used to model the gravitational-wave data noise spectrum, detector calibration uncertainties, coherent gravitational waveform deviations, and modulations to a power-law mass distribution (Farr et al. 2015; Littenberg & Cornish 2015; Edwards et al. 2018; Edelman et al. 2021, 2022). In this paper, we highlight how the use of basis splines can provide a

powerful nonparametric modeling approach to the astrophysical distributions of compact binaries. We illustrate how one can efficiently model both the mass and spin distributions of merging compact binaries in GWTC-3 with basis splines to infer compact binary population properties using hierarchical Bayesian inference. We discuss our results in the context of current literature on compact-object populations, and how this method complements the simpler lower-dimensional parametric models in the short run—and will become necessary with future catalogs. Should they appear with more observations, this data-driven approach will provide checks of our understanding by uncovering more subtle—and potentially unexpected—features. The rest of this article is structured as follows: a description of the background of basis splines in Section 2, followed by a presentation of the results of our extensive, data-driven study of the mass and spin distributions of BBHs in GWTC-3 in Section 3. We then discuss these results and their astrophysical implications in Section 4, and we finish with our conclusions in Section 5.

2. Building the Model

We construct our data-driven model with the application of basis splines, or B-splines (de Boor 1978). B-splines of order k are a set of order k polynomials that span the space of possible spline functions interpolated from a given set of knot locations. For all B-splines used in our model, we use a third-order basis, which consists of individual cubic polynomials. The basis representation of the splines allows for the computationally expensive interpolation to be done in a single preprocessing step—amortizing the cost of model evaluation during inference. To mitigate the unwanted side effects of adding extra knots, and to avoid running a model grid of differing numbers of knots (as in Edelman et al. 2022), we use the smoothing prior for Bayesian P-splines (Lang & Brezger 2004; Jullion & Lambert 2007; Eilers & Marx 2021), allowing the data to pick the optimal scale needed to fit the present features. We discuss basis splines, the smoothing prior, and our specific prior choices on hyperparameters in Appendices A, B, and D.

We parameterize each binary’s mass with the primary (more massive component) mass (m_1) and the mass ratio ($q = m_2/m_1$) with support from 0 to 1. Furthermore, we model four of the six total spin degrees of freedom of a binary merger: the component spin magnitudes a_1 and a_2 , and (the cosines of) the tilt angles of each component, $\cos \theta_1$ and $\cos \theta_2$. The tilt angle is defined as the angle between each component’s spin vector and the binary’s orbital angular momentum vector. We assume the polar spin angles are uniformly distributed in the orbital plane. For the primary mass distribution, we model the log probability with a B-spline interpolated over knots linearly spaced in $\log(m_1)$ from a minimum black hole mass, which we fix to $5M_\odot$, and a maximum mass, which we set to $100M_\odot$. We then have the hyperprior on primary mass with log probability density $\log(p(m_1|\mathbf{c})) \propto B_{k=3}(\log(m_1)|\mathbf{c})$, where $B_{k=3}$ is the cubic B-spline function with a vector of basis coefficients \mathbf{c} . We follow the same procedure for the models in mass-ratio and spin distributions with knots spaced linearly across each domain so that we have $\log(p(\theta|\mathbf{c})) \propto B_{k=3}(\theta|\mathbf{c})$, where θ can be q , a_1 , a_2 , $\cos \theta_1$, or $\cos \theta_2$. For the spin magnitude and tilt distributions, we construct two versions of the model: first, we model each component’s distribution as independent and identically distributed (IID), where we have a single B-spline model and parameters (coefficients) for each binary spin.

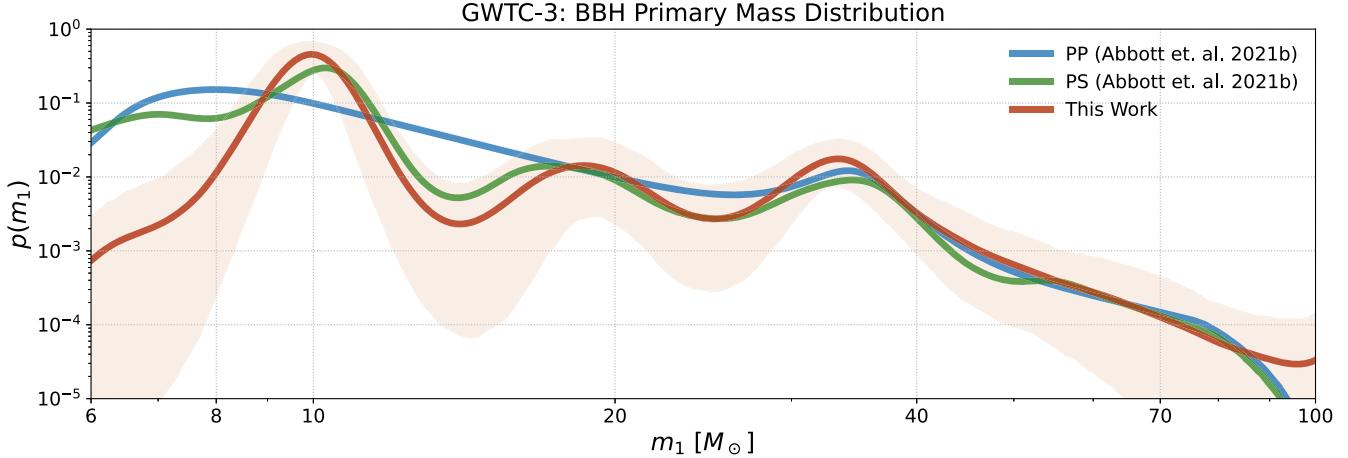


Figure 1. The marginal primary mass distribution inferred with the B-spline model (red), with 64 knots spaced linearly in $\log m_1$, from $5 M_\odot$ to $100 M_\odot$. The solid line shows the population predictive distribution (PPD), and the shaded region the 90% credible interval. We show the inferred PPD from the POWERLAWPEAK (blue) and POWERLAWSPINE (green) models from the LVK’s GWTC-3 population analyses (The LIGO Scientific Collaboration et al. 2021b). The icons link to the code used to generate this figure and to Zenodo entries of any public data used.   

Second, we model each component’s distribution to be unique, fitting separate sets of coefficients for the B-spline models of the primary and secondary spin distributions. Last, we fit a population model on the redshift or luminosity distance distribution of BBHs, assuming a Λ CDM cosmology defined by the parameters from the Planck 2015 results (Planck Collaboration et al. 2016). This defines an analytical mapping between each event’s inferred luminosity distance and its redshift, which we now use interchangeably. We take a semiparametric approach to model the evolution of the merger rate with redshift, following Edelman et al. (2022), that parameterizes modulations to an underlying model with splines (in our case, basis splines). We use the POWERLAWREDSHIFT model as the underlying distribution to modulate, which has a single hyperparameter, λ_z , and a probability density defined as $p(z|\lambda_z) \propto \frac{dV_c}{dz}(1+z)^{\lambda_z-1}$ (Fishbach et al. 2018). For more detailed descriptions of each model and the specific prior choices used for the hyperparameters, see Appendix D. Now that we have our comprehensive data-driven population model built, we simultaneously fit the basis spline models on the BBH masses, spins, and redshift. We use the usual hierarchical Bayesian inference framework (see Appendix C for a review; Abbott et al. 2019b, 2021b; The LIGO Scientific Collaboration et al. 2021b), to perform the most extensive characterization of the population of BBHs to date, using the most recent catalog of GW observations, GWTC-3 (The LIGO Scientific Collaboration et al. 2021a).

3. Binary Black Hole Population Inference with GWTC-3

We use hierarchical Bayesian inference (see Appendix C) to simultaneously infer the astrophysical mass, spin, and redshift distributions of binary black holes (BBHs) given a catalog of gravitational-wave observations. Following the same cut on the recent GWTC-3 catalog done in the LVK’s accompanying BBH population analysis, we have 70 possible BBH mergers with false-alarm rates less than 1 per year (The LIGO Scientific Collaboration et al. 2021a; LVK Collaboration 2021a; The LIGO Scientific Collaboration et al. 2021b). Because it was concluded to be an outlier of the rest of the BBH population in both GWTC-2 and GWTC-3, we choose to omit the poorly

understood event, GW190814 (Abbott et al. 2020b, 2021b; The LIGO Scientific Collaboration et al. 2021b; Essick et al. 2022). This leaves us with a catalog of 69 confident BBH mergers, observed over a period of about 2 yr, from which we want to infer population properties. Following what was done in The LIGO Scientific Collaboration et al. (2021b), for events included in GWTC-1 (Abbott et al. 2019a), we use the published samples that equally weight samples from analyses with the IMRPHENOMPV2 (Hannam et al. 2014) and SEOBNRv3 (Pan et al. 2014; Taracchini et al. 2014) waveforms. For the events from GWTC-2 (Abbott et al. 2021a), we use samples that equally weight all available analyses using higher-order mode waveforms (PRECESSINGIMRPHM). Finally, for new events reported in GWTC-2.1 and GWTC-3 (The LIGO Scientific Collaboration et al. 2021a, 2021d), we use combined samples, equally weighted, from analyses with the IMRPHENOMXPHM (Pratten et al. 2021) and the SEOBNRv4PHM (Ossokine et al. 2020) waveform models. Our study provides the first comprehensive data-driven investigation, simultaneously inferring all the BBH population distributions (i.e., mass, spin, and redshift), uncovering new insights and corroborating those found with other methods. We start with our inference of the mass distribution.

3.1. Binary Black Hole Masses

Figure 1 shows the primary mass distribution inferred with our B-spline model (red), where we see features consistent with those inferred by the POWERLAWPEAK and POWERLAWSPINE mass models (Talbot & Thrane 2018; LVK Collaboration 2021b; Abbott et al. 2021b; The LIGO Scientific Collaboration et al. 2021b; Edelman et al. 2022). In particular, our B-spline model finds peaks in merger rate density as a function of primary mass at both $\sim 10 M_\odot$ and $\sim 35 M_\odot$, agreeing with those reported using the same data set in The LIGO Scientific Collaboration et al. (2021b). The B-spline model finds the same feature at $\sim 18 M_\odot$ as the POWERLAWSPINE model, but remains consistent with the POWERLAWPEAK model; the mass distribution is more uncertain in this region. For each of these features, we find the local maximums occurring at primary masses of $9.9^{+0.67}_{-0.48} M_\odot$, $19^{+3.2}_{-2.3} M_\odot$, and $33^{+2.1}_{-3.0} M_\odot$, all at 90% credibility. We find the largest disagreement at low masses, where the power-law-based

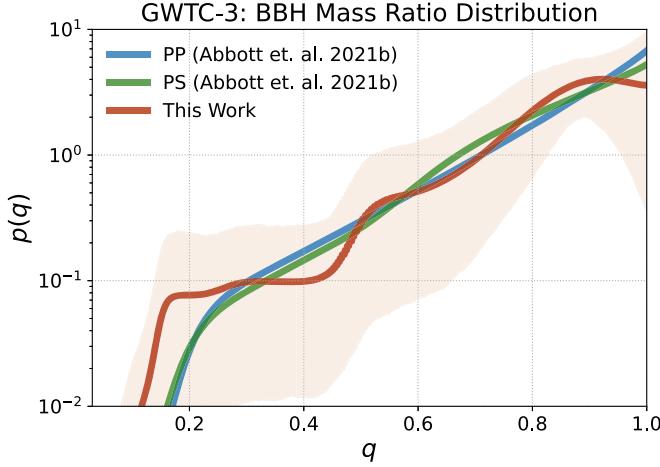


Figure 2. The marginal mass-ratio distribution inferred with the B-spline model (red), with 18 knots spaced linearly in q , from 0.05 to 1. The solid line shows the population predictive distribution (PPD), and the shaded region shows the 90% credible interval. We show the inferred PPD from the POWERLAWPEAK (blue) and POWERLAWSPLINE (green) models from the LVK’s GWTC-3 population analyses (The LIGO Scientific Collaboration et al. 2021b). The icons link to the code used to generate this figure and to Zenodo entries of any public data used. 

models show a higher rate below $\sim 8\text{--}9 M_{\odot}$. This is partly due to the minimum mass hyperparameter (where the power law “begins”) serving as the minimum allowable primary and secondary masses of the catalog. This leads to inferences of m_{\min} below the minimum observed primary mass in the catalog, which is $\sim 6.4 M_{\odot}$, to account for secondary BBH masses lower than that. We choose to fix the minimum black hole mass for both primary and secondaries to $5 M_{\odot}$, similar to the inferred minimum mass in The LIGO Scientific Collaboration et al. (2021b). The lack of observations of binaries with low primary mass makes rate estimates in this region strongly model-dependent, while our flexible model provides an informed upper limit on the rate in this region, given the selection effects and that there are no observations. We could be seeing signs of a decrease in merger rate from a “lower-mass gap” between neutron star and BH masses, or we could be seeing fluctuations due to low-number statistics (van Son et al. 2022a). Either way, we expect this to be resolved with future catalog updates. We also find no evidence for a sharp fall-off in merger rate either following the pileup at $\sim 35 M_{\odot}$ —expected if such a pileup were due to pulsational pair-instability supernovae (PPISNe)—or where the maximum-mass truncations of the power-law models are inferred. The lack of any high-mass truncation, along with the peak at $\sim 35 M_{\odot}$ (significantly lower than expected from PPISNe) may pose challenges for conventional stellar evolution theory. This could be hinting at the presence of subpopulations that avoid pair-instability supernovae during binary formation, but the confirmation of the existence of such subpopulations cannot be determined with the current catalog.

The marginal mass-ratio distribution inferred by the B-spline model is shown in Figure 2. These results suggest we may be seeing the first signs of departure from a simple power-law behavior. We find potential signs of a plateau or decrease in the merger rate near equal mass ratios, as well as a broader tail toward unequal mass ratios than the power-law-based models find, although a smooth power law is still consistent with these results, given the large uncertainties. Our results also suggest a

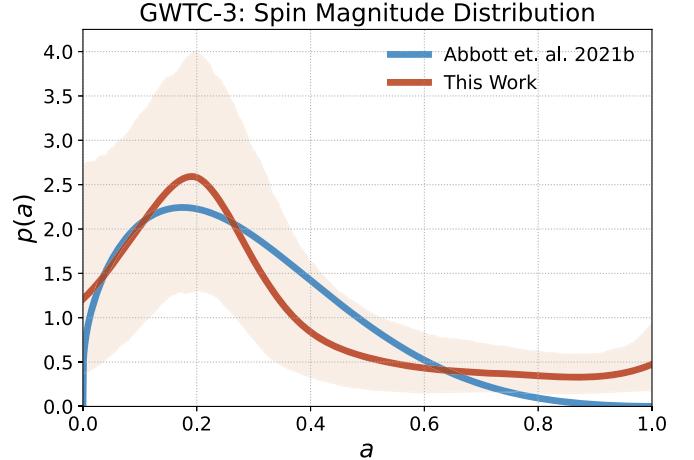


Figure 3. The spin magnitude distribution inferred with the B-spline model (red) with 16 knots spaced linearly from 0 to 1, assuming the components are IID. The solid line shows the population predictive distribution (PPD), and the shaded region shows the 90% credible interval. For comparison, we show the PPD inferred via the DEFAULT (blue) model from the LVK’s GWTC-3 population analyses (The LIGO Scientific Collaboration et al. 2021b). The icons link to the code used to generate this figure and to Zenodo entries of any public data used. 

shallower slope from $q \sim 0.3$ to $q \sim 0.7$, though uncertainty is larger in this region. The sharp decrease in rate just below $q \sim 0.5$ is due to the minimum mass-ratio truncation defined by $q_{\min} = \frac{m_{\min}}{m_1}$. When marginalizing over the primary mass distribution with a strong peak at $10 M_{\odot}$, the mass-ratio distribution truncates at $q \sim 0.5$: the minimum mass, $5 M_{\odot}$, divided by the most common primary mass, $\sim 10 M_{\odot}$.

3.2. Binary Black Hole Spins

3.2.1. Spin Magnitude

The DEFAULT spin model (used by The LIGO Scientific Collaboration et al. 2021b) describes the spin magnitude of both components as identical and independently distributed (IID) nonsingular beta distributions (Talbot & Thrane 2017; Wysocki et al. 2019). The beta distribution provides a simple two-parameter model that can produce a wide range of functional forms on the unit interval. However, the constraint that keeps the beta distribution nonsingular (i.e., $\alpha > 1$ and $\beta > 1$) enforces a spin magnitude that always has $p(a_i = 0) = 0$. Recent studies have proposed the possible existence of a distinct subpopulation of nonspinning or negligibly spinning black holes that can elude discovery with such a model (Fuller & Ma 2019; Galaudage et al. 2021; Roulet et al. 2021; Callister et al. 2022; Tong et al. 2022).

We model the spin magnitude distributions as IID B-spline distributions. Figure 3 shows the spin magnitude distribution inferred with the B-spline model, compared with the DEFAULT model from The LIGO Scientific Collaboration et al. (2021b). The B-spline model results are consistent with those using the beta distribution, peaking near $a \sim 0.2$, with 90% of BBH spins below $0.71^{+0.13}_{-0.14}$ at 90% credibility. The B-spline model does not impose vanishing support at the extremal values like the beta distribution, allowing it to probe the zero-spin question. We find broad support, with large variance, for nonzero probabilities at $a_i = 0$, but cannot confidently determine the presence of a significant nonspinning subpopulation, corroborating similar

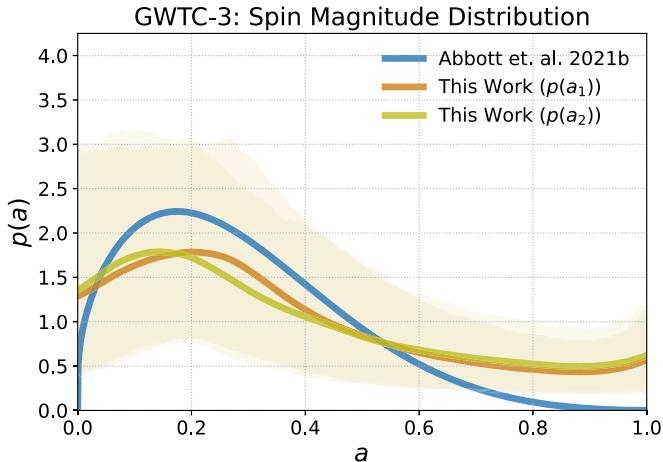


Figure 4. The primary (orange) and secondary (olive) spin magnitude distributions inferred with the B-spline model with 16 knots spaced linearly from 0 to 1. The solid line shows the population predictive distribution (PPD), and the shaded region shows the 90% credible interval. For comparison, we show the PPD inferred via the DEFAULT (blue) model from the LVK’s GWTC-3 population analyses (The LIGO Scientific Collaboration et al. 2021b). The icons link to the code used to generate this figure and to Zenodo entries of any public data used.   

recent conclusions (Galaudage et al. 2021; Callister et al. 2022; Tong et al. 2022; Mould et al. 2022). We repeat the same analysis with independent B-spline distributions for each spin magnitude component. In Figure 4, we show the inferred primary (orange), and secondary (olive) spin magnitude distributions inferred when relaxing the IID assumption along with summary statistics of the infer redspin distributions in Table 1. We find no signs that the spin magnitude distributions differ between the primary and secondary component black holes in BBHs, with both component distributions being consistent with the inferred IID model in Figure 3, with similar support at near-vanishing spins. The secondary spin magnitude distribution is more uncertain due to the higher measurement uncertainty when inferring the secondary spins of BBH systems (Vitale et al. 2014, 2017a). The PPD of the secondary distribution peaks at smaller spin magnitudes ($a \sim 0.15$) than the primary distribution or B-spline IID spin magnitude distribution in Figure 3, although the distributions are broadly consistent with each other considering the large uncertainties. Future signs of component spin magnitude distributions that are uniquely distributed can be produced through mass-ratio reversal in isolated binary evolution (Mould et al. 2022).

3.2.2. Spin Orientation

The DEFAULT spin model (used in Abbott et al. 2021b; The LIGO Scientific Collaboration et al. 2021b) also assumes the spin orientation of both components are identical and independently distributed (IID), with a mixture model over an aligned and an isotropic component. The aligned component is modeled with a truncated Gaussian distribution with mean at $\cos \theta = 1$ and variance as a free hyperparameter to be fit (Talbot & Thrane 2017; Wysocki et al. 2019; Abbott et al. 2021b; The LIGO Scientific Collaboration et al. 2021b). This provides a simple two-parameter model motivated by simple distributions expected from the two main formation scenario families, allowing for a straightforward interpretation of results. One possible limitation, however, is that by construction this distribution is forced to peak at perfectly aligned spins, i.e.,

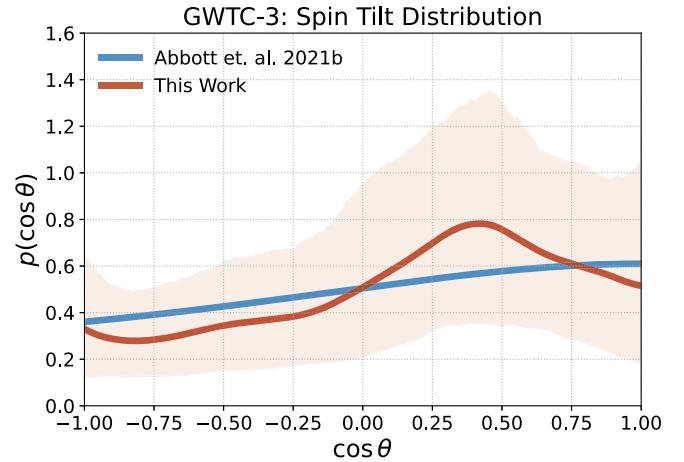


Figure 5. The spin orientation distribution inferred with the B-spline model (red), with 16 knots spaced linearly from -1 to 1 , and assuming the components are IID. The solid line shows the population predictive distribution (PPD), and the shaded region shows the 90% credible interval. For comparison, we show the PPD inferred via the DEFAULT (blue) model from the LVK’s GWTC-3 population analyses (The LIGO Scientific Collaboration et al. 2021b). The icons link to the code used to generate this figure and to Zenodo entries of any public data used.   

$\cos \theta = 1$. While this may be a reasonable assumption, Vitale et al. (2022) recently extended the model space of parametric descriptions used to model the spin orientation distribution and found considerable evidence that the distribution peaks away from $\cos \theta = 1$. Again, this provides a clear use case where data-driven models can help us understand the population.

Figure 5 shows the spin orientation distribution inferred with the IID spin B-spline model, compared with the DEFAULT model from The LIGO Scientific Collaboration et al. (2021b). The B-spline inferences have large uncertainties but start to show the same features as found and discussed in Vitale et al. (2022). As shown in Table 2, we find a distribution that, instead of intrinsically peaking at $\cos \theta = 1$, is found to peak at $\cos \theta = 0.44^{+0.56}_{-0.53}$, at 90% credibility. We find less, but still considerable, support for misaligned spins (i.e., $\cos \theta < 0$), consistent with other recent studies (Abbott et al. 2021b; The LIGO Scientific Collaboration et al. 2021b; Callister et al. 2022). Specifically, we find that the fraction of misaligned systems is $f_{\cos \theta < 0} = 0.35^{+0.11}_{-0.11}$, compared to $f_{\cos \theta < 0} = 0.44^{+0.052}_{-0.11}$ with the DEFAULT model from The LIGO Scientific Collaboration et al. (2021b). This implies the presence of an isotropic component as expected by dynamical formation channels, albeit less than with the DEFAULT model. To quantify the amount of isotropy in the tilt distribution, we calculate $\log_{10} Y$, where Y is the ratio of nearly aligned tilts to nearly antialigned ones, introduced in Vitale et al. (2022) and defined as

$$Y \equiv \frac{\int_{0.9}^{1.0} d \cos \theta p(\cos \theta)}{\int_{-1.0}^{-0.9} d \cos \theta p(\cos \theta)}. \quad (1)$$

The log of this quantity, $\log_{10} Y$, is 0 for a tilt distribution that is purely isotropic, negative when antialigned values are favored, and positive when aligned tilts are favored. We find a $\log_{10} Y = 0.24^{+0.46}_{-0.46}$, exhibiting a slight preference for aligned tilts.

We also model each component’s orientation distribution with an independent B-spline model as done above, and show the inferred primary (orange) and secondary (olive) distributions in Figure 6. The orientation distributions are broadly

Table 1
Summary of Component Spin Distributions

Model	a_{peak}	$a_{90\%}$	$\cos \theta_{\text{peak}}$	$f_{\cos \theta < 0}$	$\log_{10} Y$
B-spline IID	$0.19^{+0.12}_{-0.16}$	$0.71^{+0.13}_{-0.14}$	$0.44^{+0.56}_{-0.53}$	$0.35^{+0.11}_{-0.11}$	$0.24^{+0.46}_{-0.46}$
B-spline Ind(primary)	$0.2^{+0.24}_{-0.2}$	$0.77^{+0.11}_{-0.13}$	$0.16^{+0.84}_{-0.84}$	$0.43^{+0.19}_{-0.16}$	$0.12^{+0.5}_{-0.53}$
B-spline Ind(secondary)	$0.17^{+0.29}_{-0.17}$	$0.8^{+0.1}_{-0.14}$	$0.38^{+0.62}_{-1.0}$	$0.38^{+0.18}_{-0.15}$	$0.18^{+0.53}_{-0.54}$
DEFAULT (The LIGO Scientific Collaboration et al. 2021b)	$0.16^{+0.11}_{-0.13}$	$0.53^{+0.098}_{-0.073}$	$1.0^{+0.0}_{-0.0}$	$0.44^{+0.052}_{-0.12}$	$0.19^{+0.4}_{-0.17}$

Note. Inferred from both the independent and IID component spin B-spline models and the DEFAULT spin model from The LIGO Scientific Collaboration et al. (2021b).

Table 2
Summary of the Effective Spin Distributions

Model	$\chi_{\text{eff,peak}}$	$f_{\chi_{\text{eff}} < 0}$	$f_{\chi_{\text{eff}} < -0.3}$	f_{dyn}	f_{HM}
B-spline IID	$0.039^{+0.034}_{-0.038}$	$0.34^{+0.11}_{-0.11}$	$0.019^{+0.021}_{-0.012}$	$0.69^{+0.22}_{-0.22}$	$0.12^{+0.13}_{-0.074}$
B-spline Ind	$0.023^{+0.034}_{-0.034}$	$0.41^{+0.083}_{-0.088}$	$0.035^{+0.027}_{-0.018}$	$0.82^{+0.17}_{-0.18}$	$0.22^{+0.17}_{-0.11}$
DEFAULT (The LIGO Scientific Collaboration et al. 2021b)	$0.017^{+0.034}_{-0.022}$	$0.43^{+0.059}_{-0.13}$	$0.013^{+0.017}_{-0.0095}$	$0.87^{+0.12}_{-0.26}$	$0.081^{+0.11}_{-0.059}$
GAUSSIAN (The LIGO Scientific Collaboration et al. 2021b)	$0.06^{+0.029}_{-0.037}$	$0.28^{+0.15}_{-0.13}$	$0.00024^{+0.0081}_{-0.00024}$	$0.55^{+0.3}_{-0.26}$	$0.0015^{+0.051}_{-0.0015}$

Note. Inferred with the B-spline model variations, along with the DEFAULT and GAUSSIAN models from The LIGO Scientific Collaboration et al. (2021b).

consistent with each other and the DEFAULT model's PPD, given the wide credible intervals. We find the two distributions to peak at $\cos \theta_1 = 0.16^{+0.84}_{-0.84}$ and $\cos \theta_2 = 0.38^{+0.62}_{-1.0}$, showing that the primary distribution peak is inferred further away from the assumed $\cos \theta = 1$ with the DEFAULT model. There is also significant (albeit uncertain) evidence of spin misalignment in each distribution, with the fraction of misaligned primary and secondary components being found as $f_{\cos \theta_1 < 0} = 0.43^{+0.19}_{-0.16}$ and $f_{\cos \theta_2 < 0} = 0.38^{+0.18}_{-0.15}$. We again calculate $\log_{10} Y$ for each component distribution and find $\log_{10} Y_1 = 0.12^{+0.5}_{-0.53}$ and $\log_{10} Y_2 = 0.18^{+0.53}_{-0.54}$.

3.3. The Effective Spin Dimension

While the component spin magnitudes and tilts are more directly tied to formation physics, they are typically poorly measured. The best-measured spin quantity, which enters at the highest post-Newtonian order, is the effective spin: $\chi_{\text{eff}} = \frac{a_1 \cos \theta_1 + q a_2 \cos \theta_2}{1+q}$. There is additionally an effective precessing spin parameter, $\chi_p = \max \left[a_1 \sin \theta_1, \frac{3+4q}{4+3q} q a_2 \sin \theta_2 \right]$, that quantifies the amount of spin precession given the system's mass ratio and component spin magnitudes and orientation. Figure 7 shows the inferred effective spin and precessing spin distributions with the two versions of our B-spline models (red and purple), along with results on the DEFAULT (Talbot & Thrane 2017) and GAUSSIAN (Miller et al. 2020) models from The LIGO Scientific Collaboration et al. (2021b). We find considerable agreement among the effective spin distributions, but the more flexible B-spline models in component spins more closely resemble results from the DEFAULT model, also using the component spins. Table 2 shows summary statistics of the effective spin dimension distributions. The B-spline model finds shapes very similar to those of the other models, with a single peak centered at $\chi_{\text{eff}} = 0.039^{+0.034}_{-0.038}$, compared to $\chi_{\text{eff}} = 0.017^{+0.034}_{-0.022}$ with the DEFAULT

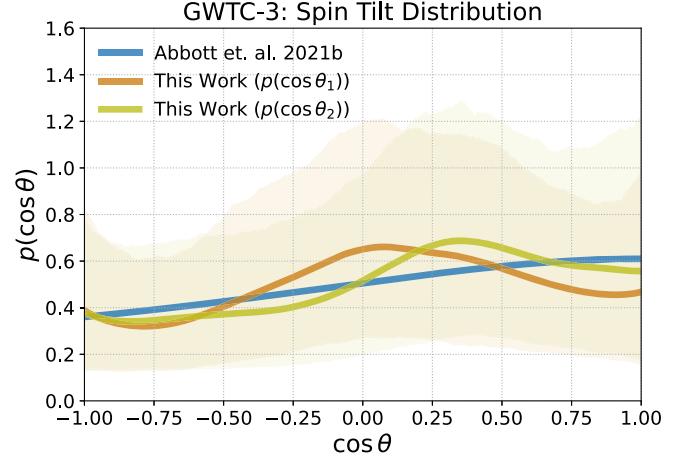


Figure 6. The primary (orange) and secondary (olive) spin orientation distributions inferred with the B-spline model with 16 knots spaced linearly from -1 to 1. The solid line shows the population predictive distribution (PPD), and the shaded region shows the 90% credible interval. For comparison, we show the PPD inferred via the DEFAULT (blue) model from the LVK's GWTC-3 population analyses (The LIGO Scientific Collaboration et al. 2021b). The icons link to the code used to generate this figure and to Zenodo entries of any public data used.   

model and $\chi_{\text{eff}} = 0.06^{+0.029}_{-0.037}$ with the GAUSSIAN χ_{eff} models from The LIGO Scientific Collaboration et al. (2021b). As for spin misalignment, we calculate the fraction of systems with effective spins that are misaligned (i.e., $\chi_{\text{eff}} < 0$) and find similar agreement with previous work (Abbott et al. 2021b; The LIGO Scientific Collaboration et al. 2021b; Callister et al. 2022). We find for the B-spline model $f_{\chi_{\text{eff}} < 0} = 0.34^{+0.11}_{-0.11}$, compared to $f_{\chi_{\text{eff}} < 0} = 0.43^{+0.059}_{-0.13}$ and $f_{\chi_{\text{eff}} < 0} = 0.28^{+0.15}_{-0.13}$ with the DEFAULT and GAUSSIAN models from The LIGO Scientific Collaboration et al. (2021b). The precessing spin distributions inferred with the B-spline models exhibit a shape similar to that of the DEFAULT

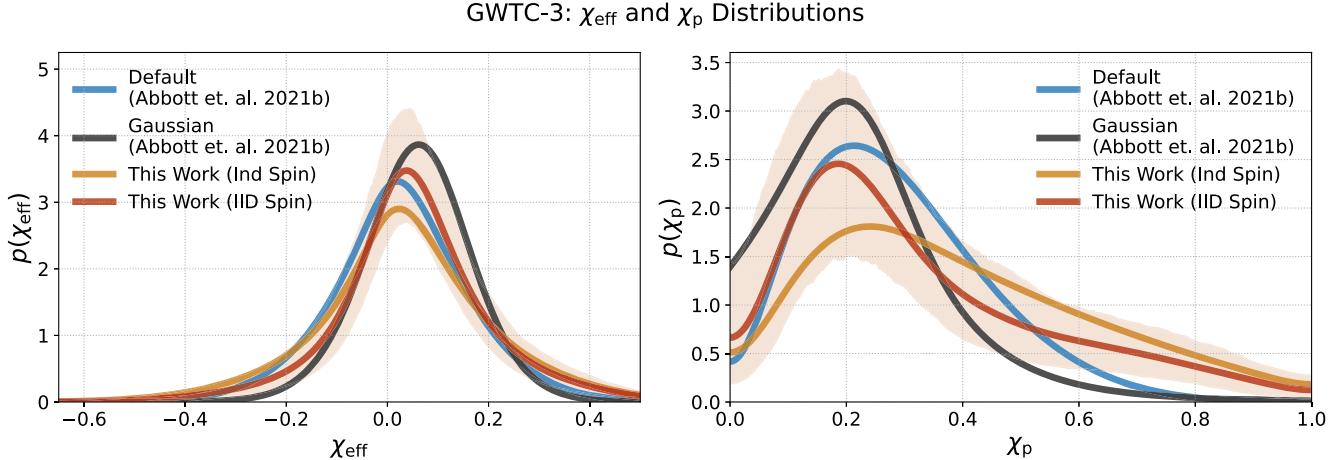


Figure 7. The effective (left) and precessing (right) spin distributions inferred with the B-spline IID spin model (red). The solid line shows the population predictive distribution (PPD), and the shaded region shows the 90% credible interval. We show the inferred PPDs from the independent component spin B-spline model (purple), and both the DEFAULT (blue) model and the Gaussian (green) model from the LVK’s GWTC-3 population analyses (The LIGO Scientific Collaboration et al. 2021b). The icons link to the code used to generate this figure and to Zenodo entries of any public data used.   

model, but with a much fatter tail toward highly precessing systems, driven by the extra support for highly spinning components seen in Figures 3 and 4.

3.4. Merger Rate Evolution with Redshift

Recent analysis of the GWTC-3 BBH population has shown evidence for an increasing merger rate with redshift, nearly ruling out a merger rate that is constant with comoving volume (Fishbach et al. 2018; The LIGO Scientific Collaboration et al. 2021b). When extending the power-law form of the previously used model to have a modulation that we model with B-splines, the merger rate as a function of redshift in Figure 8 shows mild support for features departing from the underlying power law. In particular, we see a small increase in merger rate from $z \sim 0.09$ to $z \sim 0.2$ (where we best constrain the rate), followed by a plateau in the rate from $z \sim 0.2$ to $z \sim 0.4$. At larger redshifts, where we begin to have sparse observations, we see no sign of departure from the power law as the rate continues to increase with redshift. The underlying power-law slope of our B-spline modulated model is consistent with the GWTC-3 results with the underlying model by itself: the POWERLAWREDSHIFT model found $\lambda_z = 2.7^{+1.8}_{-1.9}$ when inferred with the POWERLAWPEAK mass and DEFAULT spin models. Our more flexible model infers a power-law slope of $\lambda_z = 2.1^{+2.2}_{-2.5}$. We show the basis spline modulations or departure from the power law in Figure 9, compared to the prior—showing where we cannot constrain any significant deviations from the simpler parametric power-law model. The extra freedom of our model does inflate the uncertainty in its rate estimates, especially at $z \sim 0$, where there are not any observations in the catalog. We find a local ($z = 0$) merger rate of $\mathcal{R}_0 = 20^{+29}_{-12} \text{ Gpc}^{-3} \text{ yr}^{-1}$ using the B-spline modulation model, as compared to $\mathcal{R}_0 = 17^{+10}_{-6.7} \text{ Gpc}^{-3} \text{ yr}^{-1}$ for the GWTC-3 result.

4. Astrophysical Implications

The collective distribution of BBH source properties provides a useful probe of the complex and uncertain astrophysics that govern their formation and evolution until merger (Rodriguez et al. 2016; Farr et al. 2017; Zevin et al. 2017). Our analysis with the newly constructed B-spline

models uncovers hints of new features in the population (e.g., in mass ratio and redshift), corroborates important conclusions of recent work, and provides a robust data-driven framework for future population studies.

The results presented in Section 3.1 illustrate a mass distribution wider than that inferred with power-law-based models in The LIGO Scientific Collaboration et al. (2021b), and a suppressed merger rate at low primary masses (i.e., $\lesssim 8M_{\odot}$), showing possible signs of binary selection effects or the purported low-mass gap between neutron stars and black holes (Fishbach et al. 2020; van Son et al. 2022a; Farah et al. 2022). While isolated formation is able to predict the $10M_{\odot}$ peak (Antonini & Gieles 2020), cluster and dynamical formation scenarios struggle to predict a peak in the BH mass distribution less than $15-20M_{\odot}$ (Hong et al. 2018; Rodriguez et al. 2019). Globular cluster formation is expected to produce more top-heavy mass distributions than isolated, and recent studies have shown suppressed BBH merger rates at lower ($m \lesssim 15M_{\odot}$) masses when compared to predictions from the isolated channel (Rodriguez et al. 2015; Belczynski et al. 2016; Rodriguez et al. 2019; Bavera et al. 2021). BBHs that form near active galactic nuclei (AGN) can preferentially produce higher-mass black holes (Yang et al. 2019; Tagawa et al. 2021; Ford & McKernan 2022). We do not find any evidence for a truncation or rapid decline in the merger rate as a function of mass, which stellar evolution theory predicts due to pair-instability supernovae (PISNe; Heger & Woosley 2002; Woosley et al. 2002; Heger et al. 2003; Spera & Mapelli 2017; Stevenson et al. 2019). The original motivation for the peak in the POWERLAWPEAK model (Talbot & Thrane 2018) was to represent a possible “pileup” of masses just before such truncation, because massive stars just light enough to avoid PISN will shed large amounts of mass in a series of “pulses” before collapsing to BHs, in a process called pulsational pair-instability supernova (PPISN) (Woosley 2017; Farmer et al. 2019; Woosley 2019). While the predictions of the mass scale where pair-instability kicks in are uncertain and depend on poorly understood physics like nuclear reaction rates of carbon and oxygen in the core of stars, models have a hard time producing this peak lower than $m \sim 40M_{\odot}$ (Belczynski et al. 2016; Farmer et al. 2019; Marchant et al. 2019; Farmer et al. 2020; Renzo et al. 2020). The lack of a

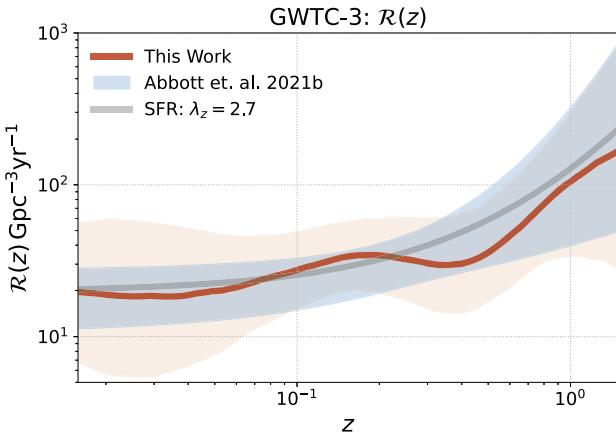


Figure 8. The BBH merger rate as a function of redshift. We show the B-spline model (red) with 16 knots spaced linearly in $\log(z)$, from the minimum to the maximum observed redshifts. The solid line shows the population predictive distribution (PPD), and the shaded region shows the 90% credible interval. We show the inferred 90% credible interval from the POWERLAWREDSHIFT model from the LVK’s GWTC-3 population analyses in blue and a power law with exponent of 2.7 in gray, representing the expected star formation rate (Madau & Dickinson 2014; The LIGO Scientific Collaboration et al. 2021b). The icons link to the code used to generate this figure and to [Zenodo](#) entries of any public data used.

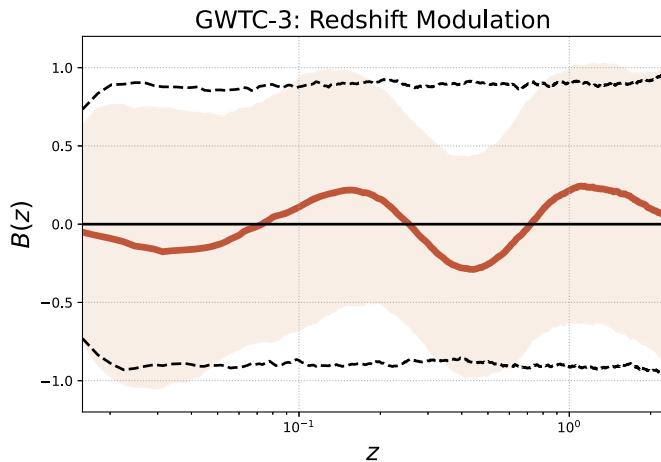


Figure 9. The B-spline modulation to the underlying power law in redshift, (red). The solid line shows the population predictive distribution (PPD), and the shaded region the 90% credible interval. We show the 90% credible interval of the prior predictive distribution in dashed black lines. The icons link to the code used to generate this figure and to [Zenodo](#) entries of any public data used.

truncation could point toward a higher prevalence of dynamical processes that can produce black holes in mass ranges stellar collapse cannot, such as hierarchical mergers of BHs (Fishbach & Holz 2017; Doctor et al. 2020; Kimball et al. 2020; Doctor et al. 2021; Kimball et al. 2021; Fishbach et al. 2022), very low-metallicity Population III stars (Belczynski 2020; Farrell et al. 2021), new beyond-standard-model physics (Croon et al. 2020; Sakstein et al. 2020), or black hole accretion of BHs in gaseous environments such as AGNs (McKernan et al. 2020; Secunda et al. 2020; Cruz-Osorio et al. 2021).

Our constraints on the mass-ratio distribution are not yet precise enough to claim definitive departures from power-law behavior, but they do suggest possible plateaus in the rate at several mass ratios, including equal mass. These features should sharpen (or resolve) with future updates to the catalog.

Section 3.2 focused on inferences of the spin distributions of black holes, observing evidence of spin misalignment, spin antialignment, and suppressed support for exactly aligned systems. These point toward a significant contribution to the population from dynamical formation processes, agreeing with conclusions drawn about the mass distribution inference of Section 3.1. While field formation is expected to produce systems with preferentially aligned spins due to tidal interactions, observational evidence suggests that tides may not be able to realign spins in all systems as some isolated population models assume. Additionally, because of uncertain knowledge of supernovae kicks, isolated formation can produce systems with negative but small effective spins. Consistent with recent studies, we report an effective spin distribution that is not symmetric about zero, disfavoring a scenario in which all BBHs are formed dynamically (Abbott et al. 2021b; The LIGO Scientific Collaboration et al. 2021b; Callister et al. 2022). Following the rules in Fishbach et al. (2022), we place conservative upper bounds on the fraction of hierarchical mergers, f_{HM} , and the fraction of dynamically formed BBHs, f_{dyn} , with the B-spline χ_{eff} model constraining $f_{\text{HM}} < 0.058$ and $f_{\text{dyn}} < 0.52$ at 90% credibility. This is consistent with the 90% credible interval found from the GWTC-2 analysis, $0.25 \leq f_{\text{dyn}} \leq 0.93$ (Abbott et al. 2021b).

Finally, Section 3.4 shows the potentially interesting evolution of the BBH merger rate with redshift. Though uncertainties are still large, we may be seeing the first signs of departure from following the star formation rate, which could help in distinguishing different subpopulations should they exist (van Son et al. 2022b). Again, we expect these features to be resolved with future catalogs.

5. Conclusions

Nonparametric and data-driven statistical modeling methods have been put to use with great success across the ever-growing field of gravitational waves (Farr et al. 2015; Littenberg & Cornish 2015; Doctor et al. 2017; Mandel et al. 2017; Edwards et al. 2018; Edelman et al. 2021; Tiwari 2021; Tiwari & Fairhurst 2021; Vitale et al. 2021; Edelman et al. 2022; Payne & Thrane 2022; Tiwari 2022). We presented a case study exploring how basis splines make for an especially powerful and efficient data-driven method of characterizing the binary black hole population observed with gravitational waves, along with the associated open-source software **GWInferno**, which implements the models described in this paper and performs hierarchical Bayesian inference with NUMPYRO and JAX³ (Phan et al. 2019; Bingham et al. 2019). Our study paves the way as the first completely nonparametric compact-object population study, employing data-driven models for each of the hierarchically modeled population distributions. A complete understanding of the population properties of compact objects will help to advance poorly understood areas of stellar and nuclear astrophysics and provide a novel independent cosmological probe. With the coming influx of new data with the LVK’s next observing run, development of model-agnostic methods, such as the one we proposed here, will become necessary to efficiently make sense of the vast amounts of data and to extract as much information as possible from the population.

We thank Tom Callister, Will Farr, Maya Fishbach, Salvatore Vitale, and Jaxen Godfrey for useful discussions during the

³ <https://github.com/google/jax>

preparation of this manuscript and/or helpful comments on early drafts. Z.D. also acknowledges support from the CIERA Board of Visitors Research Professorship. This research has made use of data, software, and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org/>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation grants PHY-0757058 and PHY-0823459. This work was supported in part by the National Science Foundation under grant PHY-2146528 and benefited from access to the University of Oregon high-performance computer, Talapas. This material is based upon work supported in part by the National Science Foundation under grant PHY-1807046 and work supported by NSF's LIGO Laboratory, which is a major facility fully funded by the National Science Foundation.

Software: SHOWYOURWORK (Luger et al. 2021), MATPLOTLIB (Hunter 2007), NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), ASTROPY (Astropy Collaboration et al. 2018), JAX (<https://github.com/google/jax>), NUMPYRO (Phan et al. 2019; Bingham et al. 2019).

Appendix A Basis Splines

A common nonparametric method used in many statistical applications is that of basis splines. A spline function of order k is a piece-wise polynomial of order k polynomials stitched together from defined “knot” locations across the domain. They provide a useful and cheap way to interpolate generically smooth functions from a finite sampling of “knot” heights. Basis splines of order k are a set of order k polynomials that form a complete basis for any spline function of order k . Therefore, given an array of knot locations, \mathbf{t} or knot vector, there exists a single unique linear combination of basis splines for every possible spline function interpolated from \mathbf{t} . To construct a basis of n components and knots, t_0, t_1, \dots, t_{i+k} , we use the Cox–de Boor recursion formula (de Boor 1978; Ramsay 1988). The recursion starts with the $k=0$ (constant) case and recursively constructs the basis components of higher orders. The base case and recursion relation that generates this particular basis are defined as

$$B_{i,0}(x|\mathbf{t}) = \begin{cases} 1, & \text{if } t_i \leq x < t_{i+1} \\ 0, & \text{otherwise} \end{cases} \quad (\text{A1})$$

$$B_{i,k+1}(x|\mathbf{t}) = \omega_{i,k}(x|\mathbf{t})B_{i,k}(x|\mathbf{t}) + [1 - \omega_{i+1,k}(x|\mathbf{t})]B_{i+1,k}(x|\mathbf{t}) \quad (\text{A2})$$

$$\omega_{i,k}(x|\mathbf{t}) = \begin{cases} \frac{x - t_i}{t_{i+k} - t_i}, & t_{i+k} \neq t_i \\ 0, & \text{otherwise} \end{cases} \quad (\text{A3})$$

This is known as the “B-spline” basis, after its inventor de Boor (de Boor 1978). The power of basis splines comes from the fact that one only has to do the somewhat-expensive interpolation once for each set of points at which the spline is evaluated. This provides a considerable computational speedup as each evaluation of the spline function becomes a simpler operation: a dot product of a matrix and a vector. This straightforward operation is also ideal for optimizations from the use of GPU accelerators, enabling our Markov Chain Monte Carlo (MCMC) based

analyses, often with hundreds of parameters, to converge in an hour or less. Basis splines can easily be generalized to their two-dimensional analog, producing tensor product basis splines that, with this computational advantage, allow for high-fidelity modeling of two-dimensional spline functions.

Another important feature of basis splines is that, under appropriate prior conditions, one can alleviate sensitivities to arbitrarily chosen prior specifications that splines commonly struggle with. Previous studies using splines had to perform multiple analyses, varying the number of spline knots, then either marginalized over the models or used model comparisons to motivate the best choice (Edelman et al. 2022). We can avoid this step with the use of penalized splines (or P-splines) (Lang & Brezger 2004; Jullion & Lambert 2007; Eilers & Marx 2021), where one adds a smoothing prior comprised of Gaussian distributions on the differences between neighboring basis spline coefficients. This allows for knots to be densely populated across the domain without the worry of extra variance in the inferred spline functions. When also fitting the scale of the smoothing prior (i.e., the width of the Gaussian distributions on the differences), the data will inform the model of the preferred scale of smoothing required. We discuss the details of our smoothing prior implementation in more detail in the next section, Appendix B, and our specific prior and the basis choices for each model in Appendix D.

Appendix B Penalized Splines and Smoothing Priors

Spline functions have been shown to be sensitive to the chosen number of knots, as well as their locations or spacing (de Boor 1978). Adding more knots increases the a priori variance in the spline function, while the space between knots can limit the resolution of features in the data the spline is capable of resolving. To ensure a spline-based model is flexible enough, one would want to add as many knots as densely as possible, but this comes with the unwanted side effect of larger variance imposed by the model. This can be fixed with the use of penalized splines (P-splines) in which one applies a prior or regularization term to the likelihood based on the difference of adjacent knot coefficients (Eilers & Marx 2021). The linear combination of spline basis components or the resulting spline function is flat when the basis coefficients are equal (see Figure 10). By penalizing the likelihood as the differences between adjacent knot coefficients get larger, one gets a smoothing effect on the spline function (Eilers & Marx 2021). With hierarchical Bayesian inference as our statistical framework, we formulate the penalized likelihood of the P-splines of Eilers & Marx (2021) with their Bayesian analog (Lang & Brezger 2004). The Bayesian P-spline prior places Gaussian distributions over the r th-order differences of the coefficients (Lang & Brezger 2004; Jullion & Lambert 2007). This is also sometimes referred to as a Gaussian random walk prior, and is similar in spirit to a Gaussian process prior used to regularize or smooth histogram bin heights as done in other nonparametric population studies (Mandel et al. 2017; The LIGO Scientific Collaboration et al. 2021b). For a spline basis with n degrees of freedom, and a difference penalty of order of r (see Eilers & Marx 2021), the smoothing prior on our basis spline coefficients, \mathbf{c} , is defined as

$$\mathbf{c} \sim \mathcal{N}(0, \sigma) \quad (\text{B1})$$

$$p(\mathbf{c}|\tau_\lambda) \propto \exp \left[-\frac{1}{2} \tau_\lambda \mathbf{c}^T \mathbf{D}_r^T \mathbf{D}_r \mathbf{c} \right]. \quad (\text{B2})$$

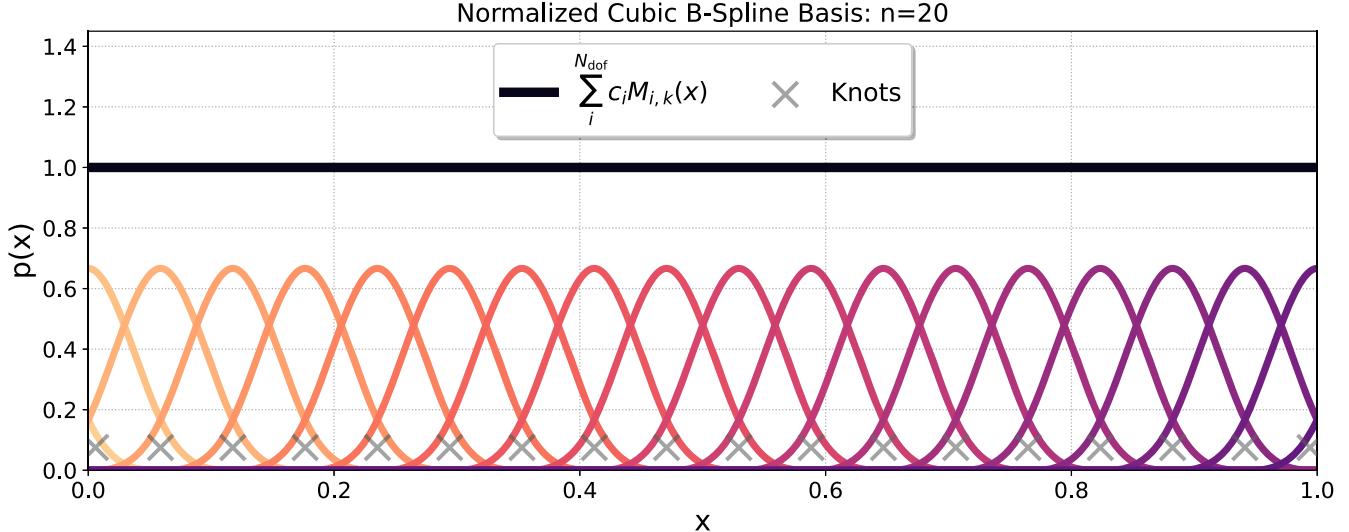


Figure 10. Plot showing a “proper” (see Appendix B) normalized B-spline basis of order 3 (cubic) with 20° of freedom and equal weights for each component. In black, we show the resulting spline function given equal weights and denote the location of the knots with gray x symbols. The icon links to the code used to generate this figure. 

Above, \mathbf{D}_r is the order- r difference matrix, of shape $(n - r \times n)$, and $\mathcal{N}(0, \sigma)$ a Gaussian distribution with zero mean and standard deviation, σ . This smoothing prior removes the strong dependence on number and location of knots that arises with using splines. The τ_λ controls the “strength” of the smoothing, or the inverse variance of the Gaussian priors on knot differences. We place uniform priors on τ_λ marginalized over this smoothing scale hyperparameter to let the data inform the optimal scale needed. When there are a very large number of knots, such that the domain is densely populated with basis coefficients, this allows the freedom for the model to find the smoothing scale that the data prefer.

This prior is imparting a natural attraction of the coefficients closer to each other in order to smooth the spline function, so one must ensure that the spline function is in fact flat given all equal coefficients. There need to be $n + k + 1$ knots to construct an order- k basis with n degrees of freedom. Some studies place knots on top of each other at hard parameter boundaries (de Boor 1978; Ramsay 1988), which may seem motivated, but this violates the above condition necessary for the P-spline prior. We follow the distinction in Eilers & Marx (2021) that such a smoothing prior is only valid with “proper” spline bases. A proper basis is where all $n + k + 1$ knots are evenly and equally spaced (see Figure 10), as opposed to stacking them at the bounds.

Appendix C Hierarchical Bayesian Inference

We use hierarchical Bayesian inference to infer the population properties of compact binaries. We want to infer the number density of merging compact binaries in the universe and how this can change with their masses, spins, etc. Oftentimes, it is useful to formulate the question in terms of the merger rate, which is the number of mergers per Gpc^3 comoving volume per year. For a set of hyperparameters, Λ , λ , and overall merger rate, \mathcal{R} , we write the overall number density

of BBH mergers in the universe as

$$\frac{dN(\theta, z | \mathcal{R}, \Lambda, \lambda)}{d\theta dz} = \frac{dV_c}{dz} \left(\frac{T_{\text{obs}}}{1+z} \right) \frac{d\mathcal{R}(\theta, z | \mathcal{R}_0, \Lambda, \lambda)}{d\theta} = \mathcal{R}p(\theta | \Lambda)p(z | \lambda), \quad (\text{C1})$$

where up above, we denote the comoving volume element as dV_c (Hogg 1999), and T_{obs} as the observing time period that produced the catalog, with the related factor of $1+z$ converting this detector-frame time to source-frame. We assume a ΛCDM cosmology using the cosmological parameters from Planck Collaboration et al. (2016). We model the merger rate evolving with redshift following a power-law distribution: $p(z | \lambda) \propto \frac{dV_c}{dz} \frac{1}{1+z} (1+z)^\lambda$ (Fishbach et al. 2018). When integrating Equation (C1) across all θ and out to some maximum redshift, z_{max} , we get the total number of compact binaries in the Universe out to that redshift. We follow previous notations, letting $\{d_i\}$ represent the set of data from N_{obs} compact binaries observed with gravitational waves. The merger rate is then described as an inhomogeneous Poisson process, and after imposing the usual log-uniform prior on the merger rate, we marginalize over the merger rate, \mathcal{R} , and arrive at the posterior distribution of our hyperparameters, Λ (Mandel et al. 2019; Vitale et al. 2021):

$$p(\Lambda, \lambda | \{d_i\}) \frac{p(\Lambda)p(\lambda)}{\xi(\Lambda, \lambda)^{N_{\text{obs}}}} \prod_{i=1}^{N_{\text{obs}}} \left[\frac{1}{K_i} \sum_{j=1}^{K_i} \frac{p(\theta^{i,j} | \Lambda)p(z^{i,j} | \lambda)}{\pi(\theta, z^{i,j})} \right], \quad (\text{C2})$$

where above, we replaced the integrals over each event’s likelihood with ensemble averages over K_i posterior samples (LVK Collaboration 2021a). Here, j indexes the K_i posterior samples from each event and $\pi(\theta, z)$ is the default prior used by parameter estimations that produced the posterior samples for each event. In the analyses of GWTC-3, either the default prior used was uniform in detector-frame masses, component spins,

and Euclidean volume, or the posterior samples were reweighted to such a prior before using them in our analysis. The corresponding prior evaluated in the parameters we hierarchically model, i.e., source-frame primary mass, mass ratio, component spins, and redshift:

$$\begin{aligned} \pi(m_1, q, a_1, a_2, \cos \theta_1, \cos \theta_2, z) \\ \propto \frac{1}{4} m_1 (1+z)^2 D_L^2(z) \frac{dD_L}{dz}. \end{aligned} \quad (\text{C3})$$

Above, D_L is the luminosity distance. To carefully incorporate selection effects to our model, we need to quantify the detection efficiency, $\xi(\Lambda, \lambda)$, of the search pipelines that were used to create GWTC-3, at a given population distribution described by Λ and λ :

$$\xi(\Lambda, \lambda) = \int d\theta dz P_{\text{det}}(\theta, z) p(\theta|\Lambda) p(z|\lambda). \quad (\text{C4})$$

To estimate this integral, we use a software injection campaign where gravitational waveforms from a large population of simulated sources, which are then put into real detector data. These data are then evaluated with the same search pipelines that were used to produce the catalog we are analyzing, giving us detection probabilities for these simulated waveforms. With these search results in hand, we use importance sampling and evaluate the integral with the Monte Carlo sum estimate, μ , and its corresponding variance and effective number of samples:

$$\xi(\Lambda, \lambda) \approx \mu(\Lambda, \lambda) \frac{1}{N_{\text{inj}}} \sum_{i=1}^{N_{\text{found}}} \frac{p(\theta^i|\Lambda) p(z^i|\lambda)}{p_{\text{inj}}(\theta, z^i)} \quad (\text{C5})$$

$$\begin{aligned} \sigma^2(\Lambda, \lambda) \equiv \frac{\mu^2(\Lambda, \lambda)}{N_{\text{eff}}} \simeq \frac{1}{N_{\text{inj}}^2} \sum_{i=1}^{N_{\text{found}}} \\ \times \left[\frac{p(\theta|\Lambda) p(z|\lambda)}{p_{\text{inj}}(\theta, z)} \right]^2 - \frac{\mu^2(\Lambda, \lambda)}{N_{\text{inj}}}, \end{aligned} \quad (\text{C6})$$

where the sum is only over the N_{found} injections that were successfully detected out of N_{inj} total injections, and $p_{\text{inj}}(\theta, z)$ is the reference distribution from which the injections were drawn. We use the LVK released injection sets that describe the detector sensitivities over the first, second, and third observing runs (LVK Collaboration 2021c). Additionally, we follow the procedure outlined in Farr (2019) to marginalize the uncertainty in our estimate of $\xi(\Lambda, \lambda)$, in which we verify that N_{eff} is sufficiently high after reweighting the injections to a given population (i.e., $N_{\text{eff}} > 4N_{\text{obs}}$). The total hyperposterior is

marginalized over the merger rate and the uncertainty in the Monte Carlo integral calculating $\xi(\Lambda, \lambda)$ (Farr 2019):

$$\begin{aligned} \log p(\Lambda, \lambda | \{d_i\}) \propto \sum_{i=1}^{N_{\text{obs}}} \log \left[\frac{1}{K_i} \sum_{j=1}^{K_i} \frac{p(\theta^{i,j}|\Lambda) p(z^{i,j}|\lambda)}{\pi(\theta^{i,j}, z^{i,j})} \right] \\ - N_{\text{obs}} \log \mu(\Lambda, \lambda) + \frac{3N_{\text{obs}} + N_{\text{obs}}^2}{2N_{\text{eff}}} \\ + \mathcal{O}(N_{\text{eff}}^{-2}). \end{aligned} \quad (\text{C7})$$

We explicitly enumerate each of the models used in this work for $p(\theta|\Lambda)$, along with their respective hyperparameters and prior distributions, in the next section. To calculate draw samples of the hyperparameters from the hierarchical posterior distribution shown in Equation (C7), we use the NUTS Hamiltonian Monte Carlo sampler in NUMPYRO and JAX to calculate likelihoods (Phan et al. 2019; Bingham et al. 2019).

Appendix D Model and Prior Specification

For each of the distributions with basis spline distributions, we have two fixed hyperparameters to specify. The number of degrees of freedom, n , and the difference penalty order for the smoothing prior, r . Additionally, one must choose a prior distribution on the smoothing prior scale hyperparameter, τ_λ , which we take to be uniform. For the primary mass distribution, we model the log probability with a B-spline interpolated in $\log(m_1)$ spaced evenly from a minimum black hole mass of $5M_\odot$, and maximum of $100M_\odot$. We follow a similar scheme for the models in mass ratio and spin, except we model the log probability with B-splines that are interpolated in q , a_i , or $\cos \theta_i$ space. The knots for the mass-ratio B-spline are linearly spaced from $\frac{m_{\text{min}}}{m_{\text{max}}} = 0.05$ to 1. There is motivation for the evolution of the merger rate with redshift to follow a power-law form because it should be related to the star formation rate (Madau & Dickinson 2014), motivating our adoption of a semiparametric approach where we use B-splines to model modulations to the simpler underlying POWERLAWREDSHIFT model (Fishbach et al. 2018; Edelman et al. 2022). We model modulations to the underlying probability density with the multiplicative factor, $e^{B(\log z)}$, where $B(\log z)$ is the B-spline interpolated from knots spaced linearly in $\log z$ space. We enumerate each of our specific model hyperparameter and prior choices in Table 3.

Table 3
All Hyperparameter Prior Choices for Each of the Newly Introduced Basis Spline Models from This Article

Model	Parameter	Description	Prior
Primary Mass Model Parameters			
B-SPLINE PRIMARY	c	Basis coefficients	$\sim \text{Smooth}(\tau_\lambda, \sigma, r, n)$
	τ_λ	Smoothing prior scale	$\sim \text{U}(2, 1000)$
	r	Order of the difference matrix for the smoothing prior	2
	σ	Width of Gaussian priors on coefficients in smoothing prior	6
	n	Number of knots in the basis spline	64
Mass Ratio Model Parameters			
B-SPLINE RATIO	c	Basis coefficients	$\sim \text{Smooth}(\tau_\lambda, \sigma, r, n)$
	τ_λ	Smoothing prior scale	$\sim \text{U}(1, 100)$
	r	Order of the difference matrix for the smoothing prior	2
	σ	Width of Gaussian priors on coefficients in smoothing prior	4
	n	Number of knots in the basis spline	18
Redshift Evolution Model Parameters			
POWERLAW+B-SPLINE	λ	Slope of redshift evolution power law $(1 + z)^\lambda$	$\sim \mathcal{N}(0, 3)$
	c	Basis coefficients	$\sim \text{Smooth}(\tau_\lambda, \sigma, r, n)$
	τ_λ	Smoothing prior scale	$\sim \text{U}(1, 10)$
	r	Order of the difference matrix for the smoothing prior	2
	σ	Width of Gaussian priors on coefficients in smoothing prior	1
	n	Number of knots in the basis spline	18
Spin Distribution Model Parameters			
B-SPLINE MAGNITUDE	c	Basis coefficients	$\sim \text{Smooth}(\tau_\lambda, \sigma, r, n)$
	τ_λ	Smoothing prior scale	$\sim \text{U}(1, 10)$
	r	Order of the difference matrix for the smoothing prior	2
	σ	Width of Gaussian priors on coefficients in smoothing prior	1
	n	Number of knots in the basis spline	18
B-SPLINE TILT	c	Basis coefficients	$\sim \text{Smooth}(\tau_\lambda, \sigma, r, n)$
	τ_λ	Smoothing prior scale	$\sim \text{U}(1, 10)$
	r	Order of the difference matrix for the smoothing prior	2
	σ	Width of Gaussian priors on coefficients in smoothing prior	1
	n	Number of knots in the basis spline	18

Note. See Appendices A and B for more detailed descriptions of basis splines and smoothing prior parameters.

Appendix E Posterior Predictive Checks

We follow the posterior predictive checking procedure done in recent population studies to validate our models inferences (Abbott et al. 2021b; Edelman et al. 2022). For each posterior sample describing our model’s inferred population, we reweigh

the observed event samples and the found injections to that population and draw a set of 69 (the size of the GWTC-3 BBH catalog) samples to construct the observed and predicted distributions we show in Figures 11 and 12. When the observed region stays encompassed within the predicted region, the model is performing well, which we see across each of the fit parameters.

GWTC-3: B-Spline Model Posterior Predictive Checks

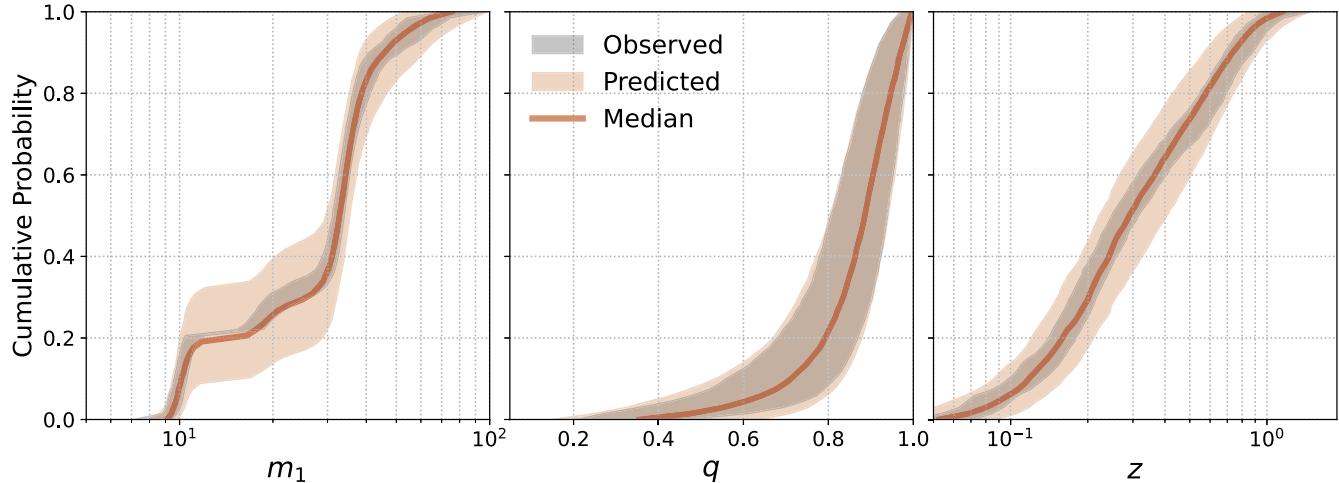


Figure 11. Posterior predictive checks showing the CDFs of the observed (black) and predicted (red) distributions of GWTC-3 sized catalogs for each posterior sample of the IID spin B-spline model. The shaded regions show 90% credible intervals, and the solid red line is the median of the predicted distribution. The icons link to the code used to generate this figure and to Zenodo entries of any public data used. [⬇️](#) [🔗](#)

GWTC-3: B-Spline Spin Model Posterior Predictive Checks

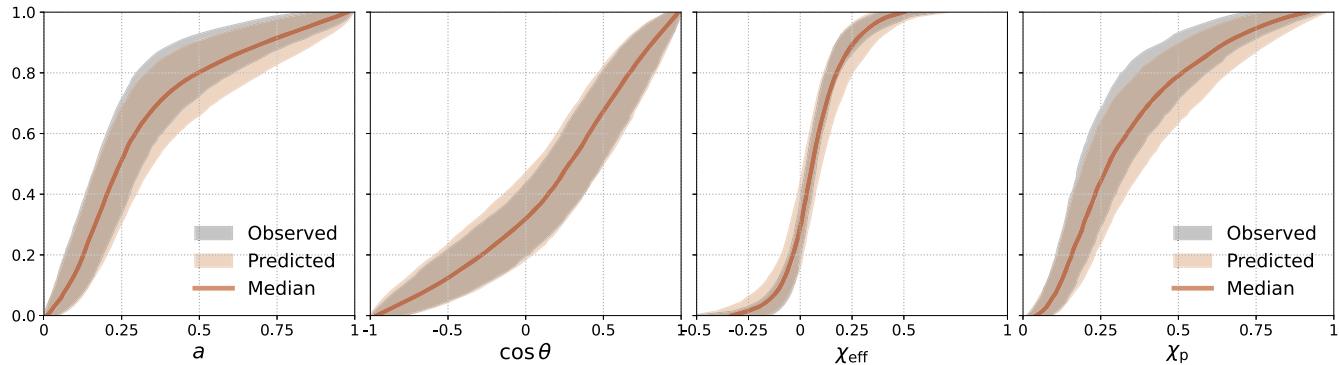


Figure 12. Posterior predictive checks showing the CDFs of the observed (black) and predicted (red) distributions of GWTC-3 sized catalogs for each posterior sample of the IID spin B-spline model. The shaded regions show 90% credible intervals, and the solid red line is the median of the predicted distribution. The icons link to the code used to generate this figure and to Zenodo entries of any public data used. [⬇️](#) [🔗](#)

Appendix F Reproducibility

In the spirit of open-source and reproducible science, this study was done using the reproducibility software `showyourwork` (Luger et al. 2021), which leverages continuous integration to programmatically download the data from <https://zenodo.org>, create the figures, and compile the manuscript. Each figure caption contains two links that point toward the data set (stored on Zenodo) used in the corresponding figure, and to the script used to make the figure (at the commit corresponding to the current build of the article). The git repository associated to this study is publicly available at <https://github.com/bruce-edelman/CoveringYourBasis>, which allows anyone to rebuild the entire article. The data sets and all analysis or figure-generating scripts are all stored at doi:[10.5281/zenodo.7566301](https://doi.org/10.5281/zenodo.7566301).

ORCID iDs

Bruce Edelman [ID](https://orcid.org/0000-0001-7648-1689) <https://orcid.org/0000-0001-7648-1689>
 Ben Farr [ID](https://orcid.org/0000-0002-2916-9200) <https://orcid.org/0000-0002-2916-9200>
 Zoheyr Doctor [ID](https://orcid.org/0000-0002-2077-4914) <https://orcid.org/0000-0002-2077-4914>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, *PhRvX*, 9, 031040
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, *ApJL*, 882, L24
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020a, *LRR*, 23, 3
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2020b, *ApJL*, 896, L44
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021a, *PhRvX*, 11, 021053
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021b, *ApJL*, 913, L7
- Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, *CQGra*, 32, 024001
- Akutsu, T., Ando, M., Arai, K., et al. 2021, *PTEP*, 2021, 05A102
- Antonini, F., & Gieles, M. 2020, *PhRvD*, 102, 123016
- Astropy Collaboration, Price-Whelan, A. M., Sipocz, B. M., et al. 2018, *AJ*, 156, 123
- Bavera, S. S., Fragos, T., Qin, Y., et al. 2020, *A&A*, 635, A97
- Bavera, S. S., Fragos, T., Zevin, M., et al. 2021, *A&A*, 647, A153
- Belczynski, K. 2020, *ApJL*, 905, L15
- Belczynski, K., Heger, A., Gladysz, W., et al. 2016, *A&A*, 594, A97
- Bingham, E., Chen, J. P., Jankowiak, M., et al. 2019, *JMLR*, 20, 1
- Callister, T. A., Miller, S. J., Chatzioannou, K., & Farr, W. M. 2022, *ApJL*, 937, L13
- Croon, D., McDermott, S. D., & Sakstein, J. 2020, *PhRvD*, 102, 115024
- Cruz-Ororio, A., Lora-Clavijo, F. D., & Herdeiro, C. 2021, *JCAP*, 2021, 032
- de Boor, C. 1978, *A Practical Guide to Splines* (Berlin: Springer)
- Doctor, Z., Farr, B., & Holz, D. E. 2021, *ApJL*, 914, L18
- Doctor, Z., Farr, B., Holz, D. E., & Pürrer, M. 2017, *PhRvD*, 96, 123011
- Doctor, Z., Wysocki, D., O'Shaughnessy, R., Holz, D. E., & Farr, B. 2020, *ApJ*, 893, 35

Edelman, B., Doctor, Z., Godfrey, J., & Farr, B. 2022, *ApJ*, **924**, 101

Edelman, B., Rivera-Paleo, F. J., Merritt, J. D., et al. 2021, *PhRvD*, **103**, 042004

Edwards, M. C., Meyer, R., & Christensen, N. 2018, *Stat. Comput.*, **29**, 67

Eilers, P. H. C., & Marx, B. D. 2021, Practical Smoothing: The Joys of P-splines (Cambridge: Cambridge Univ. Press)

Essick, R., Farah, A., Galauadage, S., et al. 2022, *ApJ*, **926**, 34

Ezquiaga, J. M., & Holz, D. E. 2022, *PhRvL*, **129**, 061102

Farah, A., Fishbach, M., Essick, R., Holz, D. E., & Galauadage, S. 2022, *ApJ*, **931**, 108

Farmer, R., Renzo, M., de Mink, S. E., Fishbach, M., & Justham, S. 2020, *ApJL*, **902**, L36

Farmer, R., Renzo, M., de Mink, S. E., Marchant, P., & Justham, S. 2019, *ApJ*, **887**, 53

Farr, B., Holz, D. E., & Farr, W. M. 2018, *ApJL*, **854**, L9

Farr, W. M. 2019, *RNAAS*, **3**, 66

Farr, W. M., Farr, B., & Littenberg, T. 2015, Modelling Calibration Errors in CBC Waveforms, LIGO Document, <https://dcc.ligo.org/LIGO-T1400682/public>, T1400682-v1

Farr, W. M., Fishbach, M., Ye, J., & Holz, D. E. 2019, *ApJL*, **883**, L42

Farr, W. M., Stevenson, S., Miller, M. C., et al. 2017, *Natur*, **548**, 426

Farrell, E., Groh, J. H., Hirschi, R., et al. 2021, *MNRAS*, **502**, L40

Finke, A., Foffa, S., Iacovelli, F., Maggiore, M., & Mancarella, M. 2022, *PDU*, **36**, 100994

Fishbach, M., Essick, R., & Holz, D. E. 2020, *ApJL*, **899**, L8

Fishbach, M., & Holz, D. E. 2017, *ApJL*, **851**, L25

Fishbach, M., Holz, D. E., & Farr, W. M. 2018, *ApJL*, **863**, L41

Fishbach, M., Kimball, C., & Kalogera, V. 2022, *ApJL*, **935**, L26

Ford, K. E. S., & McKernan, B. 2022, *MNRAS*, **517**, 5827

Fuller, J., & Ma, L. 2019, *ApJL*, **881**, L1

Galauadage, S., Talbot, C., Nagar, T., et al. 2021, *ApJL*, **921**, L15

Gerosa, D., Berti, E., O'Shaughnessy, R., et al. 2018, *PhRvD*, **98**, 084036

Golomb, J., & Talbot, C. 2022a, *ApJ*, **926**, 79

Golomb, J., & Talbot, C. 2022b, arXiv:2210.12287

Hannam, M., Schmidt, P., Bohé, A., et al. 2014, *PhRvL*, **113**, 151101

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357

Heger, A., Fryer, C. L., Woosley, S. E., Langer, N., & Hartmann, D. H. 2003, *ApJ*, **591**, 288

Heger, A., & Woosley, S. E. 2002, *ApJ*, **567**, 532

Hogg, D. W. 1999, arXiv:9905116

Hong, J., Vesperini, E., Askar, A., et al. 2018, *MNRAS*, **480**, 5645

Hunter, J. D. 2007, *CSE*, **9**, 90

Jullion, A., & Lambert, P. 2007, *Comput. Stat. Data Anal.*, **51**, 2542

Kimball, C., Talbot, C., L. Berry, C. P., et al. 2020, *ApJ*, **900**, 177

Kimball, C., Talbot, C., Berry, C. P. L., et al. 2021, *ApJL*, **915**, L35

Lagos, M., Fishbach, M., Landry, P., & Holz, D. E. 2019, *PhRvD*, **99**, 083504

Landry, P., & Read, J. S. 2021, *ApJL*, **921**, L25

Lang, S., & Brezger, A. 2004, *J. Comput. Graph. Stat.*, **13**, 183

LIGO Scientific Collaboration, Aasi, J., Abbott, B. P., et al. 2015, *CQGra*, **32**, 074001

Littenberg, T. B., & Cornish, N. J. 2015, *PhRvD*, **91**, 084034

Luger, R., Bedell, M., Foreman-Mackey, D., et al. 2021, arXiv:2110.06271

LVK Collaboration 2021a, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run – Parameter estimation data release, v1, Zenodo, doi:10.5281/zenodo.5546663

LVK Collaboration 2021b, The population of merging compact binaries inferred using gravitational waves through GWTC-3—Data release, v1, Zenodo, doi:10.5281/zenodo.5655785

LVK Collaboration 2021c, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run – O1+O2+O3 Search Sensitivity Estimates, v1, Zenodo, doi:10.5281/zenodo.5636816

Madau, P., & Dickinson, M. 2014, *ARA&A*, **52**, 415

Mancarella, M., Genoud-Prachex, E., & Maggiore, M. 2022, *PhRvD*, **105**, 064030

Mandel, I., Farr, W. M., Colonna, A., et al. 2017, *MNRAS*, **465**, 3254

Mandel, I., Farr, W. M., & Gair, J. R. 2019, *MNRAS*, **486**, 1086

Marchant, P., Renzo, M., Farmer, R., et al. 2019, *ApJ*, **882**, 36

McKernan, B., Ford, K. E. S., O'Shaughnessy, R., & Wysocki, D. 2020, *MNRAS*, **494**, 1203

Miller, S., Callister, T. A., & Farr, W. M. 2020, *ApJ*, **895**, 128

Mould, M., Gerosa, D., Broekgaarden, F. S., & Steinle, N. 2022, *MNRAS*, **517**, 2738

Ng, K., Vitale, S., Hannuksela, O. A., & Li, T. G. F. 2021a, *PhRvL*, **126**, 151102

Ng, K. K. Y., Franciolini, G., Berti, E., et al. 2022, *ApJL*, **933**, L41

Ng, K. K. Y., Hannuksela, O. A., Vitale, S., & Li, T. G. F. 2021b, *PhRvD*, **103**, 063010

Ng, K. K. Y., Vitale, S., Farr, W. M., & Rodriguez, C. L. 2021c, *ApJL*, **913**, L5

Okounkova, M., Farr, W. M., Isi, M., & Stein, L. C. 2022, *PhRvD*, **106**, 044067

Ossokine, S., Buonanno, A., Marsat, S., et al. 2020, *PhRvD*, **102**, 044055

Pan, Y., Buonanno, A., Taracchini, A., et al. 2014, *PhRvD*, **89**, 084006

Payne, E., & Thrane, E. 2022, arXiv:2210.11641

Phan, D., Pradhan, N., & Jankowiak, M. 2019, arXiv:1912.11554

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, **594**, A13

Pratten, G., García-Quirós, C., Colleoni, M., et al. 2021, *PhRvD*, **103**, 104056

Ramsay, J. O. 1988, *StaSc*, **3**, 425

Renzo, M., Farmer, R., Justham, S., et al. 2020, *A&A*, **640**, A56

Rodriguez, C. L., Morscher, M., Pattabiraman, B., et al. 2015, *PhRvL*, **115**, 051101

Rodriguez, C. L., Zevin, M., Amaro-Seoane, P., et al. 2019, *PhRvD*, **100**, 043027

Rodriguez, C. L., Zevin, M., Pankow, C., Kalogera, V., & Rasio, F. A. 2016, *ApJL*, **832**, L2

Roulet, J., Chia, H. S., Olsen, S., et al. 2021, *PhRvD*, **104**, 083010

Sakstein, J., Croon, D., McDermott, S. D., Straight, M. C., & Baxter, E. J. 2020, *PhRvL*, **125**, 261105

Secunda, A., Bellovary, J., Mac Low, M.-M., et al. 2020, *ApJ*, **903**, 133

Spera, M., & Mapelli, M. 2017, *MNRAS*, **470**, 4739

Stevenson, S., Sampson, M., Powell, J., et al. 2019, *ApJ*, **882**, 121

Tagawa, H., Kocsis, B., Haiman, Z., et al. 2021, *ApJ*, **908**, 194

Talbot, C., & Thrane, E. 2017, *PhRvD*, **96**, 023012

Talbot, C., & Thrane, E. 2018, *ApJ*, **856**, 173

Taracchini, A., Buonanno, A., Pan, Y., et al. 2014, *PhRvD*, **89**, 061502

The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. 2021a, arXiv:2111.03606

The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. 2021b, arXiv:2111.03634

The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. 2021c, arXiv:2111.03604

The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. 2021d, arXiv:2108.01045

Tiwari, V. 2021, *CQGra*, **38**, 155007

Tiwari, V. 2022, *ApJ*, **928**, 155

Tiwari, V., & Fairhurst, S. 2021, *ApJL*, **913**, L19

Tong, H., Galauadage, S., & Thrane, E. 2022, arXiv:2209.02206

van Son, L. A. C., de Mink, S. E., Renzo, M., et al. 2022a, *ApJ*, **940**, 184

van Son, L. A. C., de Mink, S. E., Callister, T., et al. 2022b, *ApJ*, **931**, 17

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261

Vitale, S., Biscoveanu, S., & Talbot, C. 2022, *A&A*, **668**, L2

Vitale, S., Haster, C.-J., Sun, L., et al. 2021, *PhRvD*, **103**, 063016

Vitale, S., Lynch, R., Raymond, V., et al. 2017a, *PhRvD*, **95**, 064053

Vitale, S., Lynch, R., Sturani, R., & Graff, P. 2017b, *CQGra*, **34**, 03LT01

Vitale, S., Lynch, R., Veitch, J., Raymond, V., & Sturani, R. 2014, *PhRvL*, **112**, 251101

Woosley, S. E. 2017, *ApJ*, **836**, 244

Woosley, S. E. 2019, *ApJ*, **878**, 49

Woosley, S. E., Heger, A., & Weaver, T. A. 2002, *RvMP*, **74**, 1015

Wysocki, D., Lange, J., & O'Shaughnessy, R. 2019, *PhRvD*, **100**, 043012

Yang, Y., Bartos, I., Haiman, Z., et al. 2019, *ApJ*, **876**, 122

Zevin, M., & Bavera, S. S. 2022, *ApJ*, **933**, 86

Zevin, M., Pankow, C., Rodriguez, C. L., et al. 2017, *ApJ*, **846**, 82