Building better spin models for merging binary black holes:
Evidence for non-spinning and rapidly spinning nearly aligned sub-populations

Shanika Galaudage,[1,2] Colm Talbot,[3] Tushar Nagar,[1,2] Deepnika Jain,[4] Eric Thrane,[1,2] and Ilya Mandel[1,2,5]

[1]School of Physics and Astronomy, Monash University, Clayton VIC 3800, Australia
[2]OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Clayton VIC 3800, Australia
[3]LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA
[4]National Institute of Technology Karnataka, Surathkal, Srinivasnagar, 57025, India
[5]Institute of Gravitational Wave Astronomy and School of Physics and Astronomy,
University of Birmingham, Birmingham, B15 2TT, United Kingdom

(Dated: September 8, 2021)

ABSTRACT

Recent work paints a conflicting portrait of the distribution of black hole spins in merging binaries measured with gravitational waves. Some analyses find that a significant fraction of merging binaries contain at least one black hole with a spin tilt $> 90°$ with respect to the orbital angular momentum vector, which has been interpreted as a signature for dynamical assembly. Other analyses find the data are consistent with a bimodal population in which some binaries contain black holes with negligible spin while the rest contain black holes with spin vectors preferentially aligned with the orbital angular momentum vector. In this work, we scrutinize models for the distribution of black hole spins to pinpoint possible failure modes in which the model yields a faulty conclusion. We reanalyze data from the second LIGO–Virgo gravitational-wave transient catalog (GWTC-2) using a revised spin model, which allows for a sub-population of black holes with negligible spins. In agreement with recent results by Roulet et al., we show that the GWTC-2 detections are consistent with two distinct sub-populations. We estimate that $70 - 90\%$ (90% credible interval) of merging binaries contain black holes with negligible spin $\chi \approx 0$. The remaining binaries are part of a second sub-population in which the spin vectors are preferentially (but not exactly) aligned to the orbital angular momentum. The black holes in this second sub-population are characterized by spins of $\chi \sim 0.5$. We suggest that the inferred spin distribution is consistent with the hypothesis that all merging binaries form via the field formation scenario.

1. INTRODUCTION

Gravitational waves from merging binaries encode information about the mass and spin of the component black holes and/or neutron stars. These properties, in turn, provides clues as to how the binary formed. Two scenarios are frequently invoked to explain merging binary black holes: field and dynamical (see Mandel & Farmer 2018; Mapelli 2021; Mandel & Broekgaarden 2021 for reviews). In the field scenario, binary black hole (BBH) systems are formed from isolated stellar binaries. In the dynamical scenario, BBH systems are assembled through interactions in dense stellar environ-

ments such as globular clusters and nuclear clusters; see Mapelli (2020) for a recent review.

These two scenarios yield distinct predictions for the distribution of black hole spin vectors. Field binaries are generally expected to form BBH systems with dimensionless spin vectors $\vec{\chi}_1, \vec{\chi}_2$ preferentially-aligned with the orbital angular momentum vector $\vec{L}$. Supernova kicks may serve to somewhat misalign the $\vec{\chi}_{1,2}$ and $\vec{L}$ vectors, but the typical misalignment angle is expected to be modest (O'Shaughnessy et al. 2017; Stevenson et al. 2017; Gerosa et al. 2018; Rodriguez et al. 2016; Bavera et al. 2019). Dynamically assembled binaries on the other hand are expected to form BBH systems with isotropically distributed spin vectors (Kalogera 2000; Mandel & O'Shaughnessy 2010; Rodriguez et al. 2016; Zevin et al. 2017; Talbot & Thrane 2017; Rodriguez et al. 2018; Doctor et al. 2019). These two predictions

Corresponding author: Shanika Galaudage
shanika.galaudage@monash.edu

can be used to estimate the fraction of merging binaries associated with each channel.

Using data from the second LIGO–Virgo gravitational-wave transient catalog (GWTC-2) (Abbott et al. 2021a), Abbott et al. (2021b) fit the distribution of $\vec{\chi}_1, \vec{\chi}_2$ with two different models, which produced qualitatively similar results. The DEFAULT model incorporates the spin magnitude model from Wysocki et al. (2019):

$$\pi(\chi_{1,2}|\alpha_\chi, \beta_\chi) = \text{Beta}(\chi_{1,2}|\alpha_\chi, \beta_\chi). \qquad (1)$$

Here, $\chi_{1,2} \equiv |\vec{\chi}_{1,2}|$ are the magnitudes of the dimensionless spin vectors, $\pi(...)$ denotes a prior probability density function, and α_χ, β_χ are hyper-parameters[1] controlling the shape of the Beta distribution, which is defined on the interval $\chi \in [0,1)$ and constrained to only allow non-singular distributions ($\alpha_\chi, \beta_\chi \geq 1$). The spin tilt distribution, meanwhile, is from Talbot & Thrane (2017):

$$\pi(z_{1,2}|\zeta, \sigma_t) = \zeta \, G_t(z_{1,2}|\sigma_t) + (1 - \zeta) \, \mathfrak{I}(z_{1,2}), \qquad (2)$$

where $z_{1,2} \equiv \cos\theta_{1,2}$ and $\theta_{1,2}$ are the misalignment angles between the orbital angular momentum vector and the respective spin vectors. Here G_t is a Gaussian distribution with a peak at $z = 1$ (aligned) and a width of σ_t, truncated at $z = [-1, 1]$. This sub-model represents a population of preferentially aligned field binaries. Meanwhile, $\mathfrak{I}$ is a uniform distribution on the interval $z \in [-1, 1]$. This sub-model represents an isotropic distribution associated with dynamical formation. The hyper-parameter ζ is a mixing fraction determining the relative importance of each sub-model.[2] The final two spin degrees of freedom ϕ_1, ϕ_2 describing the azimuthal directions of the spin vectors are assumed to be uniformly distributed.

In this paper we do not carry out calculations using the second model from Abbott et al. (2021b)—the GAUSSIAN model (Roulet & Zaldarriaga 2019; Miller et al. 2020). However, we argue below that it shares key features with the DEFAULT model. We therefore expect that our findings obtained with the DEFAULT model are applicable to results obtained with the GAUSSIAN model as well.

Using the DEFAULT model, Abbott et al. (2021b) reconstruct the distribution for the effective inspiral spin

parameter

$$\chi_{\text{eff}} \equiv \frac{\chi_1 \cos\theta_1 + q\chi_2 \cos\theta_2}{1 + q}, \qquad (3)$$

where $q \equiv m_2/m_1 \leq 1$ is the mass ratio, the distribution of which is modeled according to a power law:

$$\pi(q|\beta_q) \propto q^{\beta_q}. \qquad (4)$$

The χ_{eff} parameter is frequently used to interpret results because it is an approximate constant of motion in precessing binaries and it is a relatively well-measured quantity. Abbott et al. (2021b) find that $12 - 44\%$ of BBH systems possess $\chi_{\text{eff}} < 0$ (90% credible interval), indicating that at least one black hole tilt angle $> 90°$.

However, this finding is disputed by Roulet et al. (2021). These authors analyze a slightly different set of events (chosen to be more confidently detected) using a distinct set of posterior samples. They carry out two analyses of black hole spin. We focus for the moment on their "model-free" analysis, which assumes that χ_{eff} is uniformly distributed on the interval $[-1, 1]$. The prior for the poorly-measured variable

$$\chi_{\text{diff}} \equiv \frac{\chi_1 \cos\theta_1 - \chi_2 \cos\theta_2}{1 + q}, \qquad (5)$$

is chosen to be uniform, conditioned on χ_{eff} and subject to the constraint that $\chi_{1,2} \leq 1$ as required by general relativity.[3] This leaves four degrees of freedom χ_{ix}, χ_{iy} (where $i = 1, 2$), which are assumed to be uniformly distributed on a disk with radius $\sqrt{1 - \chi_{iz}^2}$. We call this set of assumptions ROULET+.

Using these distributions, Roulet et al. (2021) plot the credible intervals for χ_{eff} for individual events in their catalog. We reproduce this "dot plot" in Fig. 1 for the events in Abbott et al. (2021b), using posterior samples from that study [4]. Individual event posteriors computed using priors from different population models are represented with different colors: black is the fiducial LIGO–Virgo spin prior (uniform in χ_1, χ_2 with an isotropic prior for the spin directions), green is the population model from Roulet et al. (2021), and pink is the DEFAULT model from Abbott et al. (2021b) (averaged over the posterior on the population parameters).

By inspecting a version of this plot, Roulet et al. (2021) find there is not clear support for a population of events with $\chi_{\text{eff}} < 0$. Focusing on the DEFAULT model's pink points, we observe there are only two events with

---

[1] For convenience, we later re-parameterize the Beta distribution in terms of the spin magnitude mean μ_χ and standard deviation σ_χ.

[2] Other possible formation channels (e.g., hierarchical triple systems, formation in active galactic nuclei, primordial black holes) have distinct spin predictions. However, this model focuses on the isolated binary and dynamical formation channels.

[3] Given these constraints, the marginalised prior on χ_{diff} is not actually uniform.

[4] Posterior samples from dcc.ligo.org/LIGO-P2000223/public

a weak statistical preference for $\chi_{\rm eff} < 0$. We therefore concur with Roulet et al. (2021) that the BBH mergers observed to date are consistent with a sub-population of negligible-spin events $\chi_{1,2} \approx 0$ with a second sub-population of events with $\chi_{\rm eff} > 0$. The signature from Abbott et al. (2021b) for a sub-population of $\chi_{\rm eff} < 0$ seen by Abbott et al. (2021b) is likely due to a misspecification of the assumed population spin model. This paper therefore seeks to improve on the DEFAULT model by developing a more sophisticated model for black hole spin in order to better describe the data and to provide more reliable inferences.

The remainder of this paper is structured as follows. In Section 2, we discuss why checks in Abbott et al. (2021b), designed to validate the observation of a sub-population with $\chi_{\rm eff} < 0$, did not reveal the model-dependency of this result. In Section 3, we propose an improved black-hole spin model (called EXTENDED) designed to address the shortcomings of the DEFAULT model. We repeat the black hole spin distribution analysis from Abbott et al. (2021b) (Section 4) using the EXTENDED model and report results in Section 5. In agreement with Roulet et al. (2021), we show that the the observation of a sub-population with $\chi_{\rm eff} < 0$ is model-dependent. We present revised fits to the distribution of black hole spins. We conclude in Section 6 by discussing strategies for future development and testing of black hole spin models.

2. DIAGNOSING LIMITATIONS OF THE DEFAULT MODEL

Abbott et al. (2021b) include several checks, which would seem at first glance to support the claim for a sub-population of events with $\chi_{\rm eff} < 0$. First, they note the similar reconstructed distributions for $\chi_{\rm eff}$ obtained with two different models: the DEFAULT model described above and GAUSSIAN model that describes the distribution of $\chi_{\rm eff}$ and the effective precession spin parameter χ_p using a multi-variate Gaussian (see their Fig. 11b). The potential for model-induced systematic error is reduced given the signature is visible using two different models. Second, they carry out posterior predictive checks to show that mock source populations drawn from the GAUSSIAN model would have similar distributions of $\chi_{\rm eff}$ and χ_p to the observed data; see their Fig. 26. Third, they perform an analysis in which a new hyper-parameter $\chi_{\rm eff}^{\rm min}$ is added to the population model, which enforces a minimum value of $\chi_{\rm eff} > \chi_{\rm eff}^{\rm min}$; see their Fig. 27. A negative value for $\chi_{\rm eff}^{\rm min}$ is preferred over $\chi_{\rm eff}^{\rm min} = 0$, seemingly bolstering the case for a sub-population of events with $\chi_{\rm eff} < 0$.
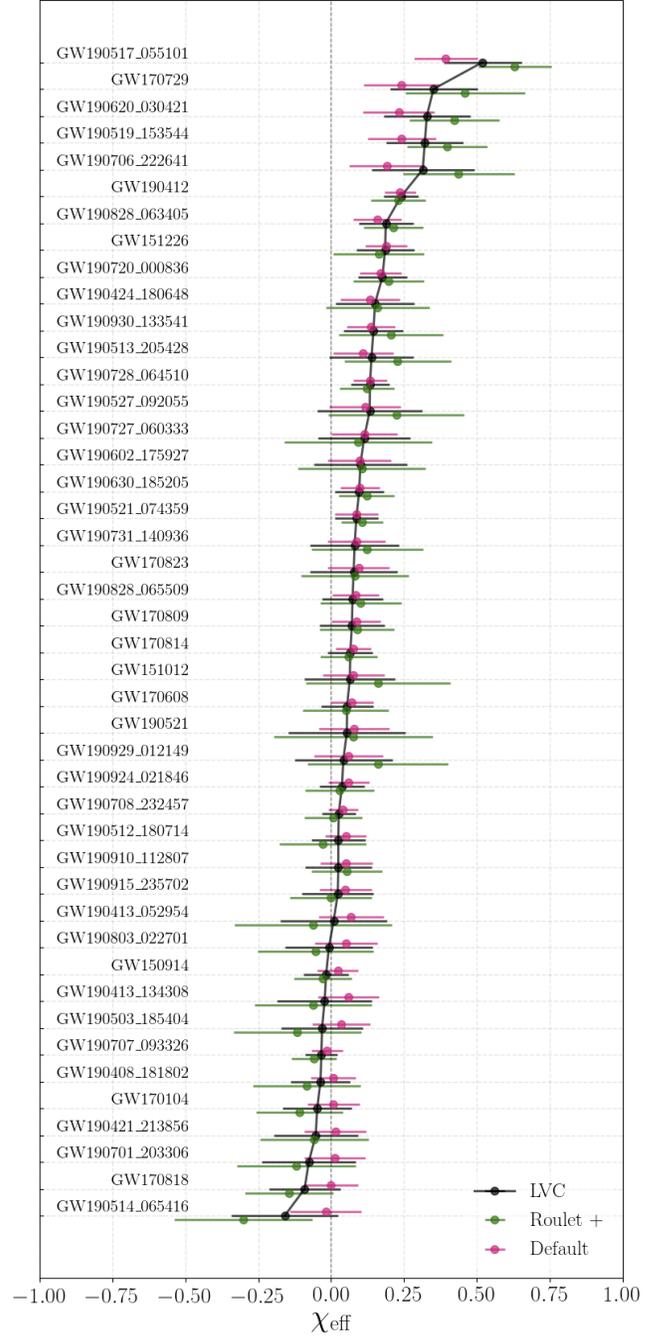

Figure 1. A "dot plot" showing the mean and one-sigma credible interval for $\chi_{\rm eff}$ assuming different population models. The fiducial model used by LIGO–Virgo (LVC, black) assumes uniform distribution for dimensionless spins χ_1, χ_2 with an isotropic distribution for the spin directions. Green (ROULET+) shows the model from Roulet et al. (2021). The DEFAULT model from Abbott et al. (2021b) is shown in pink. No models indicate support for a sub-population of events with $\chi_{\rm eff} < 0$.

It is instructive to discuss how each of these checks fails to catch the apparent lack of $\chi_{eff} < 0$ events in Fig. 1. While the agreement between the DEFAULT and GAUSSIAN models provides something of a sanity check, both models lack a key feature: the ability to account for an excess of BBH systems with $\chi_{1,2} \approx 0$. The possibility that some fraction of LIGO–Virgo binaries should merge with negligible spin is supported by theoretical studies of angular momentum transport (Fuller & Ma 2019; Belczynski et al. 2020). Moreover, the negligible spin of many/most LIGO–Virgo binaries has been noted observationally (Miller et al. 2020; Kimball et al. 2021). Neither spin model from Abbott et al. (2021b) accounts for a sub-population of binaries with $\chi \approx 0$ black holes. Since neither model can accommodate an excess of events with $\chi_{1,2} \approx 0$, the models are liable to fit such an excess with the next best thing: a subpopulation of events distributed *about* $\chi_{eff} = 0$. In this way, the DEFAULT and GAUSSIAN models can both yield false positive signals for $\chi_{eff} < 0$ when the true population contains an excess of events with $\chi_{1,2} \approx 0$.

Next we turn to the posterior predictive check from Abbott et al. (2021b), their Fig. 26. This plot compares the cumulative distribution of χ_{eff} for the data to the distribution expected from the model. The two distributions are visually consistent. While a mismatch between data and model predictions would indicate a failure of the model, a match does not prove that the model is a faithful representation of the data. Using a different population model, Roulet et al. (2021) find no evidence for a sub-population of events with $\chi_{eff} < 0$—a finding that we confirm independently below.

Finally, we consider the test from Abbott et al. (2021b) showing that the data prefer a negative values of χ_{eff}^{min} such that some events in the population are characterized by negative χ_{eff} values on the interval $(\chi_{eff}^{min}, 0)$ (see their Fig. 27). It should now be apparent that this test can be tricked if the data contain an excess of events with $\chi_{1,2} \approx 0$. In order to illustrate this—and to build an improved model for BBH spin— we introduce an EXTENDED model, which incorporates elements from Roulet et al. (2021) that enable an excess of negligible-spin events.

3. THE EXTENDED MODEL

We introduce two changes to the DEFAULT model. First, we add a new population parameter λ_0, corresponding to the fraction of BBH mergers with negligible spin. Our revised spin magnitude distribution is

$$\pi(\chi_{1,2}|\alpha_\chi, \beta_\chi, \lambda_0) = (1-\lambda_0)\text{Beta}(\chi_{1,2}|\alpha_\chi, \beta_\chi) + \lambda_0 G_t(\chi_{1,2} \mid \mu = 0, \ \sigma_0), \quad (6)$$

where G_t is a truncated Gaussian distribution peaking at $\chi_{1,2} = 0$ with width σ_0. We set $\sigma_0 = 0$ so that the spin magnitudes are precisely zero and carry out the analysis using a dedicated set of zero-spin posterior samples. Setting $\chi = 0$ is probably a good approximation since Fuller & Ma (2019) suggest that typical black-hole spin magnitudes are $\chi \approx 0.01$ (ignoring binary effects, particularly tidally induced spin-up, discussed by, e.g., Kushnir et al. 2016; Zaldarriaga et al. 2018; Qin et al. 2018; Bavera et al. 2019; Bavera et al. 2020; Belczynski et al. 2020; Mandel & Fragos 2020). It would be interesting to make σ_0 a population parameter that can be fit with the data. However, we leave this for future work as there are some technical challenges; preliminary studies with small, non-zero values of σ_0 seem to suffer from under-sampling effects. We adopt a uniform prior for λ_0 on the interval $[0, 1]$.

Our second change is to add a population parameter z_{min} to the distribution of black-hole tilt angles such that the combined distribution from Eq. 2 is forced to zero for $z_{1,2} < z_{min}$:

$$\pi(z_{1,2}|\zeta, \sigma_t, z_{min}) \propto \left(\zeta G_t(z_{1,2}|\sigma_t) + (1-\zeta) \Im(z_{1,2}) \right) \Theta(z_{1,2} - z_{min}). \quad (7)$$

We adopt a uniform prior for z_{min} on the interval $[-1, 1]$. In section 5 we present results of an analysis of events from the GWTC-2 catalog using the EXTENDED model. The full list of population hyper-parameters and priors is given in Table 1.

4. ANALYSIS

We use GWPopulation (Talbot et al. 2019) to obtain posterior distributions for the population parameters in the EXTENDED model using the same dataset used in Abbott et al. (2021b). This dataset consists of 44 confidently detected BBH mergers. It does not include GW190814 (Abbott et al. 2020a), which may be a neutron-star black-hole binary and which is a clear outlier from the rest of the population. The GWPopulation package employs Bilby (Ashton et al. 2019; Romero-Shaw et al. 2020) and dynesty (Speagle 2020). We fit the distribution of black-hole masses using the POWER LAW + PEAK model from Abbott et al. (2021b) adapted from Talbot & Thrane (2018). We fit the merger redshift distribution using the POWER-LAW EVOLUTION model from Abbott et al. (2021b), adapted from Fishbach et al.

Parameter	Description	Prior
λ_0	Mixing fraction of mergers with negligible spin, $\chi_{1,2} \lesssim \sigma_0$	$U(0,1)$
σ_0	Spread in $\chi_{1,2}$ for systems with negligible spin	$\sigma_0 = 0$
μ_χ	Mean of spin magnitude distribution	$U(0,1)$
σ_χ^2	The square of the width of the spin magnitude distribution	$U(0,0.25)$
ζ	Mixing fraction of mergers with preferentially aligned spin	$U(0,1)$
σ_t	Spread in projected misalignment for preferentially aligned black holes	$U(0,4)$
$z_{\min}$	Minimum value of the projected misalignment	$U(-1,1)$

Table 1. Summary of EXTENDED model hyper-parameters. The notation $U(a,b)$ indicates a uniform distribution on the interval ranging from a to b.

(2018). We employ two sets of posterior samples[5]: one using the standard LIGO–Virgo priors and one generated using a zero-spin prior; see (Kimball et al. 2021).[6] The zero-spin samples are necessary to avoid artifacts due to under-sampling. Additional technical details are provided in Appendix A.

We adopt the same treatment of selection effects as used in Abbott et al. (2021b), which accounts for mass-dependent Malmquist bias, but not selection effects due to spin. The authors of that work explain that they cannot reliably estimate spin-induced selection effects using the currently available injection set; see their Appendix F. It is slightly easier to detect BBH signals with $\chi_{\rm eff} > 0$, so we expect that our fit may slightly over-emphasize high positive values of $\chi_{\rm eff}$ relative to the true distribution. Based on Fig. 1 of Ng et al. (2018), we concur with Abbott et al. (2021b) that this shift is likely to be small.

5. RESULTS

We plot population predictive distributions (PPDs)[7] comparing the DEFAULT and EXTENDED models in Fig. 2. In the left panel, we show the reconstructed spin magnitude distribution, the EXTENDED model exhibits clear support for a narrow peak at $\chi_{1,2} = 0$, which is

---

[5] The samples used for our analysis are from LIGO-Virgo analyses and Kimball et al. (2021). The combination we used is available at https://github.com/shanikagalaudage/bbh_spin

[6] Astute readers may notice that Kimball et al. (2021) reports that the fraction of BBH systems with negligible spin is consistent with zero $\lambda_0 \approx 0$. In that study, the BBH events with measurable spin are likely attributed to a sub-population containing one or more rapidly spinning "second-generation" black holes (formed from previous mergers). Thus, they likely do not influence the fit for first-generation dimensionless spin parameters, α_χ, β_χ, which are found to be consistent with a population of low-spin black holes.

[7] The PPD is given by the conditional prior marginalized over the posterior distribution of the population parameter

$$p_\Lambda(\chi_{1,2}|d) = \int d\Lambda \, p(\Lambda|d) \, \pi(\chi_{1,2}|\Lambda). \qquad (8)$$

Here, Λ are population parameters described in Table 1.

Spin model	$\log_{10}\mathcal{B}$	$\Delta \log_{10}\mathcal{L}_{\max}$
DEFAULT	0.00	0.00
EXTENDED	17.55	17.32

Table 2. Log Bayes factors and maximum log likelihood differences for different spin models compared to the DEFAULT model.

not present in the DEFAULT model because the model lacks the flexibility to fit this peak. In the right panel, we show the reconstructed distribution of the (cosine of the) spin tilts, we see that the EXTENDED model distribution of $z_{1,2} \equiv \cos\theta_{1,2}$ tapers off for $\cos\theta_{1,2} \lesssim 0$ while the DEFAULT model exhibits considerably more support for binaries with $\cos\theta_{1,2} < 0$. Since the EXTENDED model includes the DEFAULT model as a special case (with $\lambda_0 = 0, z_{\min} = -1$), we conclude that the data prefer a sub-population of binaries with $\chi_{1,2} \approx 0$ over a sub-population of binaries with $\chi_{\rm eff} < 0$. This conclusion is supported by our model selection results, summarized in Table 2, which show the EXTENDED model is preferred over the DEFAULT model by $\log_{10}\mathcal{B} \approx 17.6$.

Figure 3 shows the posterior distributions for the two variables new to the EXTENDED model: $z_{\min}$ (left) and λ_0 (right). Turning first to the posterior for $z_{\min}$, we find ample posterior support for $z_{\min} = 0$, which means the data are consistent with the premise that all merging binaries have $\chi_{\rm eff} > 0$. Next we turn our attention to the posterior distribution for λ_0 (right), the parameter that controls the fraction of binaries merging with negligible black-hole spin. The distribution shows that most BBH systems merge with negligible spins: $\lambda_0 = 0.81^{+0.09}_{-0.11}$ (90% credibility).

In Fig. 4 we include a corner plot showing the posterior distribution for all the population parameters in the EXTENDED model.[8] The navy contours mark the one-, two-, and three-sigma credible intervals for the Ex-

---

[8] In the companion repository, we include an updated corner plot for the mass population parameters. The results are similar to those obtained with the DEFAULT model, though, we note that
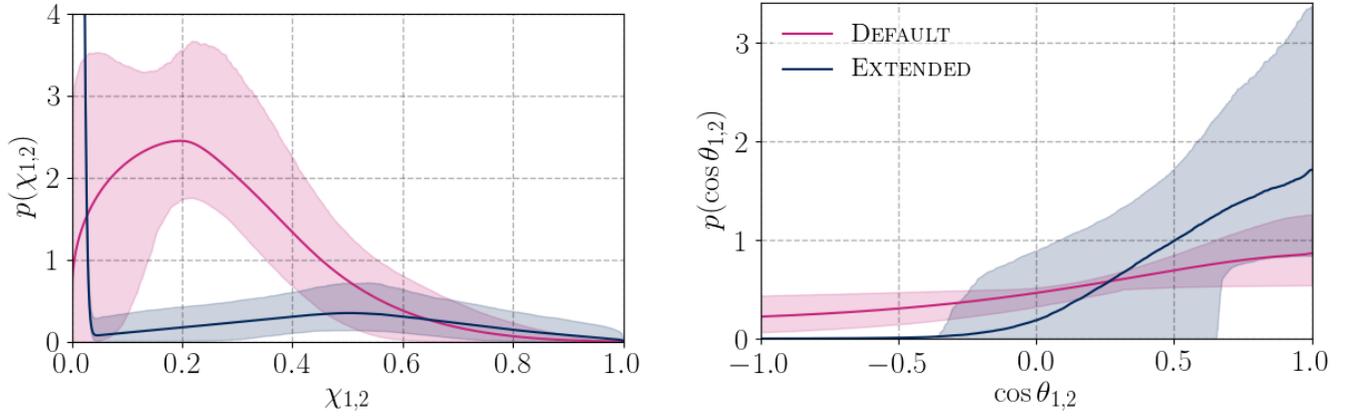
Figure 2. Population predictive distributions for the DEFAULT (pink) and the EXTENDED (navy) models. *Left:* The distribution of dimensionless spin. *Right:* The distribution of the cosine of the tilt angle. The solid curves represent the mean and the shaded region represents the 90% credible interval. In the left-hand panel, we represent the EXTENDED model's delta function at $\chi = 0$ with a narrow Gaussian with a width of $\sigma_0 = 0.01$ for visibility and is zoomed in.
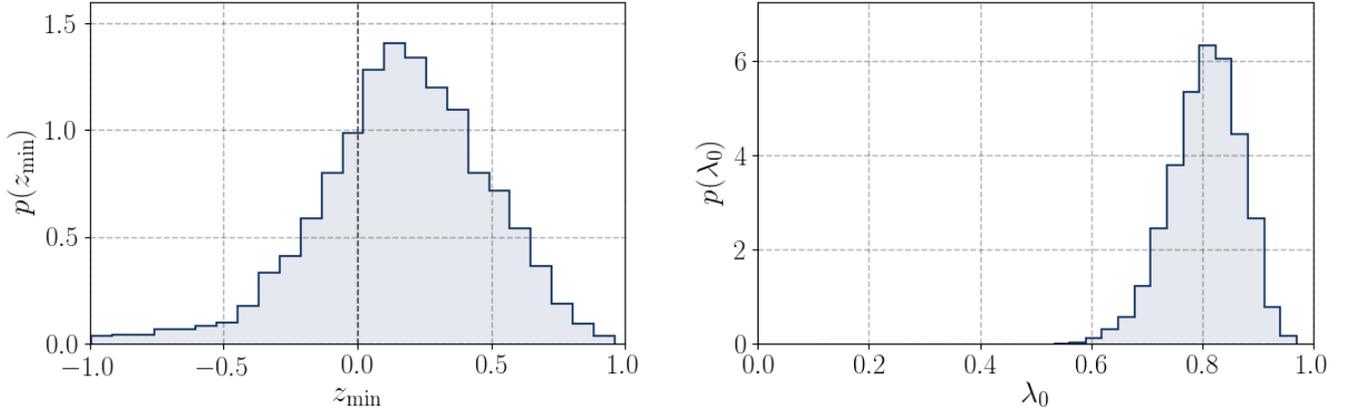

Figure 3. Posterior distributions for population parameters new to the EXTENDED model. *Left:* The posterior distributions for the population parameter $z_{\min}$, which determines the minimum cosine tilt angle such that $\cos\theta_{1,2} \geq z_{\min}$. *Right:* The posterior distribution for λ_0, the fraction of binaries with negligible black-hole spins $\chi_{1,2} = 0$.

TENDED model while the pink contours show the results for the DEFAULT model used in Abbott et al. (2021b). There are a number of interesting differences between the two models. First, the μ_χ parameter, which determines the average dimensionless spin magnitude for black holes (those with non-negligible spin in the EXTENDED model), shifts from ≈ 0.25 for DEFAULT to ≈ 0.5 for EXTENDED. This shift reflects the fact that dimensionless spin is higher when we allow for a separate population of binaries with negligible black-hole spin. Taken together, Figs. 3 and 4 suggest two subpopulations: the majority of BBH systems merge with negligible spin while a minority merge with large spins.

the peak at $\approx 35 M_\odot$ shifts higher to $\approx 38 M_\odot$ and becomes slightly narrower.

Second, the posterior for the ζ parameter (shown in Fig. 4), which determines the fraction of binaries with preferentially aligned spin, becomes broader, approaching a uniform distribution. This change is explained by the introduction of the $z_{\min}$ parameter, which provides a new means of building a population of preferentially aligned binaries.[9] Both models disfavor perfect alignment of the black-hole spin vectors with the orbital angular momentum. The DEFAULT model shows support for a preferentially aligned population ($\zeta > 0$) but with non-zero misalignment angle spread ($\sigma_t > 0$).

---

[9] Consider, for example, the case where $z_{\min} \approx 0$ and $\zeta = 0$, which is well-supported by the data. This case corresponds to a population where black hole tilts are "half-isotropically" distributed for angles $\theta_{1,2} < 90°$, but black holes never merge with tilt angles $> 90°$.

Meanwhile, the EXTENDED model favors a broad range of $z_{min} = (-0.36, 0.66)$ at 90% credibility, so the misalignment distribution is likely not cut off at near-exact alignment.

Finally, in Fig. 5, we show the χ_{eff} PPDs for the DEFAULT and EXTENDED models. The EXTENDED model is characterized by a sharp peak at $\chi_{eff} = 0$ corresponding to the sub-population of BBH systems with negligible spin. While the EXTENDED PPD does not vanish for $\chi_{eff} < 0$, there is very limited support there, with fewer than 0.2% of all binaries predicted to have $\chi_{eff} < -0.1$. This is consistent with Fig. 1, which shows no suggestion of events with negative χ_{eff}. This conclusion is also consistent with Fig. 3, which shows support for z_{min} as low as $z_{min} \approx -0.4$. However, the same plot shows that the data are consistent with $z_{min} = 0$. In other words, we cannot rule out the possibility of a small sub-population with $\chi_{eff} < 0$, but the EXTENDED model provides no evidence that such a sub-population exists.

6. DISCUSSION

Inspired by Roulet et al. (2021), we reanalyze the population of merging binary black holes from Abbott et al. (2021b) using a revised model for the distribution of black hole spins designed to allow for a sub-population with negligible spin. Using the same events and the same posterior samples as Abbott et al. (2021b), we obtain results qualitatively similar to Roulet et al. (2021), suggesting that Abbott et al. (2021b)'s finding of a sub-population of binaries merging with $\chi_{eff} < 0$ is model-dependent. Using our new EXTENDED model, we find that the data can be explained by the hypothesis that all binaries merge with black hole spin preferentially aligned with the orbital angular momentum *or* with negligible spin. These findings are consistent with GWTC-1 studies (Farr et al. 2017, 2018; Abbott et al. 2019) inferring the population of what was then 10 BBH detections is consistent with mostly negligible spins and a few events with support for $\chi_{eff} > 0$.

This result somewhat diminishes the case for dynamical mergers as a major channel for merging binaries. However, the dynamical scenario remains a plausible explanation for binary mergers with black holes in the pair instability mass gap (Abbott et al. 2020b). If GW190521 is a hierarchical merger (assembled from the products of previous mergers), perhaps black holes participating in the first generation of dynamical mergers are among the $\approx 80\%$ of with negligible spin. Indeed, studies of dynamical mergers in dense clusters find that small spins are required for first-generation black holes in order for their remnants to be retained to merge again (Kimball et al. 2020, 2021). If a large fraction of merging bina-

ries are assembled dynamically, our results suggest that most dynamically merging black holes have negligible spin.

A sub-population of binaries with negligible black-hole spins can also be accommodated within the field formation framework; see, e.g., Belczynski et al. (2020); Bavera et al. (2020). If the stellar progenitors of BBH systems efficiently shed angular momentum through mass loss, it may be common for most field BBH systems to form with negligible spin. A minority of progenitors in sufficiently tight binaries could, however, tidally spin up the stellar core of the secondary star prior to its collapse into a black hole, producing a BBH system with at least one rapidly spinning BH. In this framework, black holes with non-negligible spins could be those whose naked helium star progenitors were tidally spun up in tight binaries. This possibility can be explored in the future by relaxing the assumption that the spins and spin-orbit misalignment angles of both binary components are independently drawn from the same distributions.

Binary black hole detections in which the more massive black hole has clearly non-zero spin $\chi_1 > 0$ could be challenging to accommodate in the standard field framework as described above (Mandel & Fragos 2020; Qin et al. 2021). Several events in GWTC-2 would seem to fall into this category, including GW151226 (Abbott et al. 2016; Chia et al. 2021, but see Mateu-Lucena et al. 2021), GW190412 (Abbott et al. 2020c; Zevin et al. 2020, but see Mandel & Fragos 2020), and GW190403_051519 (Abbott et al. 2021c; Qin et al. 2021, though the high mass already makes field formation unlikely for this system). If these systems formed in the field, it is possible that the more massive black hole formed from what was initially the secondary (lower mass) star that subsequently gained mass through accretion. In this scenario, the secondary star could still be tidally spun up before collapsing into the more massive black hole. Furthermore, the progenitor of the first-formed black hole can also be tidally spun up if the binary is initially very compact, as appears to be the case with high-mass X-ray binaries such as M33 X-7 (Valsecchi et al. 2010) and Cygnus X-1 (Qin et al. 2019), though it is unclear if such systems lead to merging BBHs (Neijssel et al. 2021).

Our work highlights the subtleties in diagnosing model misspecification, which can lead to overly model-dependent conclusions. While Abbott et al. (2021b) perform a number of checks to validate their evidence for a sub-population of events with $\chi_{eff} < 0$ (see Section 2), none of these checks pinpointed how the model could misbehave when applied to a distribution with a sub-population of binaries with negligible-spin black
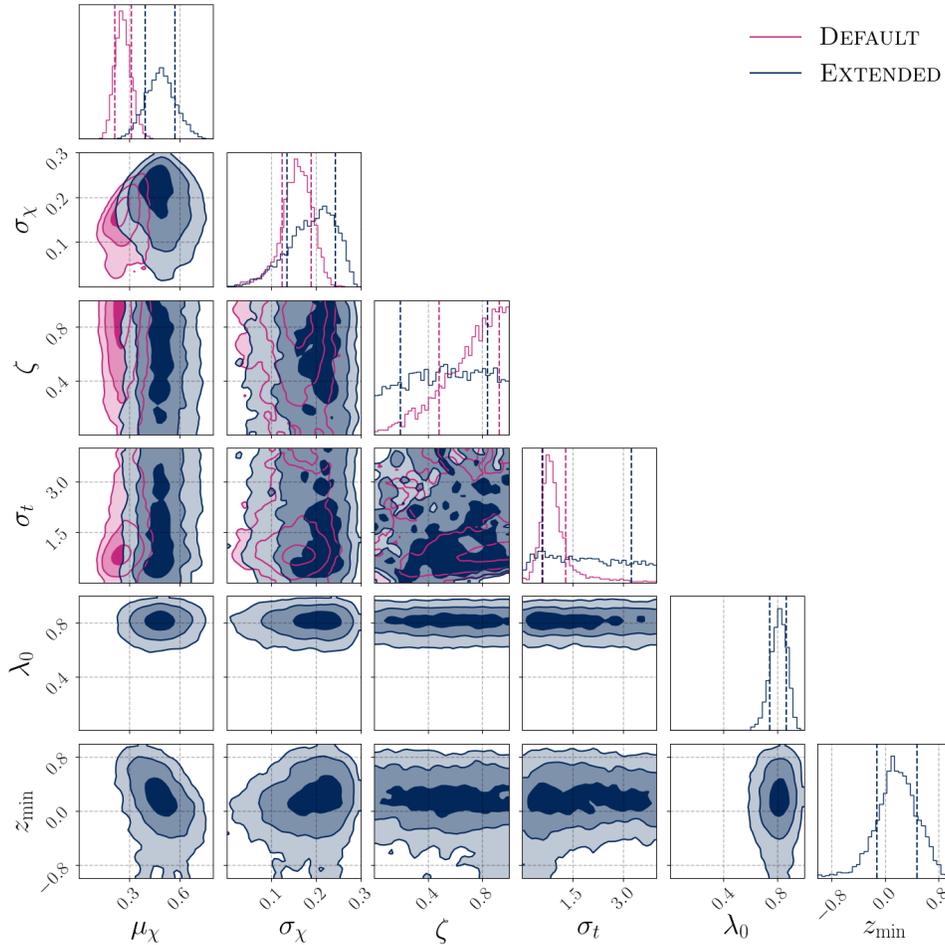
Figure 4. A corner plot showing hyper-posterior distributions for the DEFAULT (pink) and EXTENDED (navy) models. The parameters are summarized in Table 1. The shaded regions represent one-, two-, and three-sigma credible intervals.
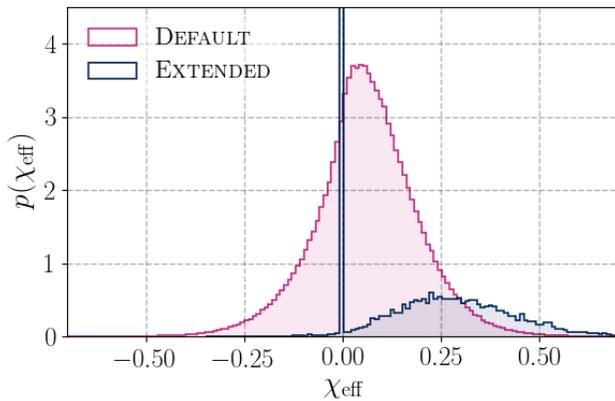

Figure 5. The population predictive distribution for the effective inspiral spin parameter $\chi_{\rm eff}$ for the DEFAULT (pink) and the EXTENDED (navy) models.

holes, as pointed out by Roulet et al. (2021). We recommend careful consideration of possible *sharp features*—in this case, a narrow peak with $\chi_{1,2} \approx 0$—since they

can yield misleading results when fit with slowly varying functions.

The "dot plot" featured in Fig. 1 of Roulet et al. (2021), and included in our own Fig. 1, provides an important visual check. However, it is not obvious that such a plot can be applied more broadly to help spot model misspecification. For one thing, the $\chi_{\rm eff}$ dot plot is fairly straightforward to interpret because the likelihood function for individual events is approximately Gaussian in this parameter. However, this is not true for other variables, for example, χ_p. Moreover, even the $\chi_{\rm eff}$ dot plot includes a hidden population model, which can affect its interpretation.

In order to obtain a prior distribution that is uniform in $\chi_{\rm eff}$, the "model-free" prescription in Roulet et al. (2021) makes implicit assumptions about the distribution of the physical spin parameters χ_1, χ_2, $\cos\theta_1$ and $\cos\theta_2$. Plotting the distribution of these spin parameters in Fig. 6, we see that some of the assumed distributions are not especially physical. In particular, the "model-
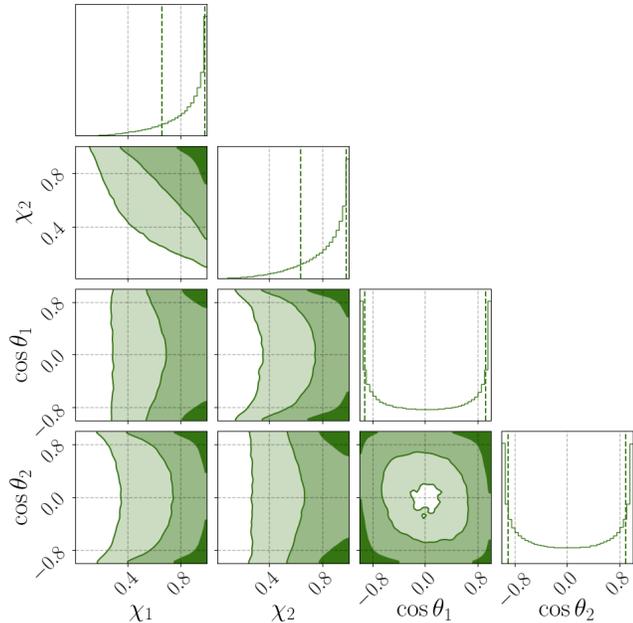
Figure 6. Prior distribution on spin magnitudes and cosine of the tilt angle for the Roulet et al. (2021) $\chi_{\rm eff}$ model.

free" analysis implicitly assumes that black holes preferentially spin near the Kerr limit $\chi_{1,2} = 1$ with a strong covariance between χ_1 and χ_2. The distribution of spin tilts, meanwhile, favours spin vectors that are either aligned or anti-aligned with the orbital angular momentum vector. Comparing the green "Roulet +" model to the black LVC model in Fig. 1, we see that the two models produce similar dot plots when one accounts for the implicit preference for larger spins in the "model-free" approach. However, the two models do not produce the same ordered list when we rank events from smallest to largest $\chi_{\rm eff}$, which indicates that other spin degrees of freedom described by the implicit model influence this plot – an effect that will become more significant as the gravitational-wave catalog grows. However, the surpris-

ing shape of the distributions in Fig. 6 reminds us that all fully specified models of black-hole spin require *some* distribution of physical parameters even if these distributions are not explicit. It therefore seems useful to cast population models in terms of physical parameters, or, at the very least, to check the suitability of the distribution of physical parameters.

There are a number of interesting extensions to this work worthy of future exploration. Our EXTENDED model is designed to investigate questions raised in Abbott et al. (2021b) and Roulet et al. (2021). A number of possible model extensions consider more astrophysically motivated distributions. As discussed above, spin magnitudes and directions may be coupled with each other, or with mass and mass ratio (see, e.g., Callister et al. 2021). For example, there could be different spin magnitude distributions in the field scenario (preferentially aligned spins) and the dynamical scenario (isotropic spins). The low-spin-magnitude subpopulation could incorporate small but non-zero spins by varying σ_0. The $z_{\rm min}$ parameter could be applied only to the sub-population of binaries with $\cos\theta_{1,2}$ described by the truncated Gaussian distribution; this would restore the mixture model choice between a preferentially aligned distribution and an isotropic distribution.

ACKNOWLEDGEMENTS

REFERENCES

Abbott, B. P., et al. 2016, Phys. Rev. Lett., 116, 241103

—. 2019, Astrophys. J. Lett., 882, L24,
doi: 10.3847/2041-8213/ab3800

Abbott, R., et al. 2020a, Astrophys. J. Lett., 896, L44

—. 2020b, Phys. Rev. Lett., 125, 101102

—. 2020c, Phys. Rev. D, 102, 043015

—. 2021a, Physical Review X, 11, 021053,
doi: 10.1103/PhysRevX.11.021053

—. 2021b, Astrophys. J. Lett., 913, L7

—. 2021c, arXiv:2108.01045

Ashton, G., et al. 2019, Astrophys. J. Supp., 241, 27

Bavera, S. S., et al. 2019, Astron. Astrophys., 635, A97

Bavera, S. S., Fragos, T., Qin, Y., et al. 2020, Astron.
Astrophys., 635, A97

Belczynski, K., Klencki, J., Fields, C. E., et al. 2020, A&A,
636, A104, doi: 10.1051/0004-6361/201936528

Callister, T. A., Haster, C.-J., Ng, K. K. Y., Vitale, S., &
Farr, W. M. 2021, arXiv e-prints, arXiv:2106.00521.
https://arxiv.org/abs/2106.00521

Chia, H. S., Olsen, S., Roulet, J., et al. 2021, arXiv e-prints, arXiv:2105.06486. https://arxiv.org/abs/2105.06486

Doctor, Z., Wysocki, D., O'Shaughnessy, R., Holz, D. E., & Farr, B. 2019, Astrophys. J., 893, 35

Farr, B., Holz, D. E., & Farr, W. M. 2018, Astrophys. J. Lett., 854, L9, doi: 10.3847/2041-8213/aaaa64

Farr, W. M., Stevenson, S., Miller, M. C., et al. 2017, 548, 426, doi: 10.1038/nature23453

Fishbach, M., Holz, D. E., & Farr, W. M. 2018, Astrophys. J. Lett., 863, L41

Fuller, J., & Ma, L. 2019, Astrophys. J. Lett., 881, L1

Gerosa, D., Berti, E., O'Shaughnessy, R., et al. 2018, Phys. Rev. D, 98, 084036

Kalogera, V. 2000, apj, 541, 319, doi: 10.1086/309400

Kimball, C., Talbot, C., Berry, C. P., et al. 2021, Astrophys. J. Lett., 915, L35

Kimball, C., Talbot, C., Berry, C. P. L., et al. 2020, Astrophys. J., 900, 177

Kushnir, D., Zaldarriaga, M., Kollmeier, J. A., & Waldman, R. 2016, MNRAS, 462, 844, doi: 10.1093/mnras/stw1684

Mandel, I., & Broekgaarden, F. S. 2021, arXiv e-prints, arXiv:2107.14239. https://arxiv.org/abs/2107.14239

Mandel, I., & Farmer, A. 2018, ArXiv e-prints. https://arxiv.org/abs/1806.05820

Mandel, I., & Fragos, T. 2020, ApJL, 895, L28, doi: 10.3847/2041-8213/ab8e41

Mandel, I., & O'Shaughnessy, R. 2010, Class. Quantum Grav., 27, 114007

Mapelli, M. 2020, Front. Astron. Space Sci., 7, 38

Mapelli, M. 2021, arXiv e-prints, arXiv:2106.00699. https://arxiv.org/abs/2106.00699

Mateu-Lucena, M., Husa, S., Colleoni, M., et al. 2021, arXiv e-prints, arXiv:2105.05960. https://arxiv.org/abs/2105.05960

Miller, S., Callister, T. A., & Farr, W. M. 2020, Astrophys. J., 895, 128

Neijssel, C. J., Vinciguerra, S., Vigna-Gómez, A., et al. 2021, ApJ, 908, 118, doi: 10.3847/1538-4357/abde4a

Ng, K. K. Y., Vitale, S., Zimmerman, A., et al. 2018, Phys. Rev. D, 98, 083007

O'Shaughnessy, R., Gerosa, D., & Wysocki, D. 2017, Phys. Rev. Lett., 119, 011101

Qin, Y., Fragos, T., Meynet, G., et al. 2018, Astron. Astrophys., 616, A28

Qin, Y., Marchant, P., Fragos, T., Meynet, G., & Kalogera, V. 2019, ApJL, 870, L18, doi: 10.3847/2041-8213/aaf97b

Qin, Y., Yuan-Zhu, Wang, Dong-Hong, & Wu. 2021, arXiv e-prints, arXiv:2108.04821. https://arxiv.org/abs/2108.04821

Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., & Rasio, F. A. 2018, Phys. Rev. Lett., 120, 151101

Rodriguez, C. L., Zevin, M., Pankow, C., Kalogera, V., & Rasio, F. A. 2016, Astrophys. J. Lett., 832, L2

Romero-Shaw, I. M., et al. 2020, Mon. Not. R. Ast. Soc., 499, 3295

Roulet, J., Chia, H. S., Olsen, S., et al. 2021

Roulet, J., & Zaldarriaga, M. 2019, Mon. Not. R. Ast. Soc., 484, 4216, doi: 10.1093/mnras/stz226

Speagle, J. S. 2020, Mon. Not. R. Ast. Soc., 493, 3132

Stevenson, S., Berry, C. P. L., & Mandel, I. 2017, Mon. Not. R. Ast. Soc., 471, 2801

Talbot, C., Smith, R. J. E., Thrane, E., & Poole, G. B. 2019, Phys. Rev. D, 100, 043030

Talbot, C., & Thrane, E. 2017, Phys. Rev. D, 96, 023012

—. 2018, Astrophys. J., 856, 173

Thrane, E., & Talbot, C. 2019, Pub. Astron. Soc. Aust., 36, E010

Valsecchi, F., Glebbeek, E., Farr, W. M., et al. 2010, Nature, 468, 77, doi: 10.1038/nature09463

Wysocki, D., Lange, J., & O'Shaughnessy, R. 2019, Phys. Rev. D, 100, 043012

Zaldarriaga, M., Kushnir, D., & Kollmeier, J. A. 2018, MNRAS, 473, 4174, doi: 10.1093/mnras/stx2577

Zevin, M., Berry, C. P. L., Coughlin, S., & Chatziioannou, K. 2020, Astrophys. J. Lett., 899, L17

Zevin, M., Pankow, C., Rodriguez, C. L., et al. 2017, Astrophys. J., 846, 82

APPENDIX

A. IMPLEMENTING THE Extended MODEL WITH ZERO-SPIN SAMPLES

In this appendix we describe our implementation of inference on the Extended model hyper-parameters using two sets of posterior samples—one allowing for black hole spins (called "fiducial" samples), the other assuming zero black hole spins. Our starting point is the likelihood for the data d associated with a single gravitational-wave event given the population hyper-parameters Λ:

$$\mathcal{L}(d|\Lambda) = \int d\chi_{1,2} \int dz_{1,2} \int d\eta \, \mathcal{L}(d|\chi_{1,2}, z_{1,2}, \eta) \, \pi(\chi_{1,2}|\alpha_\chi, \beta_\chi, \lambda_0) \, \pi(z_{1,2}|\zeta, \sigma_t, z_{\min}) \, \pi(\eta|\Lambda). \tag{A1}$$

Here, $\mathcal{L}(d|\chi_{1,2}, z_{1,2}, \eta)$ is the usual Gaussian likelihood of the data conditional on spin parameters $\chi_{1,2}, z_{1,2}$ and other parameters (mass, redshift, etc.), which we denote by η. The next term $\pi(\chi_{1,2}|\alpha_\chi, \beta_\chi, \lambda_0)$ is the Extended model for spin magnitude described in Eq. 6. This is followed by $\pi(z_{1,2}|\zeta, \sigma_t, z_{\min})$, which is the Extended model for spin tilts; see Eq. 7. The final term $\pi(\eta|\Lambda)$ describes the population model for mass and redshift as well as the usual priors for extrinsic parameters. Our mass model is the Power Law + Peak model from Abbott et al. (2021b) adopted from Talbot & Thrane (2018). Our redshift model is the Power-Law Evolution model from Abbott et al. (2021b) adopted from Fishbach et al. (2018).

The likelihood can be written in terms of a non-zero-spin likelihood and a zero-spin likelihood:

$$\mathcal{L}(d|\Lambda) = (1 - \lambda_0)\mathcal{L}_{\chi>0}(d|\Lambda) + \lambda_0 \mathcal{L}_{\chi=0}(d|\Lambda) \tag{A2}$$

where

$$\mathcal{L}_{\chi>0}(d|\Lambda) = \int d\chi_{1,2} \int dz_{1,2} \int d\eta \, \mathcal{L}(d|\chi_{1,2}, z_{1,2}, \eta) \, \pi(\chi_{1,2}|\alpha_\chi, \beta_\chi) \, \pi(z_{1,2}|\zeta, \sigma_t, z_{\min}) \, \pi(\eta|\Lambda) \tag{A3}$$

$$\mathcal{L}_{\chi=0}(d|\Lambda) = \int d\chi_{1,2} \int dz_{1,2} \int d\eta \, \mathcal{L}(d|\chi_{1,2}, z_{1,2}, \eta) \, \delta(\chi_{1,2}) \, \pi(z_{1,2}|\zeta, \sigma_t, z_{\min}) \, \pi(\eta|\Lambda) \tag{A4}$$

$$= \int d\eta \, \mathcal{L}(d|\chi_{1,2} = 0, z_{1,2}, \eta) \, \pi(\eta|\Lambda). \tag{A5}$$

Using the standard trick for "recycling" posterior samples; see, e.g., (Thrane & Talbot 2019). In order to recycle, we rewrite the likelihood in terms of the posterior p the fiducial evidence $\pi(\theta|\varnothing)$, and the fiducial evidence $\mathcal{Z}_\varnothing$ so that

$$\mathcal{L}_{\chi>0}(d|\Lambda) = \frac{p_{\chi>0}(\Lambda|d)}{\pi(\theta|\varnothing)} \mathcal{Z}_\varnothing. \tag{A6}$$

Then we use the fact that

$$\int d\theta \, p(\theta) f(\theta) = \frac{1}{n_s} \sum_k f(\theta_k), \tag{A7}$$

to rewrite the integral as a sum over posterior samples. We rewrite $\mathcal{L}_{\chi>0}(d|\Lambda)$ as a sum over n_s fiducial ($\chi > 0$) posterior samples:

$$\mathcal{L}_{\chi>0}(d|\Lambda) = \frac{1}{n_s} \sum_k \left(\frac{\text{Beta}(\chi_{1,2}^k|\alpha_\chi, \beta_\chi)}{\pi(\chi_{1,2}^k|\varnothing)} \right) \left(\frac{\pi(z_{1,2}^k|\zeta, \sigma_t, z_{\min})}{\pi(z_{1,2}^k|\varnothing)} \right) \left(\frac{\pi(\eta^k, ...|\Lambda)}{\pi(\eta^k, ...|\varnothing)} \right) \mathcal{Z}_\varnothing. \tag{A8}$$

Here, $\mathcal{Z}_\varnothing$ is the fiducial evidence obtained using the fiducial ($\chi > 0$) prior. Likewise, the zero-spin likelihood can be written in terms of the n_0 zero-spin samples:

$$\mathcal{L}_{\chi=0}(d|\Lambda) = \sum_j \left(\frac{\pi(\eta^j, ...|\Lambda)}{\pi(\eta^j, ...|\varnothing)} \right) \mathcal{Z}_0. \tag{A9}$$

Here, $\mathcal{Z}_0$ is the zero-spin evidence obtained using the zero-spin ($\chi = 0$) prior. Putting everything together, we obtain

$$\mathcal{L}(d|\Lambda) = \frac{1 - \lambda_0}{n_s} \sum_k \left(\frac{\text{Beta}(\chi_{1,2}^k|\alpha_\chi, \beta_\chi)}{\pi(\chi_{1,2}^k|\varnothing)} \right) \left(\frac{\pi(z_{1,2}^k|\zeta, \sigma_t, z_{\min})}{\pi(z_{1,2}^k|\varnothing)} \right) \left(\frac{\pi(\eta^k, ...|\Lambda)}{\pi(\eta^k, ...|\varnothing)} \right) \mathcal{Z}_\varnothing + \frac{\lambda_0}{n_0} \sum_j \left(\frac{\pi(\eta^j, ...|\Lambda)}{\pi(\eta^j, ...|\varnothing)} \right) \mathcal{Z}_0 \tag{A10}$$