Draft version October 31, 2024
Typeset using LaTeX twocolumn style in AASTeX631

Hints of spin-magnitude correlations and a rapidly spinning subpopulation of binary black holes

Asad Hussain,[1] Maximiliano Isi,[2] and Aaron Zimmerman[1]

[1] Weinberg Institute, University of Texas at Austin, Austin, TX 78712, USA
[2] Center for Computational Astrophysics, Flatiron Institute, NY

(Dated: October 31, 2024)

ABSTRACT

The complex astrophysical processes leading to the formation of binary black holes and their eventual merger are imprinted on the spins of the individual black holes. We revisit the astrophysical distribution of those spins based on gravitational waves from the third gravitational wave transient catalog (GWTC-3, Abbott et al. 2023a), looking for structure in the two-dimensional space defined by the dimensionless spin magnitudes of the heavier (χ_1) and lighter (χ_2) component black holes. We find support for two distinct subpopulations with greater than 95% credibility. The dominant population is made up of black holes with small spins, preferring $\chi_1 \approx 0.2$ for the primary and $\chi_2 \approx 0$ for the secondary; we report signs of an anticorrelation between χ_1 and χ_2, as well as as evidence against a subpopulation of binaries in which both components are nonspinning. The subdominant population consists of systems in which both black holes have relatively high spins and contains $20^{+18}_{-18}\%$ of the binaries. The binaries that are most likely to belong in this subpopulation are massive and slightly more likely to have spin-orientations aligned with the orbital angular momentum—potentially consistent with isolated binary formation channels capable of producing large spins, like chemically homogeneous evolution. This hint of a rapidly spinning subpopulation hinges on GW190517, a binary with large and well-measured spins. Our results, which are enabled by novel hierarchical inference methods, represent a first step towards more descriptive population models for black hole spins, and will be strengthened or refuted by the large number of gravitational wave detections expected in the next several years.

1. INTRODUCTION

Gravitational wave (GW) detections by the LIGO-Virgo-KAGRA (LVK) Collaboration (Aasi et al. 2015; Acernese et al. 2015; Akutsu et al. 2021) have opened a unique window onto compact objects like black holes (BHs) and neutron stars, as well as the massive stars that produce them. In particular, the vast majority of GW detections are of binary black holes (BBHs) (Abbott et al. 2019a, 2021a, 2024, 2023a; Nitz et al. 2019, 2020, 2021, 2023; Zackay et al. 2019; Venumadhav et al. 2019, 2020; Zackay et al. 2021; Olsen et al. 2022; Mehta et al. 2023), which are otherwise invisible.

The distribution of spins of the individual BHs in these binaries may hold clues about their origin, e.g., whether they evolve from an isolated stellar binary or they are dynamically formed in dense environments (see, e.g., reviews by Mapelli 2020; Mandel & Farmer 2022). The dimensionless spin magnitudes, in particular, may reveal how angular momentum is distributed in the stellar progenitors and captured by the BHs at birth, as well as carry imprints of binary interactions after the first BH forms (Belczynski et al. 2020; Qin et al. 2018; Fuller et al. 2019; Fuller & Ma 2019; Ma & Fuller 2019; Bavera et al. 2020; Bavera et al. 2021; Steinle & Kesden 2021; Zevin & Bavera 2022). Spin magnitudes may additionally identify hierarchical BBHs, whose component BHs are themselves the product of previous mergers (Gerosa & Berti 2017; Rodriguez et al. 2019; Kimball et al. 2020, 2021; Doctor et al. 2020; Gerosa & Fishbach 2021; McKernan & Ford 2024; Payne et al. 2024).

Past studies of LVK data have explored the distribution of BH spins under different, more or less restrictive, assumptions. Since measuring individual component spins can be difficult (van der Sluys et al. 2008; Raymond et al. 2010; Cho et al. 2013; O'Shaughnessy et al. 2014; Vitale et al. 2014; Ghosh et al. 2016; Chatziioannou et al. 2018; Pratten et al. 2020; Green et al. 2021; Biscoveanu et al. 2021b,a; Varma et al. 2022; Miller et al. 2024a,b), many analyses have looked at derived quantities like the effective spin $\chi_{\rm eff}$, which is a mass-weighted average of the spin components along the orbital angular momentum (Damour 2001; Ajith et al. 2011), finding that this quantity must be small but likely positive in most systems (e.g., Abbott et al. 2021b, 2023b,b; Miller et al. 2020; Callister et al. 2021b; Roulet et al. 2021; Adamcewicz & Thrane 2022; Biscoveanu et al. 2022; Franciolini & Pani 2022; García-Bellido et al. 2021).

Other works have directly tackled the individual spin magnitudes χ_i of the heavier ($i = 1$) and lighter ($i = 2$) components of the binary, typically assuming that they are independently and identically drawn from a unimodal distribution (Wysocki et al. 2018; Abbott et al. 2019b, 2021b, 2023b); those measurements constrain spin magnitudes to be small, $\chi_i \approx 0.2$, but with wide uncertainties. Motivated by predictions like Fuller & Ma (2019), such models have been enhanced to look for a subpopulation of nonspinning BBHs: while earlier studies found evidence of two populations, one with negligibly small spins and the other with larger spins (Galaudage et al. 2021; Roulet et al. 2021; Hoy et al. 2022; Kimball et al. 2021), reanalyses with more events show no clear evidence for or against it (Tong et al. 2022; Callister et al. 2022). Finally, a few studies have modeled the spins of the primary and secondary objects as drawn from distinct, independent distributions (Tong et al. 2022; Mould et al. 2022; Adamcewicz et al. 2024; Golomb & Talbot 2023; Edelman et al. 2023); these measurements agree that the component spins have typical values $\sim$0.2 with a wide spread, and Mould et al. (2022) finds hints that the secondary could tend to have lower spins.

In this paper, we take another look at the population of BH spin magnitudes, this time studying the structure in the joint distribution of the component spins. Our main motivation is to look for features in the two-dimensional χ_1–χ_2 plane that may have escaped previous analyses because of their assumption of independent components: information about χ_1–χ_2 correlations is destroyed, and evidence of subdominant populations may be washed away, when the spins are treated as independent. Additionally, we implement a novel technical framework that allows us to model arbitrarily narrow features in the population and treat boundary effects in the spin magnitude domain without bias. This allows us to revisit the existence of a subpopulation of nonspinning BHs while overcoming some of the technical hurdles that have challenged previous studies.

In what follows, we describe our population model and dataset in Sec. 2, our population inferences in Sec. 3, and the astrophysical implications of our results in Sec. 4. We discuss our conclusions and future prospects in Sec. 5. Additional details on our methods are given in Appendix A and further results are given in Appendix B. More details about methodology for BBH population inference, related methods, and additional applications are described in a companion paper Hussain et al. (2024).

2. METHODS AND POPULATION MODELS

We use hierarchical Bayesian inference to infer the population properties of BBHs (e.g., Loredo 2004; Mandel et al. 2019; Thrane & Talbot 2019; Vitale et al. 2022). The goal is to compute posteriors over the hyperparameters $\mathbf{\Lambda}$ of our chosen population model. In this study we adopt a flexible, two-population model for the BBH spin magnitudes, $\boldsymbol{\chi} = (\chi_1, \chi_2)$, drawing them from a mixture of two correlated and truncated 2D Gaussians (indexed by a and b) with a mixing fraction η,

$$p(\boldsymbol{\chi}) = \eta N_{[\mathbf{0},\mathbf{1}]}\left(\boldsymbol{\chi} \mid \boldsymbol{\mu}^a, \mathbf{\Sigma}^a\right) + (1-\eta) N_{[\mathbf{0},\mathbf{1}]}\left(\boldsymbol{\chi} \mid \boldsymbol{\mu}^b, \mathbf{\Sigma}^b\right), \tag{1}$$

where the $[\mathbf{0},\mathbf{1}]$ subscript indicates truncation of our domain to the $[0,1]\times[0,1]$ unit square, while both $\mathbf{\Sigma}^a$ and $\mathbf{\Sigma}^b$ independently have the general form

$$\mathbf{\Sigma} = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}. \tag{2}$$

Since we have two identical Gaussians, we face a label-switching degeneracy (e.g. Buscicchio et al. 2019), which we break by assigning an identity to the dominant population (a), requiring $\eta \in [0.5, 1]$. We use truncated Gaussians rather than the Beta distributions used in Abbott et al. (2023b) to more easily generalize to two dimensions and to better represent the edges of our truncated domain without systematics. Aside from the use of truncated Gaussians, Eq. (1) encompasses a wide variety of spin magnitude models used in previous studies.

For the remaining BBH parameters, we use the fiducial mass and redshift models defined in Abbott et al. (2023b), and the fiducial model for the tilt angles of the spins with respect to the orbital angular momentum, which assumes both spins are drawn from a two-population model: one uniform in tilt angle (isotropic spins) and one drawn from a half-normal peaking at aligned spins. Together with Eq. (1), this makes up our full population likelihood $\mathcal{L}(\{d_i\} \mid \mathbf{\Lambda})$, where d_i represents the data for the ith detection; with a population prior $p(\mathbf{\Lambda})$ and an estimate of the detection efficiency $\xi(\mathbf{\Lambda})$, we can then compute the population posterior $p(\mathbf{\Lambda} \mid \{d_i\})$ while accounting for selection effects (Mandel et al. 2019; Loredo 2004).

Population inference requires estimating high-dimensional integrals, which can be challenging for standard Monte Carlo methods (e.g., Farr 2019) when population features are narrow or concentrated at the edges of the domain. This is the case when looking for a subpopulation of BBHs with negligible spins

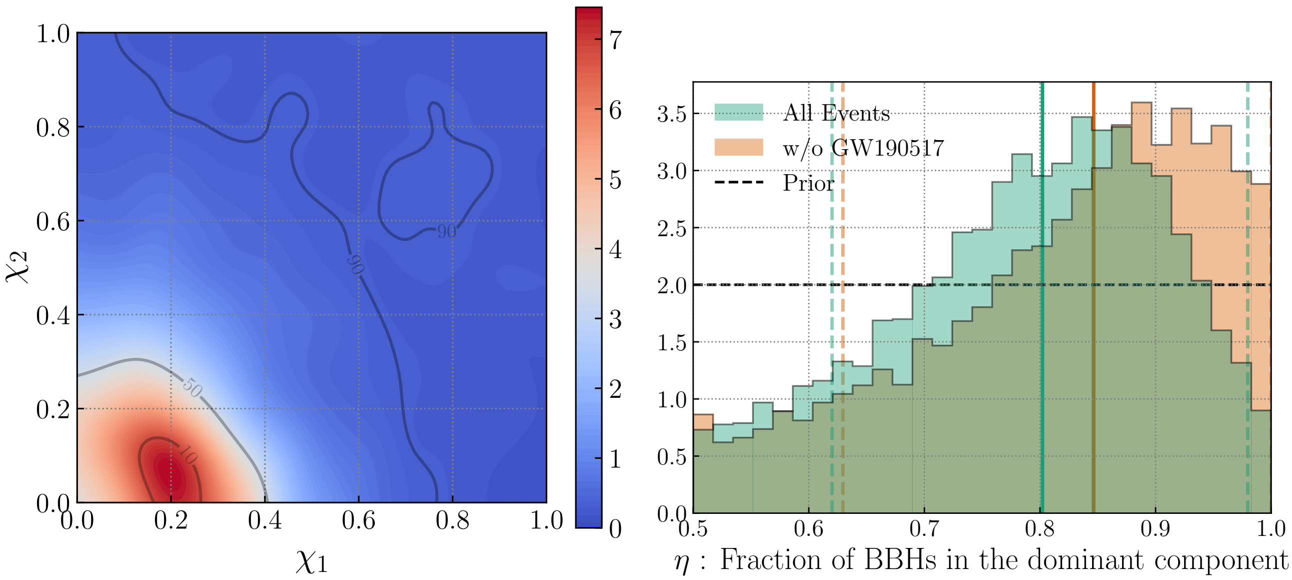

Figure 1. *Left:* PPD of the spin magnitudes in our two-component model. A hint of the subdominant component is visible at at high χ_1, χ_2. This component contains $20^{+18}_{-18}\%$ of the BBHs and is diffuse, hence is not very apparent in the PPD despite its statistical significance. The dominant, slowly spinning component shows hints of the anticorrelation between the primary and secondary spin magnitudes. *Right:* Marginalized posterior over the fraction of BBHs in the dominant, slowly spinning component. We also show the median and 90% HPDI. When all events are included the data prefers the existence of the highly spinning component ($\eta \to 1$ is disfavored). However, this subpopulation is disfavored when GW190517 is removed.

($\chi_i \to 0$). To accurately compute these integrals, we first represent both single-event posteriors and detection-probability estimates as truncated Gaussian mixture models (TGMMs). This allows us to leverage properties of Gaussians to analytically evaluate integrals over the spin sector (magnitudes and tilts), while using Monte Carlo averages over the remaining parameters. Our methods help to control the variance in the estimates of our likelihood integrals across hyper-parameter space, and so we do not apply data-dependent priors to exclude regions of high variance (unlike, e.g., Abbott et al. 2023b). We discuss our strategy in detail in a companion paper (Hussain et al. 2024), and summarize it in App. A.

We visualize the result of our fits by plotting the PPD for the spin magnitudes, which represents the inferred distribution of spins marginalized over all our population parameters, i.e.,

$$p(\boldsymbol{\chi} \mid \{d_i\}) = \int p(\boldsymbol{\chi} \mid \boldsymbol{\Lambda})\, p(\boldsymbol{\Lambda} \mid \{d_i\})\, \mathrm{d}\boldsymbol{\Lambda}\,. \tag{3}$$

We also show projections of the population posterior for different hyper-parameters. Additionally, to perform model comparison of our fiducial spin-magnitude model against lower-dimensional subcases, we use the Savage Dickey density ratio (SDDR) for nested models (Dickey 1971). To compute this, we use tailored methods to construct unbiased truncated kernel density estimates (KDEs) of our population posteriors to evaluate them at the limiting points (Hussain et al. 2024), and bootstrap over multiple hyper-posterior draws to report a median estimate and a 90%-confidence highest-density interval (see also Appendix A). Unbiased density estimation on boundaries is a well-known challenge in many settings, and our methods have benefits over standard solutions like reflective KDEs, for example not imposing a zero derivative at the boundary (see also Appendix A).

For our dataset we use the 69 confidently detected (false alarm rates below 1/yr) events used by Abbott et al. (2023b) for BBH population inferences. We use posterior samples produced using the IMRPhenomX-PHM waveform model (Pratten et al. 2021) released a part of the GWTC-2.1 and GWTC-3 catalogs (Abbott et al. 2024, 2023a), and available as open data (Abbott et al. 2021c, 2023c) at LIGO Scientific, Virgo, and KAGRA Collaborations (2024). To incorporate selection effects, we use the sensitivity estimates described in Abbott et al. (2023b) and provided by LIGO Scientific, Virgo and KAGRA Collaborations (2021). We sample our hyper-posteriors using the no-U-turn sampler (Hoffman & Gelman 2011) Hamiltonian Monte Carlo (Neal 2011; Betancourt 2017) implemented in numpyro (Phan et al. 2019; Bingham et al. 2019).

3. RESULTS

We present our main result in Fig. 1, showing the PPD over χ_1 and χ_2 (left) and the posterior on the mixture fraction η (right). The PPD reveals a bimodal distribution of spins, with a dominant component peaking at $\chi_1 \approx 0.2$ and $\chi_2 \approx 0$ and a subdominant component peaking at $\chi_i \approx 0.75$. In Fig. 2, we isolate the contributions of each component, making it clear that the dominant mode consists of low spins BHs (top), while the subdominant mode mostly supports high spins (bottom). While the dominant mode is quite well measured, the subdominant mode is localized more diffusely, as might be expected given its lower occupancy. The subdominant component makes a relatively small contribution to the PPD in Fig. 1, but this is a function of the smaller number of events that are assigned to this mode and not a measure of our certainty in its existence.

To asses the significance of the second mode, the right panel of Fig. 1 shows our inferred posterior over the fraction of systems in the dominant subpopulation. The fraction of BBHs in the subdominant component is $1-\eta = 0.2^{+0.18}_{-0.18}$ quoting the HPDI around the median. In other words, $\sim$20% ($\sim$80%) of the BBHs are favored to be in the rapidly (slowly) spinning subpopulation. The posterior on η rules out a single population of spin-magnitudes ($\eta = 1$) at better than 95% credibility. The Bayes factor (BF) in favor of the two-component model is $\mathcal{B}(2 \text{ vs } 1) = 2.1^{+0.3}_{-0.3}$ (see Sec. 2 and Appendix A.2).

The evidence for the subdominant mode is highly sensitive to the rapidly spinning event GW190517 (Abbott et al. 2021a). Removing it from our set, the data no longer support the a subpopulation, with $1-\eta = 0.15^{+0.22}_{-0.15}$ encompassing zero within 90% credibility, and instead yielding a BF against the existence of this subpopulation of $\mathcal{B}(1 \text{ vs } 2) \approx 1.4^{+0.1}_{-0.1}$. The source of GW190517 is decisive due to the fact that its spin magnitudes are confidently measured to be large (see, e.g., Fig. 10 of Abbott et al. 2021a or Fig. 2 of Qin et al. 2022), and cannot be accommodated by the dominant population alone. On the other hand, our results are insensitive to the exclusion of other GW events from highly spinning BBHs, such as GW191109 (Abbott et al. 2023a). No data quality issues have been reported for GW190517.

Although previous works have reported hints of a potential subpopulation of rapidly spinning BHs (Galaudage et al. 2021; Roulet et al. 2021; Hoy et al. 2022), it would be difficult for such studies to clearly identify the secondary mode in Fig. 1 because such a subpopulation is not apparent in the χ_1 or χ_2 marginals (see Fig. 8 in Appendix B), as it is obfuscated by the tails of the low-spin mode. On the other hand, the subpopulation stands out more clearly in 2D because those high-spin systems would otherwise have to be accommodated by the tail of the dominant component in both χ_1 and χ_2 *simultaneously*, which is made difficult by the fact that the majority of events constrain the bulk of the posterior to be in a compact ball near the origin (in other words, the secondary mode lies well beyond the 90% credibility contour in 2D, but not in the 1D marginals).

When comparing to studies that allow for a subpopulation with negligible spins, we must look at two cases, one where each of our Gaussian components separately concentrates at the origin. Using the SDDR, we find a BF against a dominant population with negligible spin of $\mathcal{B} = 7^{+71}_{-4}$ and a BF against a subdominant population with negligible spins $\mathcal{B} = 6^{+4}_{-3}$. As compared to our other SDDR-based BF comparisons, these are especially uncertain because we must extrapolate our hyper-posterior samples to a corner in 4D space ($\mu^a_1 = \sigma^a_1 = \mu^a_2 = \sigma^a_2 = 0$), but in both cases we clearly disfavor a subpopulation with negligible spins.

We next describe the features of each subpopulation in more detail. Additional corner-plots of the hyper-posteriors are shown in Appendix B.

3.1. *The dominant population: slow and anticorrelated spins*

The dominant mode in the population has a number of interesting features: (1) both component spins are well constrained to be low, (2) there are differences between the primary and secondary spins, and (3) there is a hint of anticorrelation between the spins. We discuss each in turn.

The fact that BHs in this subpopulation tend to spin slowly, regardless of whether they are the lighter or heavier component in the binary, is evident from the PPD in the top panel of Fig. 2; since this mode dominates the population, it can also be gleaned from Fig. 1. Additionally, the preference for low spins can be seen in the inferred population mean and scale parameters for this mode (Fig. 3), which strongly favor low values for both components in the binary. The uncertainties on χ_2 are not significantly worse than χ_1. Since the dominant mode contains most BBHs, it has properties similar to those inferred in past spin population studies.

Next, we find interesting differences between χ_1 and χ_2. While the spin of the primary BH peaks at $\chi_1 \approx 0.2$, as expected from previous studies, the secondary BH population is consistent with identically vanishing spins. Not only does the PPD peak at $\chi_2 \approx 0$ but also the population posterior favors a delta function at $\chi_2 = 0$ (the purple distribution peaks at $\mu^a_2 = \sigma^a_2 = 0$ in Fig. 3), with a BF of $\mathcal{B} = 15^{+2}_{-2}$ in favor of identically

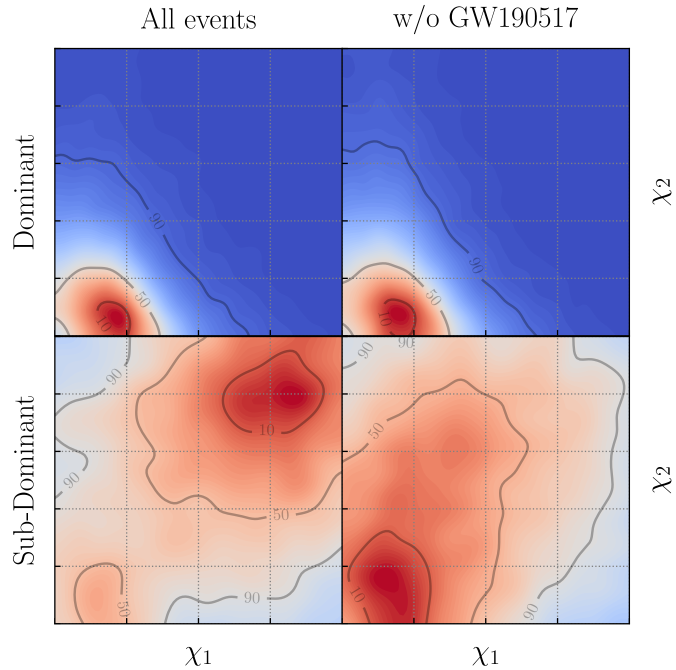

Figure 2. Separate PPDs of the spin magnitudes of the dominant and subdominant populations, shown with and without the inclusion of GW190517. One can see that the subdominant component prefers higher spins when GW190517 is included. Without GW190517 the peak moves to lower spin values and is degenerate with the dominant component, indicating that the data disfavors a subpopulation. Note that the color scale normalization varies from panel to panel to resolve the variations in the PPD.

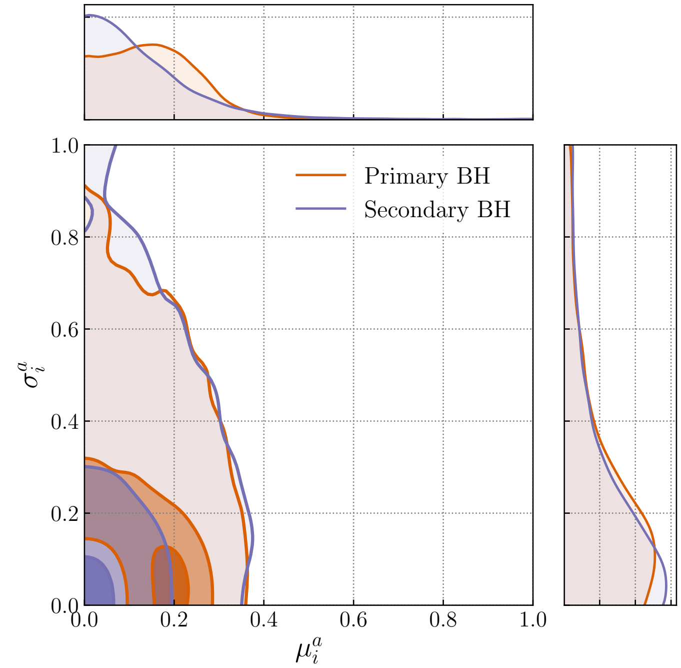

Figure 3. Dominant subpopulation. Posteriors for the mean (μ_i^a) and scale (σ_i^a) population parameters of the primary BH (orange) and secondary BH (purple) in the dominant subpopulation with 90%, 50% and 10% credible intervals marked. The data prefers that for the secondaries $\mu_2^a \to 0$, $\sigma_2^a \to 0$ (i.e. all secondary BHs nonspinning). In addition there is a degeneracy between μ_1^a and σ_1^a for the primary's spin. The data supports a sharp peak with all primaries having $\chi_1 \approx 0.2$, and is also consistent with χ_1 drawn from a half-normal peaking at $\chi_1 = 0$ with a spread of ≈ 0.2

nonspinning secondaries. Conditioned on $\eta \to 1$, (i.e., only one population exists) we still find that the data favors all of the secondaries to be nonspinning with a BF of $\mathcal{B} = 10^{+6}_{-3}$.

Previous studies have looked for differences in the χ_1 and χ_2 distributions, assuming independence. For example, Adamcewicz et al. (2024) compared a model where both BHs are spinning to one where either the primary, secondary, or both are nonspinning by repeating parameter estimation over the catalog of BBHs for each case; we corroborate their result favoring the case where only the primaries are spinning, without the need for additional, computationally expensive parameter estimation. Hints of the support for lower secondary spins can also be gleaned from Fig. A1 in Mould et al. (2022); on the other hand, Edelman et al. (2023) found no visible difference between χ_1 and χ_2 with a more flexible model.

Unlike the secondaries, the primaries cannot all be nonspinning. This is clear from Fig. 3, where μ_1^a and σ_1^a are not allowed to simultaneously vanish (orange contours). The distribution of χ_1 most likely peaks at $\mu_1^a \approx 0.2$, as expected from Fig. 2, but may peak at zero *as long as the spread is sufficiently large* ($\sigma_1^a \approx 0.2$). This means that the data show evidence of at least one event with a primary spin of $\chi_1 \approx 0.2$.

Finally, returning to the top panels of Fig. 2, we see that the PPD suggests an anticorrelation between the spin magnitudes. While we do not rule out zero correlation ($\rho^a = 0$) at the 90% credible level, the BF against an uncorrelated distribution is $\mathcal{B} = 2.9^{+0.1}_{-0.1}$. Assessment of the correlation is complicated by the fact that a uniform prior on ρ^a induces a prior on the Pearson correlation coefficient $\hat{\rho} = \mathrm{Corr}[\chi_1, \chi_2]$ which is strongly peaked at zero. This is because for truncated Gaussians, a large range of means, scale parameters, and ρ values result in a small empirical correlation in the bounded χ_1–χ_2 domain. This can be seen clearly in Fig. 4, where we show our posteriors and priors over $\hat{\rho}$ for both subpopulations. We see that there are hints of an anticorrelation between the spins of the dominant population, as it is able to overcome this strong prior on $\hat{\rho}^a$. We also plot the case where we fix $\eta = 1$, so that we have only a sin-

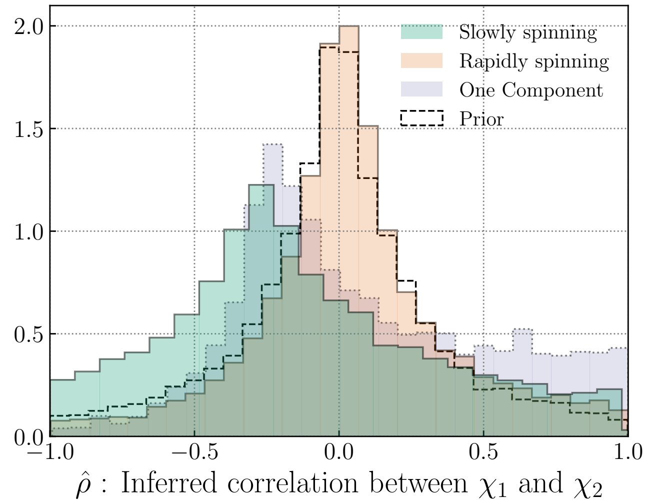

Figure 4. Posterior over the pearson correlation coefficient $\hat{\rho}$ between the spin magnitude of the primary and secondary black hole for BBHs in the different components.

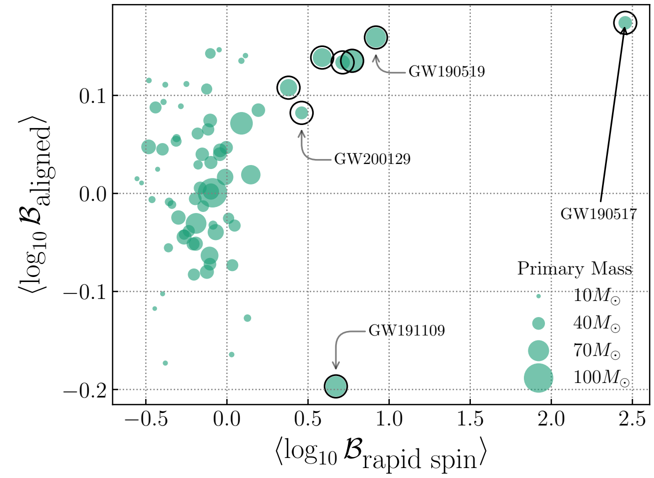

Figure 5. Log BFs for all 69 events, comparing the hypotheses that each BBH comes from the rapidly spinning subpopulation versus the dominant slowly spinning population (abscissa) and that each comes from the predominantly aligned-spin subpopulation versus the isotropic subpopulation (ordinate). These BFs are marginalized over our hyperposterior samples in the manner discussed in Appendix A.3. The trend implies that events more likely to lie in the rapidly spinning subpopulation are more likely to come from the predominantly aligned subpopulation and tend to have higher masses. Events with a median BF greater than 2 are marked with black circles, and some events of special interest are labeled.

gle spin-magnitude population. The peak is then similar to our fiducial model, favoring weak anticorrelation, but with a heavy tail towards positive $\hat{\rho}$ values. This can be interpreted as the imprint of the highly spinning subpopulation on our single-population model, since some posterior weight is pulled towards the large χ_1–χ_2 region while requiring the bulk of the spin population to lie at small spins.

In any case, the correlation structure apparent in the PPD of Fig. 2 suggests a preference for pairing higher spinning primaries with lower spinning secondaries and vice versa, such that systems where both BHs are nonspinning ($\chi_1 = \chi_2 = 0$) are measurably disfavored.

3.2. *The subdominant population: relatively high spins*

The PPD of the subdominant population (bottom left of Fig. 2) peaks at large spins, and slightly disfavors cases where one of the two BHs is rapidly spinning while the other has negligible spin. However, the subpopulation is broad, and the posteriors on the hyperparameters are weakly informed by the data. Its empirical correlation $\hat{\rho}^b$ is consistent with the prior, as seen in Fig. 4. Further, we see some contamination at small spin-magnitudes in the PPD from the remaining degeneracy between the dominant and subdominant components, which occurs when $\eta \to 0.5$. The bottom right plot in Fig. 2 shows that once we remove the highly spinning event GW190517, this subpopulation becomes degenerate with the dominant component in terms of its location but also captures the tail of the distribution of spin magnitudes towards higher values of χ_i.

Additional clues about the origin of this subdominant mode can be found in possible correlations with other BBH parameters, for example the masses or spin tilts. In Fig. 5, we visualize the estimated log BF for each BBH we analyzed to lie in the rapidly spinning versus the slowly spinning population (abscissa), as well as the log BF that to lie in the mostly-aligned-spin versus isotropic-spin tilt population (ordinate). Although the spin-tilt BFs are very weak, it seems that the rapidly-spinning BBHs tend to also fall in the aligned-spin population, with the exception of GW191109 which is considerably anti-aligned (Udall et al. 2024); for BHs with smaller spins it is harder to determine the tilts, leading to correspondingly larger scatter in the spin-tilt BF. In addition to tilts, Fig. 5 also encodes the binary primary mass (maker size) revealing that those systems most likely to lie in the rapidly spinning population are also more massive, consistent with a correlation between mass and spin previously identified in, e.g., Tiwari & Fairhurst (2021); Hoy et al. (2022); Franciolini & Pani (2022); Callister et al. (2021b); Adamcewicz & Thrane (2022); Biscoveanu et al. (2022).

4. ASTROPHYSICAL IMPLICATIONS

The spin magnitudes of BBHs are determined by a number of factors, ranging from binary interactions which may tidally spin up the progenitor stars, the angular momentum distribution and transport within the progenitor stellar cores during their evolution, the uncertain details of core collapse, and the effects of any fallback accretion (e.g., Mandel & Farmer 2022, and references therein). At face value, the bimodality that we infer in BBH spin magnitudes suggests that multiple formation channels may be at play (already supported by the inferred distribution of spin-tilt angles, Abbott et al. 2023b). This could be the case if LVK BBHs are a mixture of dynamically-formed systems and systems from isolated stellar binaries. Alternatively, isolated binaries alone could accommodate this bimodality if, e.g., angular momentum is transported from the stellar core to the envelope leading to slowly spinning BHs (Fuller & Ma 2019), except when stars are tidally spun up (e.g. Kushnir et al. 2016; Fuller & Lu 2022; Ma & Fuller 2023) or are chemically homogeneous without a core-envelope structure (e.g. Mandel & de Mink 2016; de Mink & Mandel 2016; Marchant et al. 2016), resulting in one or both BHs with large spins (e.g., Qin et al. 2019, although see Riley et al. 2021).

With this context, we first consider interpretations for the dominant population of slowly spinning BBHs, where we find the secondary BHs have negligible spin and the primaries have spins concentrated near $\chi_1 \approx 0.2$, with hints of an anticorrelation between the component spin magnitudes, and strong support against both BHs being nonspinning. Given the wide range of angular momentum values that stellar cores can possess before collapse (e.g., Qin et al. 2019), it is challenging from first principles to produce natal BH spins that are not either very high or negligible, and doubly so to have two distinct cases for the primary and secondary. As seen in Fig. 3, another acceptable scenario is one where the secondary spin is negligible and the primary distribution peaks at zero with a width $\sigma_1^a \approx 0.25$, allowing for a range of spin magnitudes for the primary. Both cases remain challenging to explain, since a standard binary evolution scenario might leave the first born, presumably more massive, BH with a small spin while the secondary can be spun up by tidal effects (Hotokezaka & Piran 2017; Zaldarriaga et al. 2018; Qin et al. 2018). Alternatively, BH spins of $\chi_i \approx 0.2$ can be explained by moderately efficient angular momentum transport mechanisms that result in such natal spins (Belczynski et al. 2020), or through the accretion of a portion of the highly convective envelope onto a BH with negligible natal spin (Antoni & Quataert 2021). However, it is not clear how to explain the distinction between primary and secondary spins in these scenarios.

One resolution to the χ_1 versus χ_2 asymmetry would be for a mass-ratio reversal to occur prior to the formation of the first BH, followed by spin-up of the now more massive star through tidal interactions with the first BH (e.g., Gerosa et al. 2013; Olejak & Belczynski 2021; Zevin & Bavera 2022; Broekgaarden et al. 2022), yielding a spinning primary and a secondary with a lower spin. The models considered by Broekgaarden et al. (2022) indicate that mass-ratio reversal can be common among detectable BBHs, but also predict some systems with spinning secondaries and a significant number of systems with negligible spins. However, the spin-tilt distribution provides evidence that a fraction of BBHs have isotropically distributed spins (Abbott et al. 2023b). Since we find that the rapidly spinning subpopulation prefers more aligned spins, we expect that some of the low-spin systems arise from the isotropic spin distribution, indicative that some BBHs form outside of the isolated binary evolution channel, or strong natal kicks (e.g., Callister et al. 2021a).

The highly spinning subpopulation is intriguing. A key feature is that both primary and secondary spins tend to be large in this population, although with a wide range of uncertainties. We disfavor the case where the primary has a large spin $\chi_1 \approx 0.7$ while the secondary has a small spin, disfavoring a population of mergers between first and second-generation BHs in a dense stellar environment; since such mergers would be more common than mergers between two second-generation BHs (e.g., Kimball et al. 2020), it is unlikely that these large BH spins are produced by previous BBH mergers. The correlation we find between probability of being highly-spinning and probability of having relatively aligned spins for the detected BBHs further hints that the large spins may arise from binary interactions. Together with the fact that the highly spinning population appears to be made up of more massive BHs, the homogeneous evolution of low-metallicity binaries would appear to be a reasonable scenario for this population.

5. CONCLUSIONS

In this work we have investigated the astrophysical distribution of BH spin magnitudes based on 69 BBHs detected with high significance in GW observations by the LVK. Unlike previous work, we have explored the two-dimensional space of component-spin magnitudes directly, using a model that subsumes those of many previous studies and allows for two distinct populations in spin-magnitude space. We have confirmed that the bulk of BBHs have small but non-negligible spin magnitudes, with primaries favoring $\chi_1 \approx 0.2$, while finding new evidence that secondary BHs are consistent with having identically zero spin. We have also found evidence for a weak anticorrelation between the spin components and a we disfavor models in which the majority of BBHs have negligible spins. The latter result was enabled in this study by novel methods to make inferences about populations with narrow features without recourse to tailored parameter estimation for the GW events.

We have uncovered hints of a second, subdominant spin population containing $20^{+18}_{-18}\%$ of the BBHs, ruling out a single component with better than 95% credibility. This second population, although broad, peaks at large spin magnitudes for both the primary and secondary BHs. The evidence for this subpopulation is largely driven by a single GW event, GW190517, whose components have high and relatively well-measured spin magnitudes. We identify the probable GW events which arise from the rapidly spinning subpopulation, and find that these are preferentially massive. As in previous population studies, (e.g. Abbott et al. 2023b), we allow for two populations of spin orientations, one isotropic and one peaking towards alignment with the orbital angular momentum. We find that the BBHs identified with the rapidly spinning population are somewhat more likely to be in the aligned spin population, perhaps suggesting an origin in field binaries composed of massive, rapidly rotating stars. Note that spin measurements can be impacted by GW modeling systematics, and future work should assess the impact of such systematic errors on the properties and significance of the subdominant population found here.

The analysis reported here is only the first allowing for multiple populations of BBHs with correlated spin-magnitudes. In order to better constrain the possible origin of the rapidly spinning BBHs, analyses including correlations between the spin magnitudes and spin orientations, binary masses, or redshift would be of great interest. Targeted analysis exploring the dominant, slowly spinning population would also be highly valuable, particularly those that address the possibility that mass ratio reversal may play a role in forming binaries with small but non-negligible primary spin magnitudes and secondary spins consistent with zero.

The evidence for the rapidly spinning subpopulation and its properties remains tentative. With only 69 events in our dataset, our ability to infer fine details and isolate subpopulations is limited. At the time of writing, however, the LVK is in the midst of its fourth observing run, at even greater sensitivity than the previous campaign. To date over 100 public alerts have been issued reporting GW event candidates with false alarm rates less than 2/yr, with many more BBH detections expected as GW detectors reach and exceed their design sensitivities in the coming years (Abbott et al. 2016). This growing dataset should confirm the rapidly spinning subpopulation if it exists, allow for more detailed inferences about the properties of BBHs in it, and, in so doing, help to uncover the origin and formation channels of merging BBHs.

We thank Tom Callister, Will Farr, Maya Fishbach, Simona Miller, Matthew Mould, and Colm Talbot for valuable discussions. A.H and A.Z. were supported by NSF Grants PHY-1912578, PHY-2207594 and PHY-2308833 over the course of this work. The Flatiron Institute is a division of the Simons Foundation. This work has preprint numbers LIGO-P2400522 and UTWI-33-2024.

This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by NSF Grants PHY-0757058 and PHY-0823459. This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation, as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan.

APPENDIX

Model / Sub-model	Schematic	Prior/Limits	Comments/Bayes Factor
Two Components $p(\boldsymbol{\chi}) = \eta N_{[\mathbf{0},\mathbf{1}]}\left(\boldsymbol{\chi} \mid \boldsymbol{\mu}^a, \boldsymbol{\Sigma}^a\right) +$ $(1-\eta) N_{[\mathbf{0},\mathbf{1}]}\left(\boldsymbol{\chi} \mid \boldsymbol{\mu}^b, \boldsymbol{\Sigma}^b\right)$	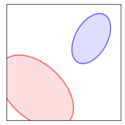	$\eta \sim U(0.5, 1)$ $\mu_i^{a,b} \sim U(0,1)$ $\sigma_i^{a,b} \sim U(0,1)$ $\rho^{a,b} \sim U(-1,1)$	Fraction of BBHs in the dominant component Mean parameters of each component Scale parameters of each component Correlation parameter of each component
Two uncorrelated components	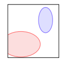	$\hat{\rho}^{a,b} = 0$	$\mathcal{B} \approx 2.5^{+0.2}_{-0.2}$ **against** this model.
Slow-spin subpopulation *(dominant component, isotropic, peaked at zero)*	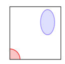	$\mu_1^a = \mu_2^a = 0,$ $\sigma_1^a = \sigma_2^a = \sigma_0,$ $\rho^{a,b} = 0$	$\mathcal{B} \approx 29^{+12}_{-12}$ **in favor** of this model compared to two uncorrelated components (above). Similar to the `BetaSpike-Galaudage` analysis of Callister et al. (2022), except there $\sigma_0 \leq 0.1$ which is a disfavored regime, and their bulk spins are identically distributed.
Zero spin subpopulation	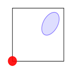	$\mu_1^a = \mu_2^a = 0,$ $\sigma_1^a = \sigma_2^a = 0$	$\mathcal{B} \approx 7^{+71}_{-4}$ **against** this model if dominant, and $\mathcal{B} \approx 6^{+4}_{-3}$ **against** this model if subdominant (swap components $a \leftrightarrow b$), compared to the full two component model. Similar to the Tong et al. (2022) `NONIDENTICAL` model and Mould et al. (2022) `NONIDENTICAL + ZEROS` model, except here the bulk has correlations.
One Component $p(\boldsymbol{\chi}) = N_{[\mathbf{0},\mathbf{1}]}\left(\boldsymbol{\chi} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}\right)$ $\mathcal{B} \approx 2.1^{+0.3}_{-0.3}$ **against** this model compared to the two component model.	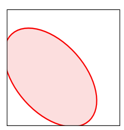	$\mu_i \sim U(0,1)$ $\sigma_i \sim U(0,1)$ $\rho \sim U(-1,1)$	Mean parameter of the truncated normal Scale parameters of the truncated normal Correlation parameter of the truncated normal
Primary and secondary uncorrelated	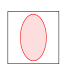	$\hat{\rho} = 0$	$\mathcal{B} \approx 2.2^{+0.1}_{-0.1}$ **against** this model when compared to the one component model. Similar to the Mould et al. (2022) `Nonidentical` model.
Primary and secondary are IID	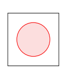	$\hat{\rho}$, $\mu_1 = \mu_2$ $\sigma_1 = \sigma_2$	$\mathcal{B} \approx 1.3^{+0.2}_{-0.3}$ **in favor** this model when compared to the one component model. This model is similar to the IID models often used, for example by Abbott et al. (2023b) in their fiducial model.
All secondaries not spinning	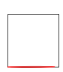	$\mu_2 = \sigma_2 = 0$	$\mathcal{B} \approx 10^{+6}_{-3}$ **in favor** of all secondaries nonspinning over the two component model. ($\mathcal{B} \approx 13^{+3}_{-2}$ if compared to the one component model)

Table 1. Summary of two-component and one-component models, their priors, and limiting cases. We include the BFs computed using the SDDR for each sub-model, compared to either the full two-component model with correlations or the one-component model (for the second section). Additionally, we note if some cases resemble those explored in the literature, modulo interchange of a beta distribution with a truncated normal.

A. DETAILS ON METHODS AND POPULATION MODELS

To allow us to probe the population properties of astrophysical BBHs, we start by using the fiducial mass, redshift and spin orientation models as defined in Abbott et al. (2023b). More specifically we use the `PowerLawPlusPeak` model for the masses, `PowerLaw` model for redshift and we use only the part of the `Default` spin model that pertains to the spin orientation given by,

$$p\left(\boldsymbol{z}_t \mid \zeta, \sigma_t\right) = \zeta N_{[\mathbf{0},\mathbf{1}]}\left(\boldsymbol{z}_t \mid \sigma_t \boldsymbol{I}_2\right) + (1-\zeta)\frac{1}{4}\,, \tag{A1}$$

where $z_{t,i} = \cos\theta_i$ and θ_i is the tilt angle between each spin and the orbital angular momentum of the binary. In other words, the cosine tilts are either drawn independently from a half-normal with scale parameter σ_t (aligned) or from a uniform distribution (isotropic), with the fraction of aligned spin binaries given by ζ. The spin magnitude model is given in Eq. (1) and the priors on the population parameters for it are given in Table 1. The priors for the rest of the population parameters are set as the same as those used in Abbott et al. (2023b), with the exception of the width of the isotropic spin distribution, where we allow the prior to go all the way to zero [$\sigma_t \sim U(0, 4)$, as opposed to $\sigma_t \sim U(0.1, 4)$].

We use truncated normal distributions to model the spin magnitudes as opposed to beta distributions, since the truncated normal distribution is better at recovering sharp distributions near the edges (e.g. a sharp peak at $\chi \approx 0$) than a beta distribution. In addition, a beta distribution may become singular for certain parameter values ($\alpha, \beta > 1$), causing peaks at both $\chi \to 0$ and $\chi \to 1$ if these are allowed. When using a hyper-prior to remove these regions the resulting prior predictive distribution is not flat in the spin magnitudes, i.e., it gives some preference to $\chi \approx 0.5$ over $\chi \approx 0$ or $\chi \approx 1$. Our hyper-priors result in very nearly flat prior predictive distributions over the spin magnitudes.

A.1. *Summary of TGMM population analysis method*

Here we briefly summarize the TGMM population analysis method. We refer to Hussain et al. (2024) for the full details of the method. As part of a hierarchical Bayesian inference procedure on GW data (Mandel et al. 2019; Thrane & Talbot 2019; Vitale et al. 2022), we need to efficiently estimate certain marginalized likelihoods using importance sampling. This requires samples from individual-event posteriors (with prior weights) and samples from an injection campaign to estimate detection sensitivity in different parts of parameter space. Those ingredients can allow us to infer the distribution of the population parameters in the presence of selection effects.

In general, importance sampling is needed to estimate integrals of the form

$$I(\boldsymbol{\Lambda}) = \int p(\boldsymbol{\theta}|\boldsymbol{\Lambda})\frac{p(\boldsymbol{\theta} \mid \cdot)}{W(\boldsymbol{\theta})}\mathrm{d}\boldsymbol{\theta} \approx \left\langle \frac{p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda})}{W(\boldsymbol{\theta})} \right\rangle_{\boldsymbol{\theta}\sim p(\boldsymbol{\theta}|\cdot)}\,, \tag{A2}$$

when we have samples from a distribution $p(\boldsymbol{\theta} \mid \cdot)$ conditional on some assumptions or observations (represented by "$\cdot$"), as well as access to sampling or prior weights $W(\boldsymbol{\theta})$. In particular, the posterior over the population parameters using N events is given by

$$p(\boldsymbol{\Lambda} \mid \{d_i\}) = \pi(\boldsymbol{\Lambda})\,\xi(\boldsymbol{\Lambda})^{-N}\prod_i^N \mathcal{L}(d_i \mid \boldsymbol{\Lambda})\,, \tag{A3}$$

where $\pi(\boldsymbol{\Lambda})$ is the hyperprior, $\xi(\boldsymbol{\Lambda})$ is the detection efficiency (defined below) and the $\mathcal{L}(d_i \mid \boldsymbol{\Lambda})$ are the individual event-level marginalized likelihoods, given by

$$\mathcal{L}(d_i \mid \boldsymbol{\Lambda}) = \int p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda})\frac{p(\boldsymbol{\theta} \mid d_i)}{\pi(\boldsymbol{\theta} \mid \emptyset)}\mathrm{d}\boldsymbol{\theta} \approx \left\langle \frac{p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda})}{\pi(\boldsymbol{\theta} \mid \emptyset)} \right\rangle_{\boldsymbol{\theta}\sim p(\boldsymbol{\theta}|d_i)}\,, \tag{A4}$$

where in turn $p(\boldsymbol{\theta} \mid d_i)$ is the posterior for event i, and $\pi(\boldsymbol{\theta} \mid \emptyset)$ is the sampling prior used in the original Bayesian inference for that analysis, and $p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda})$ is the population model whose parameters we wish to infer. The population-averaged detection efficiency is calculated using

$$\xi(\boldsymbol{\Lambda}) = \int p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda})P_{\mathrm{det}}(\boldsymbol{\theta})\,\mathrm{d}\boldsymbol{\theta} \approx \int p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda_0})\,P_{\mathrm{det}}(\boldsymbol{\theta})\frac{p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda})}{p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}_0)}\mathrm{d}\boldsymbol{\theta} \approx \frac{N_{\mathrm{det}}}{N_{\mathrm{draw}}}\left\langle \frac{p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda})}{p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}_0)} \right\rangle_{\boldsymbol{\theta}\sim p_{\mathrm{det}}(\boldsymbol{\theta}|\boldsymbol{\Lambda}_0)}\,, \tag{A5}$$

where $P_{\mathrm{det}}(\boldsymbol{\theta})$ is the probability of detecting a signal with parameters $\boldsymbol{\theta}$, $\boldsymbol{\Lambda}_0$ represents some fiducial population from which N_{draw} signals are simulated to estimate detection efficiency and obtain N_{det} samples of detected signals from $p_{\mathrm{det}}(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}_0)$, which is proportional, but not equal, to the fiducial population density times the detection probability, i.e., $p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda_0}) \propto p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda_0})\,\mathbf{P}_{\mathrm{det}}(\boldsymbol{\theta})$ with proportionality constant $N_{\mathrm{det}}/N_{\mathrm{draw}}$. We can see that Eqs. (A5) and (A4) are special cases of Eq. (A2).

For an integral of the form of Eq. (A2), we can take the samples from $p(\boldsymbol{\theta} \mid \cdot)$, representing the event posteriors or the detected injection set, and fit to them a weighted sum of *truncated* multivariate Gaussians

$$p(\boldsymbol{\theta} \mid \cdot) = \sum_k w_k N_{[\mathbf{a},\mathbf{b}]}(\boldsymbol{\theta} \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)\,, \tag{A6}$$

where w_k are mixture weights, and $\boldsymbol{a}$ and $\boldsymbol{b}$ are the two bounding corners of the hypercube defining the limits of the parameters $\boldsymbol{\theta}$. The fitting procedure is described in Lee & Scott (2012), with improvements drawn from Naim & Gildea (2012); Frisch & Hanebeck (2021); Salakhutdinov et al. (2003); Hussain et al. (2024). We have released a `Julia` package to perform this fitting procedure, `TruncatedGaussianMixtures.jl` and its corresponding `python` wrapper `truncatedgaussianmixtures`.

Now we assume that the population model is separable, such that the mass and redshift sector (with parameters denoted $\boldsymbol{\theta}^{m,z}$) separate from the spin sector (covering both spin magnitudes and spin tilts, denoted $\boldsymbol{\theta}^{\chi}$),

$$p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}) = p(\boldsymbol{\theta}^{\chi} \mid \boldsymbol{\Lambda}^{\chi})\, p(\boldsymbol{\theta}^{m,z} \mid \boldsymbol{\Lambda}^{m,z})\,, \tag{A7}$$

and that, in its most general form, the population model in the spin sector is some mixture of truncated multivariate Gaussians (or uniform distributions). For example, our population model for the spin magnitude and spin orientation is of this form, as seen in Eqs. (1) and (A1). In general, we can allow for any weighted sum of such separable sub-models using our methods.

To leverage this separability, while fitting the TGMMs of Eq. (A6) we require that the covariance matrix for each component does not create correlations between the spin sector and the other sectors, i.e.,

$$p(\boldsymbol{\theta} \mid \cdot) = \sum_k w_k N_{[\mathbf{a},\mathbf{b}]}(\boldsymbol{\theta}^{\chi} \mid \boldsymbol{\mu}_k^{\chi}, \boldsymbol{\Sigma}_k^{\chi})\, N_{[\mathbf{a},\mathbf{b}]}(\boldsymbol{\theta}^{m,z} \mid \boldsymbol{\mu}_k^{m,z}, \boldsymbol{\Sigma}_k^{m,z})\,. \tag{A8}$$

This does not impose strong restrictions in the distributions that can be fit and, in particular, does not imply that we cannot capture correlations across the two sectors: although each individual TGMM component cannot have cross-sector correlations across, cross-sector correlation structure can still be captured by the joint arrangement of multiple TGMM components. Indeed, the action of several uncorrelated components together can construct correlations (e.g., as an extreme case, KDEs with uncorrelated bandwidth matrices can easily represent distributions with large-scale covariances).

By substituting Eqs. (A7) and (A8) into (A2) with $W(\boldsymbol{\theta})$ set to the sampling prior for concreteness (equivalently, the sampling distribution for sensitivity injections), we get,

$$I(\boldsymbol{\Lambda}) = \sum_k w_k \int N_{[\mathbf{a},\mathbf{b}]}(\boldsymbol{\theta}^{\chi} \mid \boldsymbol{\mu}_k^{\chi}, \boldsymbol{\Sigma}_k^{\chi})\, p(\boldsymbol{\theta}^{\chi} \mid \boldsymbol{\Lambda}^{\chi})\, \mathrm{d}\boldsymbol{\theta}^{\chi} \int N_{[\mathbf{a},\mathbf{b}]}(\boldsymbol{\theta}^{m,z} \mid \boldsymbol{\mu}_k^{m,z}, \boldsymbol{\Sigma}_k^{m,z})\, \frac{p(\boldsymbol{\theta}^{m,z} \mid \boldsymbol{\Lambda}^{m,z})}{\pi(\boldsymbol{\theta}^{m,z} \mid \emptyset)}\, \mathrm{d}\boldsymbol{\theta}^{m,z}\,, \tag{A9}$$

where we have assumed for now that the sampling prior in spin space is flat, i.e., $\pi(\boldsymbol{\theta}^{\chi} \mid \emptyset) = 1$, which is true of standard LVK priors on the spin magnitudes and cosines of the spin tilts (this constraint is removed in the full description of our methodology in Hussain et al. 2024). We then rewrite the above as

$$I(\boldsymbol{\Lambda}) = \sum_k w_k\, I(\boldsymbol{\mu}_k^{\chi}, \boldsymbol{\Sigma}_k^{\chi}, \boldsymbol{\Lambda}_{\chi}) \left\langle \frac{p(\boldsymbol{\theta}^{m,z} \mid \boldsymbol{\Lambda}^{m,z})}{\pi(\boldsymbol{\theta}^{m,z} \mid \emptyset)} \right\rangle_{\boldsymbol{\theta}_{m,z} \sim N_{[\mathbf{a},\mathbf{b}]}(\boldsymbol{\theta}^{m,z} \mid \boldsymbol{\mu}_k^{m,z}, \boldsymbol{\Sigma}_k^{m,z})}\,. \tag{A10}$$

Here the integral $I(\boldsymbol{\mu}_k^{\chi}, \boldsymbol{\Sigma}_k^{\chi}, \boldsymbol{\Lambda}^{\chi})$ can be handled semi-analytically using propoerties of Gaussians (using custom numerical routines we implement), while the expectation over the mass and redshift can be approximated using Monte Carlo estimation, as is standard int the LVK literature. To handle cuts and reduce stabilize the TGMM fit in the mass and redshift domains, we note that each TGMM fit allows us to extract an assignment of each sample $\boldsymbol{\theta}_i$ to a specific TGMM component k. Armed with this assignment, we can then rewrite the expectation over the mass and redshift as an expectation over the *original* posterior samples that were assigned to each component k. This leaves us with the expression

$$I(\boldsymbol{\Lambda}) = \sum_k w_k\, I(\boldsymbol{\mu}_k^{\chi}, \boldsymbol{\Sigma}_k^{\chi}, \boldsymbol{\Lambda}^{\chi}) \left[\frac{1}{N_k} \sum_j^{N_k} \frac{p(\boldsymbol{\theta}_{j,k}^{m,z} \mid \boldsymbol{\Lambda}^{m,z})}{\pi(\boldsymbol{\theta}_{j,k}^{m,z} \mid \emptyset)} \right]\,, \tag{A11}$$

where $\boldsymbol{\theta}_{j,k}$ is the jth sample assigned to the kth component. Equation (A11) is efficient to evaluate.

We explore the variance properties of the above estimator (A11) for $I(\mathbf{\Lambda})$ in the companion paper Hussain et al. (2024), and find no singular behaviour of the estimator's variance as population features become narrow. Since we only perform the Monte Carlo estimate over the mass and redshift sectors, the variance of this estimator is lower than in analyses using Monte Carlo methods across all sectors. We have not computed the variance of the population likelihood estimators explicitly in this study, nor applied any cuts associated with the variance of the estimator at given hyper-parameter values (e.g., the data-dependent priors discussed in Abbott et al. 2023b).

A.2. *Details on the computation of SDDR*

To compute a SDDR, we need an estimate of the posterior density at a point of interest that often lies at the edges of our domain (e.g., $\chi = 0$). We use a custom multivariate KDE to get an estimate of the marginalized hyper-posterior from samples. We describe this multivariate KDE in Hussain et al. (2024) but note here that it does not impose a zero derivative at the edge and has no bias at the boundary at $\mathcal{O}(b^0)$, where b is the kernel bandwidth, an issue that often plagues other KDE techniques and makes them unsuitable for SDDRs. Since the bias in our method is $\mathcal{O}(b)$, we can achieve a better estimate of the density at the edge of parameter space by increasing the number of samples (reducing the bandwidth). Briefly, our KDE maps each point to a truncated multivariate normal with some uncorrelated bandwidth vector $\mathbf{b}$ ($b = |\mathbf{b}|$), but the position of this multivariate normal is moved away from the location of the sample in such a way that the overall KDE estimate has a bias of only order $\mathcal{O}(b)$. This is similar to what can be achieved using KDEs with reflective boundary conditions, such as those used by Callister et al. (2022) to compute SDDRs. A benefit of our method as compared to reflective KDEs is that we do not require the derivative of the kernel be zero at the boundaries, and we do not need the additional kernels to enforce reflective boundaries (which can require a large number of additional kernels in higher dimensions).

We highlight our procedure of extracting BFs and their associated bootstrap uncertainties using the SDDR with an illustrative example. Consider computing the SDDR for our fiducial spin-magnitude model, which supports two subpopulations, in the limit that all the secondary BHs in the dominant population are nonspinning. This corresponds to $\mu_2^a \to 0$ and $\sigma_2^a \to 0$. We first draw N samples with replacement from our hyper-posterior samples. In this two-dimensional space we then get a KDE estimate of the marginal hyper-posterior over μ_2^a and σ_2^a, $\hat{p}(\mu_2^a, \sigma_2^a)$. Then we evaluate the estimator for the BF, $\hat{\mathcal{B}}$, given by

$$\hat{\mathcal{B}} = \frac{\hat{p}(\mu_2^a = 0, \sigma_2^a = 0)}{\pi(\mu_2^a = 0, \sigma_2^a = 0)}\,, \tag{A12}$$

where $\pi(\mu_2^a, \sigma_2^a)$ is the marginalized hyper-prior for our analysis. This gives us an estimate of the BF from one sampling of our hyper-posterior. We subsequently repeat this process $\mathcal{O}(100)$ times, and get a series of BF estimates. With these estimates we compute the bootstrapped median and 90% highest-confidence interval for our estimate of the BF.

A.3. *Details on the computation of BF for event assignments*

We use the following method to compute the BF for a given event in favor of the hypothesis that the even belongs to a given of a particular subpopulation, which is used in Fig. 5. Assume that the population model breaks into two subpopulations, subpopulation A and subpopulation B, with the fraction η of the population in A. The evidence of some event i under the population prior described by $\mathbf{\Lambda}$ is given by $Z_i(\mathbf{\Lambda}) = \mathcal{L}(d_i|\mathbf{\Lambda})$ as defined in (A4).

We take our hyper-posterior samples, and for each sample $\mathbf{\Lambda}_j$, we compute the evidence $Z_i(\mathbf{\Lambda}_j^{\eta=1})$ under the hypothesis that only subpopulation A exists, $\eta \to 1$, and then compute the evidence $Z_i(\mathbf{\Lambda}_j^{\eta=0})$ under the hypothesis that only subpopulation B exists, $\eta \to 0$. Here $\mathbf{\Lambda}_j^{\eta=0}$ simply means we set $\eta = 0$ for that sample. The ratio gives our estimate of the BF between the two population hypotheses for event i

$$\mathcal{B}_{i,j} = \frac{Z_i(\mathbf{\Lambda}_j^{\eta=1})}{Z_i(\mathbf{\Lambda}_j^{\eta=0})}\,, \tag{A13}$$

giving us our posterior over BFs for event i. We use these to compute the expected $\log_{10}\mathcal{B}_i$ from our posterior over $\mathcal{B}_i$ using the samples $\mathcal{B}_{i,j}$ as reported in Fig. 5 for the spin-orientation subpopulations and the spin-magnitude subpopulations.

B. FURTHER RESULTS: CORNER PLOTS

In this Appendix we show corner plots for the parameters governing each subpopulation in our fiducial two-population spin-magnitude analysis. Figure 6 shows the two-dimensional and one-dimensional marginals for the parameters of the

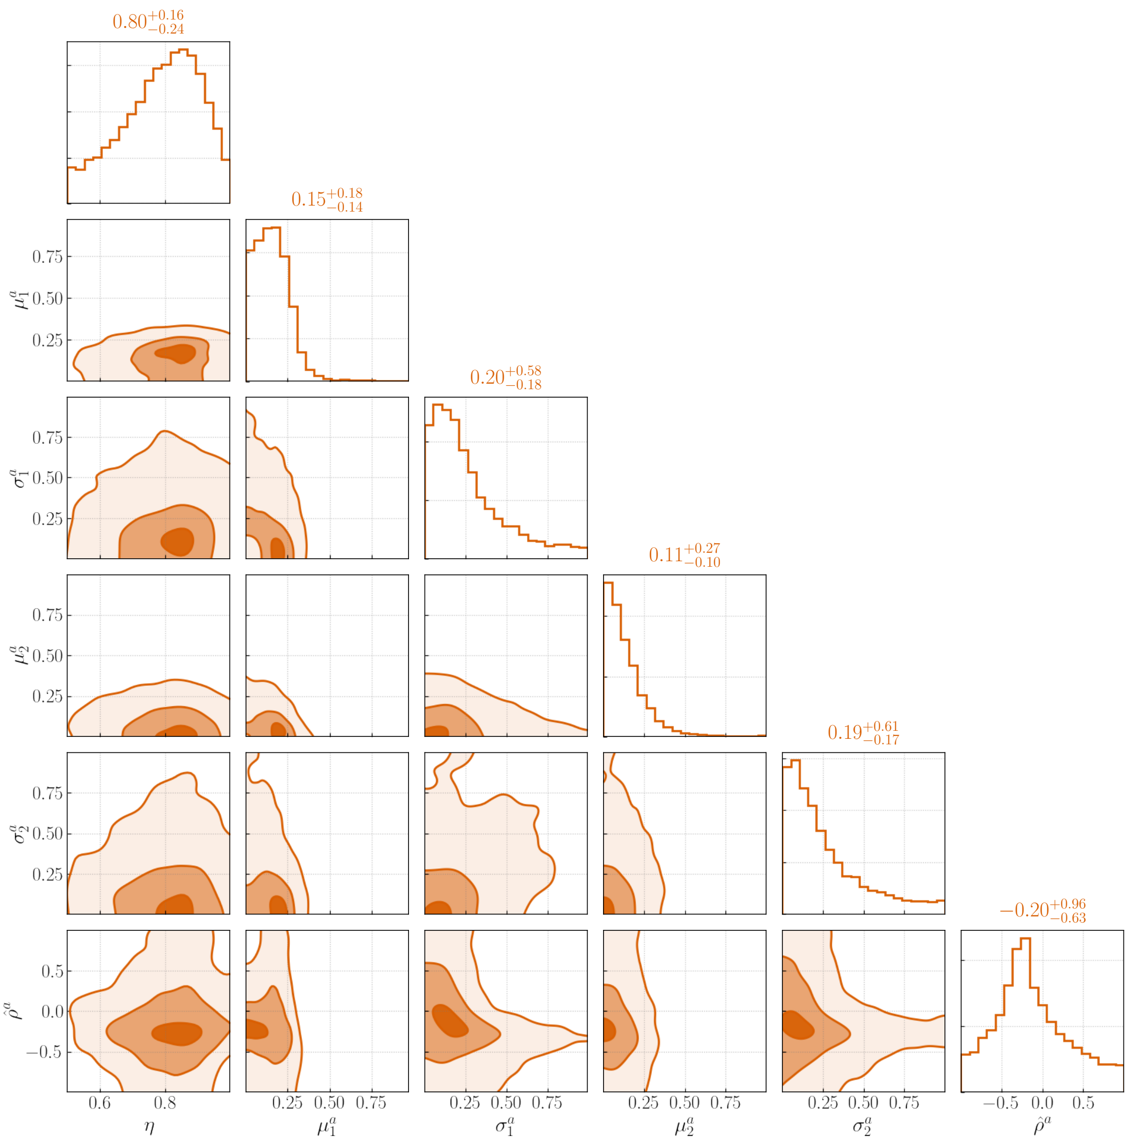

Figure 6. Corner plot for the posteriors of the hyper-parameters pertaining to the dominant subpopulation.

dominant, slowly spinning population. Figure 7 shows the same marginals but for the parameters of the subdominant, rapidly spinning population. As compared to the parameters of the dominant population, the parameters of this component are less informed. The parameters controlling the mean μ_i^b display the bi-modality discussed above, with a preferred mode at large μ_1^b–μ_2^b values and a less significant mode at small μ_1^b–μ_2^b values. This second mode is due to residual degeneracy between the two subpopulations, with these small values possible only when the fraction η is close to 1/2.

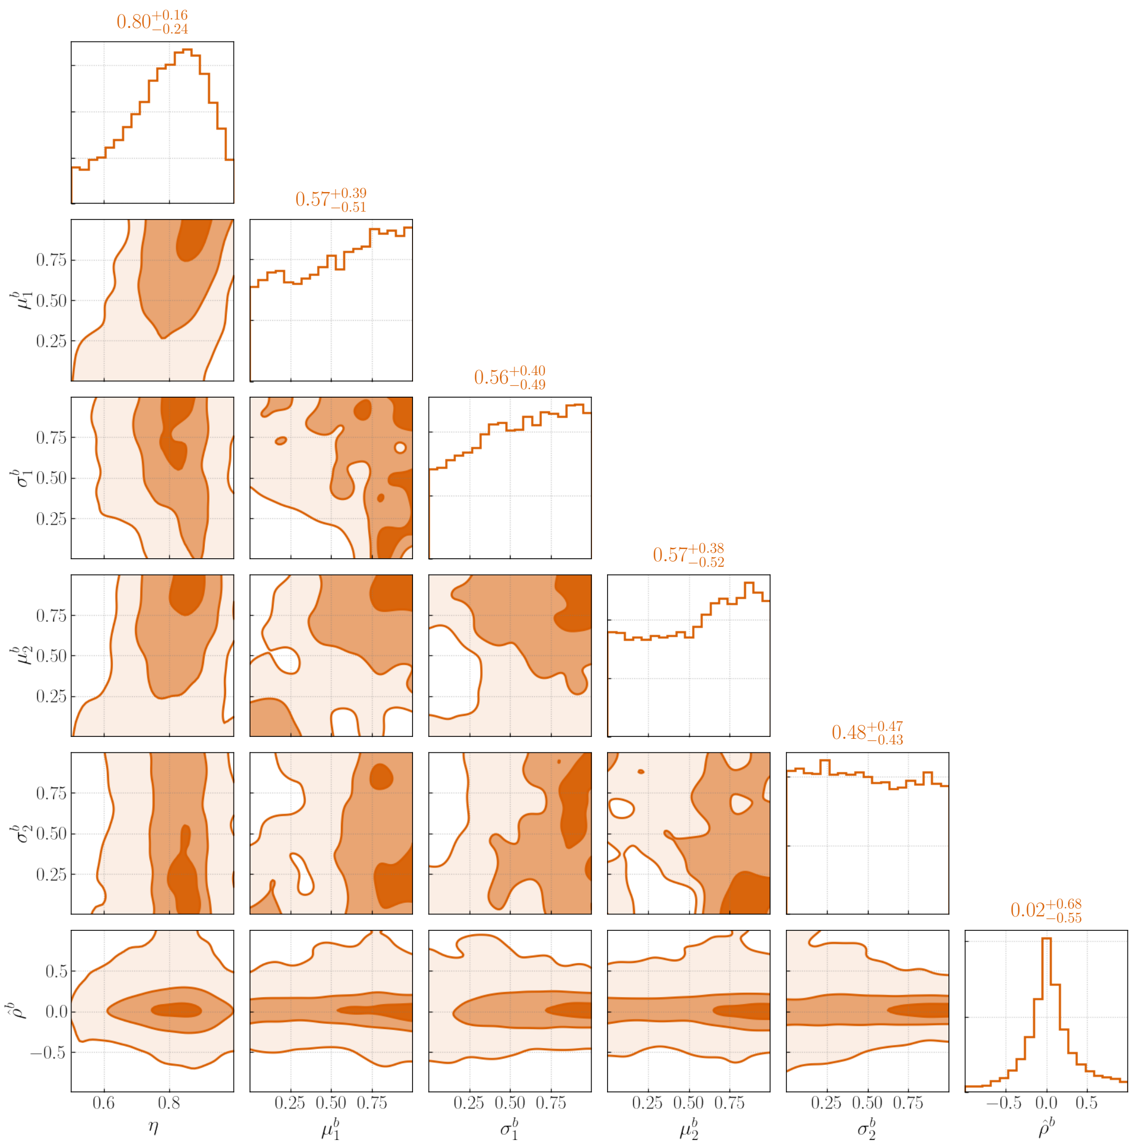

Figure 7. Corner plot for the posteriors of the hyper-parameters pertaining to the subdominant subpopulation.

Figure 8 shows the PPDs of the individual spins from our fiducial two-population analysis, with and without the event GW190517 included. The existence of a highly spinning subpopulation is difficult to infer from these individual marginal PPDs, appearing only as a heavier tail in the χ_1 PPD when GW190517 is included.

REFERENCES

Aasi, J., et al. 2015, Class. Quant. Grav., 32, 074001, doi: 10.1088/0264-9381/32/7/074001

Abbott, B. P., et al. 2016, Living Rev. Rel., 19, 1, doi: 10.1007/s41114-020-00026-9

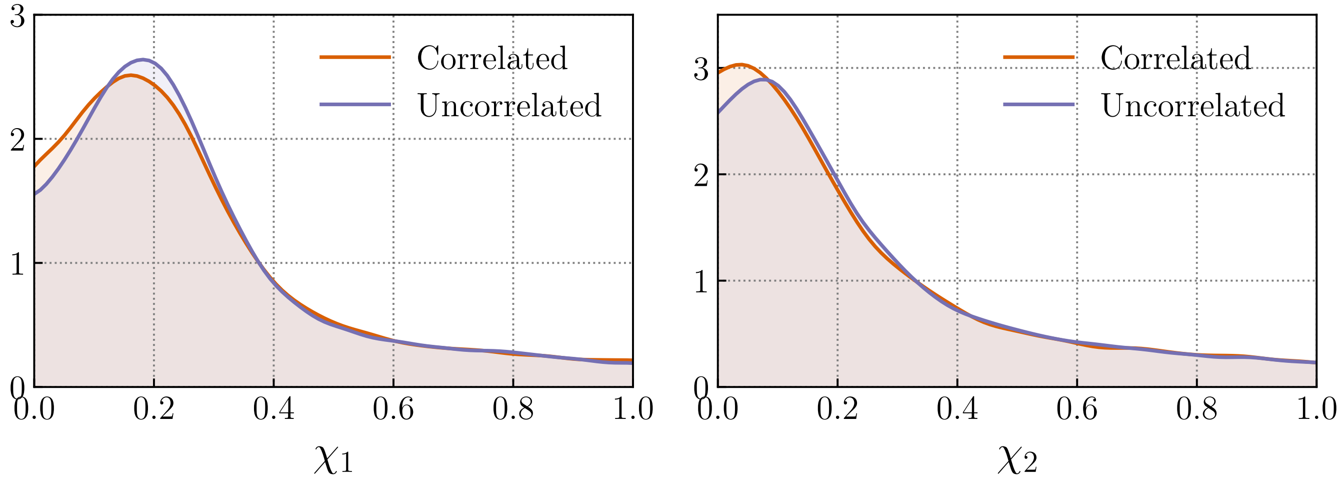

Figure 8. Marginals over the PPD of the primary and secondary spin magnitude χ_1 and χ_2. We compare the analysis allowing sub-populations to have correlations between χ_1 and χ_2, to one where they are not correlated. One can see that the introduction of correlations to the model does not leave a very significant imprint on the marignal distributions, but can be seen much more clearly on the two dimensional PPD (Fig. 1).

—. 2019a, Phys. Rev. X, 9, 031040, doi: 10.1103/PhysRevX.9.031040

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

Abbott, R., et al. 2021a, Phys. Rev. X, 11, 021053, doi: 10.1103/PhysRevX.11.021053

—. 2021b, Astrophys. J. Lett., 913, L7, doi: 10.3847/2041-8213/abe949

—. 2021c, SoftwareX, 13, 100658, doi: 10.1016/j.softx.2021.100658

—. 2023a, Phys. Rev. X, 13, 041039, doi: 10.1103/PhysRevX.13.041039

—. 2023b, Phys. Rev. X, 13, 011048, doi: 10.1103/PhysRevX.13.011048

—. 2023c, Astrophys. J. Suppl., 267, 29, doi: 10.3847/1538-4365/acdc9f

—. 2024, Phys. Rev. D, 109, 022001, doi: 10.1103/PhysRevD.109.022001

Acernese, F., et al. 2015, Class. Quant. Grav., 32, 024001, doi: 10.1088/0264-9381/32/2/024001

Adamcewicz, C., Galaudage, S., Lasky, P. D., & Thrane, E. 2024, The Astrophysical Journal Letters, 964, L6, doi: 10.3847/2041-8213/ad2df2

Adamcewicz, C., & Thrane, E. 2022, Mon. Not. Roy. Astron. Soc., 517, 3928, doi: 10.1093/mnras/stac2961

Ajith, P., et al. 2011, Phys. Rev. Lett., 106, 241101, doi: 10.1103/PhysRevLett.106.241101

Akutsu, T., et al. 2021, PTEP, 2021, 05A101, doi: 10.1093/ptep/ptaa125

Antoni, A., & Quataert, E. 2021, Monthly Notices of the Royal Astronomical Society, 511, 176, doi: 10.1093/mnras/stab3776

Bavera, S. S., Fragos, T., Qin, Y., et al. 2020, Astron. Astrophys., 635, A97, doi: 10.1051/0004-6361/201936204

Bavera, S. S., Fragos, T., Zevin, M., et al. 2021, A&A, 647, A153, doi: 10.1051/0004-6361/202039804

Belczynski, K., et al. 2020, Astron. Astrophys., 636, A104, doi: 10.1051/0004-6361/201936528

Betancourt, M. 2017, arXiv e-prints, arXiv:1701.02434, doi: 10.48550/arXiv.1701.02434

Bingham, E., Chen, J. P., Jankowiak, M., et al. 2019, J. Mach. Learn. Res., 20, 28:1. http://jmlr.org/papers/v20/18-403.html

Biscoveanu, S., Callister, T. A., Haster, C.-J., et al. 2022, Astrophys. J. Lett., 932, L19, doi: 10.3847/2041-8213/ac71a8

Biscoveanu, S., Isi, M., Varma, V., & Vitale, S. 2021a, Phys. Rev. D, 104, 103018, doi: 10.1103/PhysRevD.104.103018

Biscoveanu, S., Isi, M., Vitale, S., & Varma, V. 2021b, Phys. Rev. Lett., 126, 171103, doi: 10.1103/PhysRevLett.126.171103

Broekgaarden, F. S., Stevenson, S., & Thrane, E. 2022, Astrophys. J., 938, 45, doi: 10.3847/1538-4357/ac8879

Buscicchio, R., Roebber, E., Goldstein, J. M., & Moore, C. J. 2019, PhRvD, 100, 084041, doi: 10.1103/PhysRevD.100.084041

Callister, T. A., Farr, W. M., & Renzo, M. 2021a, Astrophys. J., 920, 157, doi: 10.3847/1538-4357/ac1347

Callister, T. A., Haster, C.-J., Ng, K. K. Y., Vitale, S., & Farr, W. M. 2021b, Astrophys. J. Lett., 922, L5, doi: 10.3847/2041-8213/ac2ccc

Callister, T. A., Miller, S. J., Chatziioannou, K., & Farr, W. M. 2022, Astrophys. J. Lett., 937, L13, doi: 10.3847/2041-8213/ac847e

Chatziioannou, K., Lovelace, G., Boyle, M., et al. 2018, Phys. Rev. D, 98, 044028, doi: 10.1103/PhysRevD.98.044028

Cho, H.-S., Ochsner, E., O'Shaughnessy, R., Kim, C., & Lee, C.-H. 2013, Phys. Rev. D, 87, 024004, doi: 10.1103/PhysRevD.87.024004

Damour, T. 2001, Phys. Rev. D, 64, 124013, doi: 10.1103/PhysRevD.64.124013

de Mink, S. E., & Mandel, I. 2016, Mon. Not. Roy. Astron. Soc., 460, 3545, doi: 10.1093/mnras/stw1219

Dickey, J. M. 1971, The Annals of Mathematical Statistics, 204

Doctor, Z., Wysocki, D., O'Shaughnessy, R., Holz, D. E., & Farr, B. 2020, ApJ, 893, 35, doi: 10.3847/1538-4357/ab7fac

Edelman, B., Farr, B., & Doctor, Z. 2023, Astrophys. J., 946, 16, doi: 10.3847/1538-4357/acb5ed

Farr, W. M. 2019, Research Notes of the AAS, 3, 66, doi: 10.3847/2515-5172/ab1d5f

Franciolini, G., & Pani, P. 2022, Phys. Rev. D, 105, 123024, doi: 10.1103/PhysRevD.105.123024

Frisch, D., & Hanebeck, U. D. 2021, Gaussian Mixture Estimation from Weighted Samples, arXiv, doi: 10.48550/ARXIV.2106.05109

Fuller, J., & Lu, W. 2022, Mon. Not. Roy. Astron. Soc., 511, 3951, doi: 10.1093/mnras/stac317

Fuller, J., & Ma, L. 2019, Astrophys. J. Lett., 881, L1, doi: 10.3847/2041-8213/ab339b

Fuller, J., Piro, A. L., & Jermyn, A. S. 2019, Mon. Not. Roy. Astron. Soc., 485, 3661, doi: 10.1093/mnras/stz514

Galaudage, S., Talbot, C., Nagar, T., et al. 2021, The Astrophysical Journal Letters, 921, L15, doi: 10.3847/2041-8213/ac2f3c

García-Bellido, J., Nuño Siles, J. F., & Ruiz Morales, E. 2021, Physics of the Dark Universe, 31, 100791, doi: https://doi.org/10.1016/j.dark.2021.100791

Gerosa, D., & Berti, E. 2017, PhRvD, 95, 124046, doi: 10.1103/PhysRevD.95.124046

Gerosa, D., & Fishbach, M. 2021, Nature Astronomy, 5, 749, doi: 10.1038/s41550-021-01398-w

Gerosa, D., Kesden, M., Berti, E., O'Shaughnessy, R., & Sperhake, U. 2013, Phys. Rev. D, 87, 104028, doi: 10.1103/PhysRevD.87.104028

Ghosh, A., Del Pozzo, W., & Ajith, P. 2016, Phys. Rev. D, 94, 104070, doi: 10.1103/PhysRevD.94.104070

Golomb, J., & Talbot, C. 2023, Phys. Rev. D, 108, 103009, doi: 10.1103/PhysRevD.108.103009

Green, R., Hoy, C., Fairhurst, S., et al. 2021, Phys. Rev. D, 103, 124023, doi: 10.1103/PhysRevD.103.124023

Hoffman, M. D., & Gelman, A. 2011, arXiv e-prints, arXiv:1111.4246, doi: 10.48550/arXiv.1111.4246

Hotokezaka, K., & Piran, T. 2017, Astrophys. J., 842, 111, doi: 10.3847/1538-4357/aa6f61

Hoy, C., Fairhurst, S., Hannam, M., & Tiwari, V. 2022, ApJ, 928, 75, doi: 10.3847/1538-4357/ac54a3

Hussain, A., Isi, M., & Zimmerman, A. 2024, Reconstructing the Edges: Testing for compact population features at the edges of parameter space, In preparation

Kimball, C., Talbot, C., Berry, C. P. L., et al. 2020, ApJ, 900, 177, doi: 10.3847/1538-4357/aba518

—. 2021, ApJL, 915, L35, doi: 10.3847/2041-8213/ac0aef

Kimball, C., et al. 2021, Astrophys. J. Lett., 915, L35, doi: 10.3847/2041-8213/ac0aef

Kushnir, D., Zaldarriaga, M., Kollmeier, J. A., & Waldman, R. 2016, Mon. Not. Roy. Astron. Soc., 462, 844, doi: 10.1093/mnras/stw1684

Lee, G., & Scott, C. 2012, Computational Statistics and Data Analysis, 56, 2816–2829, doi: 10.1016/j.csda.2012.03.003

LIGO Scientific, Virgo and KAGRA Collaborations. 2021, 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, Zenodo, doi: 10.5281/zenodo.5636816

LIGO Scientific, Virgo, and KAGRA Collaborations. 2024, Gravitational Wave Open Science Center, https://www.gw-openscience.org/

Loredo, T. J. 2004, AIP Conf. Proc., 735, 195, doi: 10.1063/1.1835214

Ma, L., & Fuller, J. 2019, Mon. Not. Roy. Astron. Soc., 488, 4338, doi: 10.1093/mnras/stz2009

—. 2023, Astrophys. J., 952, 53, doi: 10.3847/1538-4357/acdb74

Mandel, I., & de Mink, S. E. 2016, Mon. Not. Roy. Astron. Soc., 458, 2634, doi: 10.1093/mnras/stw379

Mandel, I., & Farmer, A. 2022, Phys. Rept., 955, 1, doi: 10.1016/j.physrep.2022.01.003

Mandel, I., Farr, W. M., & Gair, J. R. 2019, Mon. Not. Roy. Astron. Soc., 486, 1086, doi: 10.1093/mnras/stz896

Mapelli, M. 2020, Proc. Int. Sch. Phys. Fermi, 200, 87, doi: 10.3254/ENFI200005

Marchant, P., Langer, N., Podsiadlowski, P., Tauris, T. M., & Moriya, T. J. 2016, Astron. Astrophys., 588, A50, doi: 10.1051/0004-6361/201628133

McKernan, B., & Ford, K. E. S. 2024, MNRAS, 531, 3479, doi: 10.1093/mnras/stae1351

Mehta, A. K., Olsen, S., Wadekar, D., et al. 2023. https://arxiv.org/abs/2311.06061

Miller, S., Callister, T. A., & Farr, W. 2020, Astrophys. J., 895, 128, doi: 10.3847/1538-4357/ab80c0

Miller, S. J., Isi, M., Chatziioannou, K., Varma, V., & Mandel, I. 2024a, Phys. Rev. D, 109, 024024, doi: 10.1103/PhysRevD.109.024024

Miller, S. J., Ko, Z., Callister, T., & Chatziioannou, K. 2024b, Phys. Rev. D, 109, 104036, doi: 10.1103/PhysRevD.109.104036

Mould, M., Gerosa, D., Broekgaarden, F. S., & Steinle, N. 2022, Mon. Not. Roy. Astron. Soc., 517, 2738, doi: 10.1093/mnras/stac2859

Naim, I., & Gildea, D. 2012, Convergence of the EM Algorithm for Gaussian Mixtures with Unbalanced Mixing Coefficients, arXiv, doi: 10.48550/ARXIV.1206.6427

Neal, R. 2011, in Handbook of Markov Chain Monte Carlo (Chapman & Hall/CRC), 113–162, doi: 10.1201/b10905

Nitz, A. H., Capano, C., Nielsen, A. B., et al. 2019, Astrophys. J., 872, 195, doi: 10.3847/1538-4357/ab0108

Nitz, A. H., Capano, C. D., Kumar, S., et al. 2021, Astrophys. J., 922, 76, doi: 10.3847/1538-4357/ac1c03

Nitz, A. H., Kumar, S., Wang, Y.-F., et al. 2023, Astrophys. J., 946, 59, doi: 10.3847/1538-4357/aca591

Nitz, A. H., Dent, T., Davies, G. S., et al. 2020, Astrophys. J., 891, 123, doi: 10.3847/1538-4357/ab733f

Olejak, A., & Belczynski, K. 2021, Astrophys. J. Lett., 921, L2, doi: 10.3847/2041-8213/ac2f48

Olsen, S., Venumadhav, T., Mushkin, J., et al. 2022, Phys. Rev. D, 106, 043009, doi: 10.1103/PhysRevD.106.043009

O'Shaughnessy, R., Farr, B., Ochsner, E., et al. 2014, Phys. Rev. D, 89, 102005, doi: 10.1103/PhysRevD.89.102005

Payne, E., Kremer, K., & Zevin, M. 2024, Astrophys. J. Lett., 966, L16, doi: 10.3847/2041-8213/ad3e82

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

Pratten, G., Schmidt, P., Buscicchio, R., & Thomas, L. M. 2020, Phys. Rev. Res., 2, 043096, doi: 10.1103/PhysRevResearch.2.043096

Pratten, G., et al. 2021, Phys. Rev. D, 103, 104056, doi: 10.1103/PhysRevD.103.104056

Qin, Y., Fragos, T., Meynet, G., et al. 2018, Astron. Astrophys., 616, A28, doi: 10.1051/0004-6361/201832839

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

Qin, Y., et al. 2022, Astrophys. J., 941, 179, doi: 10.3847/1538-4357/aca40c

Raymond, V., van der Sluys, M. V., Mandel, I., et al. 2010, Class. Quant. Grav., 27, 114009, doi: 10.1088/0264-9381/27/11/114009

Riley, J., Mandel, I., Marchant, P., et al. 2021, Mon. Not. Roy. Astron. Soc., 505, 663, doi: 10.1093/mnras/stab1291

Rodriguez, C. L., Zevin, M., Amaro-Seoane, P., et al. 2019, PhRvD, 100, 043027, doi: 10.1103/PhysRevD.100.043027

Roulet, J., Chia, H. S., Olsen, S., et al. 2021, Phys. Rev. D, 104, 083010, doi: 10.1103/PhysRevD.104.083010

Salakhutdinov, R., Roweis, S., & Ghahramani, Z. 2003, in Proceedings of the Twentieth International Conference on International Conference on Machine Learning, ICML'03 (AAAI Press), 672–679

Steinle, N., & Kesden, M. 2021, PhRvD, 103, 063032, doi: 10.1103/PhysRevD.103.063032

Thrane, E., & Talbot, C. 2019, Publ. Astron. Soc. Austral., 36, e010, doi: 10.1017/pasa.2019.2

Tiwari, V., & Fairhurst, S. 2021, Astrophys. J. Lett., 913, L19, doi: 10.3847/2041-8213/abfbe7

Tong, H., Galaudage, S., & Thrane, E. 2022, Physical Review D, 106, doi: 10.1103/physrevd.106.103019

Udall, R., Hourihane, S., Miller, S., et al. 2024. https://arxiv.org/abs/2409.03912

van der Sluys, M. V., Röver, C., Stroeer, A., et al. 2008, Astrophys. J. Lett., 688, L61, doi: 10.1086/595279

Varma, V., Isi, M., Biscoveanu, S., Farr, W. M., & Vitale, S. 2022, Phys. Rev. D, 105, 024045, doi: 10.1103/PhysRevD.105.024045

Venumadhav, T., Zackay, B., Roulet, J., Dai, L., & Zaldarriaga, M. 2019, Phys. Rev. D, 100, 023011, doi: 10.1103/PhysRevD.100.023011

—. 2020, Phys. Rev. D, 101, 083030, doi: 10.1103/PhysRevD.101.083030

Vitale, S., Gerosa, D., Farr, W. M., & Taylor, S. R. 2022, in Handbook of Gravitational Wave Astronomy, ed. C. Bambi, S. Katsanevas, & K. D. Kokkotas (Springer, Singapore), 45, doi: 10.1007/978-981-15-4702-7_45-1

Vitale, S., Lynch, R., Veitch, J., Raymond, V., & Sturani, R. 2014, Phys. Rev. Lett., 112, 251101, doi: 10.1103/PhysRevLett.112.251101

Wysocki, D., Gerosa, D., O'Shaughnessy, R., et al. 2018, PhRvD, 97, 043014, doi: 10.1103/PhysRevD.97.043014

Zackay, B., Dai, L., Venumadhav, T., Roulet, J., & Zaldarriaga, M. 2021, Phys. Rev. D, 104, 063030, doi: 10.1103/PhysRevD.104.063030

Zackay, B., Venumadhav, T., Dai, L., Roulet, J., & Zaldarriaga, M. 2019, Phys. Rev. D, 100, 023007, doi: 10.1103/PhysRevD.100.023007

Zaldarriaga, M., Kushnir, D., & Kollmeier, J. A. 2018, Mon. Not. Roy. Astron. Soc., 473, 4174, doi: 10.1093/mnras/stx2577

Zevin, M., & Bavera, S. S. 2022, Astrophys. J., 933, 86, doi: 10.3847/1538-4357/ac6f5d