LIGO SURF Interim Report 2:

Inferring the Population of Merging Black Holes with Astrophysically Motivated Models

September 23, 2023

Author: April Qiu Cheng (aqc@mit.edu)
Direct Supervisor: Jacob Golomb (jgolomb@caltech.edu)
Faculty Supervisor: Alan Weinstein (ajw@caltech.edu)

1 Introduction

Gravitational waves (GWs)—first prediced by Einstein shortly after his introduction of General Relativity but at the time predicted to be unobservable—are propagating perturbations in the spacetime metric. In 2015, they were directly detected by the Laser Interferometer Gravitational-wave Observatory (LIGO) [1]. Since gravitational waves are produced by a time-varying mass quadrupole moment, thus far only merging binary compact objects (black holes (BHs) and neutron stars) have produced gravitational waves detectable by our current gravitational wave experiments LIGO and Virgo. Nonetheless, this has opened up an entirely new field of GW astronomy. With the third LIGO-Virgo GW transient catalog (GWTC-3), we have now $\mathcal{O}(100)$ detections of gravitational waves, most of which are binary black holes (BBHs). It is now possible to not only infer individual source properties, but also to probe the underlying population from which these BBHs arise, placing constraints on models of stellar evolution and astrophysics. Interesting population properties that can be probed with the GW dataset include the shape of the BH mass distribution [23, 2], the BH spin distribution [22, 7, 12, 29, 2], the distribution of BBHs across the sky [18], the evolution of merger rate and other properties with redshift [11, 7, 2], correlations between properties such as mass and spin [7, 13, 30, 4, 2, 16], and even branching ratios between different BBH formation channels [31, 15, 28, 19].

Such analyses are usually done with a method called hierarchical Bayesian inference. BBH systems are completely described by 15 parameters: two characterizing the mass of each BH, 6 characterizing the 3D spin vectors of each BH, and seven extrinsic parameters that describe the sky location, distance, and orientation of the system. However, these parameters cannot be extracted straightforwardly from the GW strain data; for example, the majority of the spin information we obtain from a detection is from a measurement of the effective dimensionless spin parameter (χ_{eff}), which is a mass-weighted combination of the spin magnitude projected onto the orbital angular momentum, since it is the leading-order spin term in the post-Newtonian expansion of the strain. Furthermore, there exist degeneracies between some

of the parameters. In the GW literature, these parameters are extracted for each system using Bayesian inference in a process called parameter estimation (PE). Therefore, what we have for each detection is not an exact measurement of any parameter, but rather a 15-dimensional *posterior distribution* describing the probability that the system has a certain set of parameters, given the detected signal. It is these posterior distributions that are then used to infer properties of the underlying population. One can then choose an astrophysically-motivated parameterization of the population and fit the posterior samples to the model in another layer of Bayesian inference in a process called hierarchical Bayesian analysis, the output of which is a posterior distribution on the value of the population parameters; it is also possible to compare the likelihoods between the models to discern which one is more favored. Analyses using flexible approaches (e.g. splines, Gaussian processes), which have the advantage of not being model-dependent, have also been explored. See [27] for a pedagogical introduction to both PE and hierarchical inference in the context of GW astronomy.

One of the primary goals of hierarchical inference in GW astronomy is to learn about the formation origins of the detected BBHs. Several different formation mechanisms have been explored in the literature, each of which have different signatures in BBH parameter space, subject to modelling uncertainties. These formation channels can be roughly grouped into two categories: dynamical formation channels, which involve BBHs formed via random dynamical capture in dense stellar environments, and field channels, which involve the isolated co-evolution of a binary star system. Recent literature has suggested that BBHs that form via field channels have mass-spin correlations; in particular, there is a negative correlation between the mass and the spin magnitude of the secondary star, which can be spun up via tidal forces [5], although the strength of this correlation is uncertain [17]. Hierarchical mergers in dynamical channels, involving the mergers of merger remnants, also have a mass-spin correlation, as the 2nd-generation component should have both a higher mass and a characteristic high spin-magnitude of around $a \sim 0.7$ [15]. Additionally, there are detections in our current catalog with extreme mass-spin properties that are difficult to explain via isolated evolution [30]. The goal of this project is to employ hierarchical Bayesian inference to fit a spin-mass correlated model, motivated by the astrophysical literature, to gravitational-wave data to learn about the formation mechanisms of BBHs. In particular, recent models suggest that in isolated formation channels, there is a correlation between the mass and spin magnitude of the secondary BH [17], which is spun-up through tidal interactions during the second mass-transfer phase. The goal of this project is to fit a heuristic model to recover this correlation in both real and injected data, and possibly extend this model to a correlation in redshift to model the delay time distribution as a function of the masses or spins.

2 Method

To do the hierarchical inference, we use the Python hierarchical inference code `gwpopulation` [24], with the `dynesty` nested sampler [20] as implemented by `bilby` [3]. To account for selection effects, we use the method described in [10] to compute the selection term $\alpha(\Lambda)$. We use the search sensitivity samples from the publicly-released simulated injection campaign [9], and consider a sample found using the same criteria as the GWTC-3 data release, i.e. if the optimal SNR $\rho_{\rm opt} > 10$ for O1 and O2, or if $\mathrm{FAR}_{\rm min} < 1\ \mathrm{yr}^{-1}$ for O3, where $\mathrm{FAR}_{\rm min}$ is

the minimum False Alarm Rate (FAR) across all search pipelines.

We have found that there are significant regions of parameter space that have a very low number of effective Monte-Carlo samples ($N_{\rm eff}$), especially at negative values of $\delta \log \sigma$, whose narrow Gaussians are poorly sampled from the individual event $\chi_{\rm eff}$ posteriors. To remedy this issue, we assign a likelihood of 0 to regions of parameter space with $N_{\rm eff} < N$, where N is the number of events, and switch from using `C01:Mixed` to `C01:IMRPhenomXPHM` PE samples for our analysis of GWTC-3 data in order to have more samples. We implement this with the `MinimumEffectiveSamplesLikelihood` class from `gwpopulation_pipe` [21], the interface package to `gwpopulation`.

3 Models

In the standard literature, labels for each BH are assigned based on mass-sorting, i.e. m_1 and a_1 are the mass and spin of the BH with greater mass, while m_2 and a_2 correspond to the secondary BH. In our models, we employ the spin-sorting method pioneered by [6], such that m_A and a_A refer to the higher-*spinning* BH while m_B and a_B refer to the lower-spinning BH, such that $a_A > a_B$. We use this approach because we would like to target the correlation between the mass and spin of the star that is spun-up. While models suggest that this should be the secondary BH, it is possible that the second mass-transfer phase can result in mass reversals. Additionally, many BBH systems have mass ratios close to unity, which makes it difficult to discern which is truly the secondary star. Furthermore, this approach is more agnostic towards the mechanism behind the mass-spin correlation and simply targets a correlation between the mass and spin of the higher-spinning star.

In all of our models, we use the source-frame masses. We use the default mass and redshift models (i.e. `Power Law + Peak` to model the masses, and a power law model for the redshift), which have 9 hyperparameters across both models. We simultaneously fit both models with a spin model. Due to the limited data quality and number of current detections, we only allow for a linear dependence of the spin on the masses. Assuming that the majority of the detected BBHs originate from the field, this should capture, to first order, the correlation described in the previous section. In the following subsections, we layout the different spin models which we employ, all of which model spin parameters as truncated Gaussians with log standard deviations and means allowed to vary linearly as a function of mass.

3.1 a_A spin model

The first model, which we call `a_A`, is astrophysically motivated and specifically targets the correlation between the mass and spin of the spun-up star described in [17]. We model a_A as a Gaussian truncated on $[0, 1]$ that is allowed to vary linearly with m_A, and model a_B as a Gaussian truncated on $[0, a_A]$ (hence enforcing the condition $a_A > a_B$) with a fixed width and standard deviation that we fit for:

$$\pi(a_A \,|\, \Lambda, m_A) = \mathcal{N}(a_A;\, \mu_A(m_A, \Lambda), 10^{\log \sigma_A(m_A, \Lambda)}, 0, 1) \tag{1}$$
$$\pi(a_B \,|\, \Lambda, a_A) = \mathcal{N}(a_B;\, \mu_B, 10^{\log \sigma_B}, 0, a_A) \tag{2}$$

where $\mathcal{N}(x;\, \mu, \sigma, c, d)$ is a Gaussian on x truncated on $[c, d]$ with mean μ and standard deviation σ, $\Lambda = [\mu_{A0}, \delta_{\mu,\, AA}, \log \sigma_{A0}, \delta_{\log \sigma,\, AA}, \mu_B, \log \sigma_B]$ are the 6 spin hyperparameters, and

$$\mu_A(m_A, \Lambda) = \mu_{A0} + \delta_{\mu,\, AA} \left(\frac{m_A}{10\, M_\odot} - 1 \right) \tag{3}$$
$$\log \sigma_A(m_A, \Lambda) = \log \sigma_{\mathrm{A0}} + \delta_{\log \sigma,\, AA} \left(\frac{m_A}{10\, M_\odot} - 1 \right). \tag{4}$$

Additionally, we enforce the following conditions on the truncated Gaussian:

$$\mathcal{N}(x;\, \mu, \sigma, c, d) = \begin{cases} 0 & \text{if } \sigma < 10^{-3} \text{ or } \mu \notin [-0.2, 1.2] \\ \mathcal{G}(\mu, \sigma, c, d) & \text{otherwise} \end{cases}, \tag{5}$$

where $\mathcal{G}(\mu, \sigma, c, d)$ is an ordinary Gaussian truncated on $[c, d]$. We enforce this cutoff to prevent infeasibly narrow Gaussians and models that peak significantly outside the bounds, which lead to numerical difficulties and have ambiguous physical interpretations.

3.2 χ_{eff} spin model

Because spin magnitudes are not measured precisely, we probe the same correlation using the effective dimensionless spin parameter χ_{eff}, which is measured with much greater certainty. If we assume, as in our astrophysical models, that only the spun-up star has significant spin, then χ_{eff} is a good proxy for a_A. This model, which we call `chieff`, models χ_{eff} as a Gaussian on $[-1, 1]$ that is again allowed to vary linearly with m_A:

$$\pi(\chi_{\mathrm{eff}} \,|\, \Lambda, m_A) = \mathcal{N}(\chi_{\mathrm{eff}};\, \mu(m_A, \Lambda), 10^{\log \sigma(m_A, \Lambda)}, -1, 1) \tag{6}$$
$$\mu(m_A, \Lambda) = \mu_0 + \delta_\mu \left(\frac{m_A}{10\, M_\odot} - 1 \right) \tag{7}$$
$$\log \sigma(m_A, \Lambda) = \log \sigma_0 + \delta_{\log \sigma} \left(\frac{m_A}{10\, M_\odot} - 1 \right). \tag{8}$$

We enforce a similar cutoff as in Equation 5, except we set $\mathcal{N}(\chi_{\mathrm{eff}}) = 0$ when $\mu \notin [-1, 1]$. This model has 4 spin hyperparameters, $\Lambda = [\mu_0, \delta_\mu, \log \sigma_0, \delta_{\log \sigma}]$.

3.3 Alternate χ_{eff} spin models

Finally, we explore two alternative models for χ_{eff}. The model is the same as the `chieff` model, but we instead allow χ_{eff} to vary with the total mass m_{tot} (`chieff_totalmass`) and

the primary mass m_1 (`chieff_m1`). The `chieff_m1` model is identical to one of the models explored in [7], and we include this model as a test of validation for our results.

The equations governing this model are identical to Equations 6-8, except the $\chi_{\rm eff}$ `_totalmass` has a pivot mass of $m_{\rm tot} = 20\,M_\odot$, such that

$$\mu(m_{\rm tot}, \Lambda) = \mu_0 + \delta_\mu \left(\frac{m_{\rm tot}}{20\,M_\odot} - 1 \right) \tag{9}$$

$$\log \sigma(m_{\rm tot}, \Lambda) = \log \sigma_0 + \delta_{\log \sigma} \left(\frac{m_{\rm tot}}{20\,M_\odot} - 1 \right). \tag{10}$$

4 Results for GWTC-3

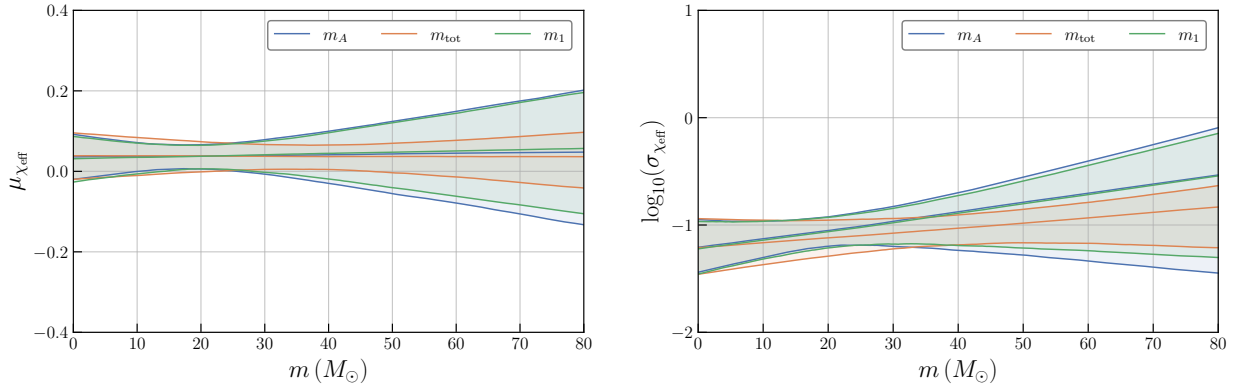

(a) (b)

Figure 1: The mean (a) and log standard deviation (b) of the $\chi_{\rm eff}$ truncated Gaussian for all three models, `chieff`, `chieff_totalmass`, and `chieff_mass1`, on GWTC-3 data as a function of m_A, $m_{\rm tot}$, and m_1, respectively. Solid lines show the median and 90% symmetric credible intervals.

We perform hierarchical inference on the 69 BBHs with FAR < 1 yr^{-1} [2]. We use the `C01:IMRPhenomXPHM` datasets from the publicly-released individual event posterior samples published in the GWTC-2.1 [25] and GWTC-3 [26] catalogs. We calculate the PE prior at the posterior sample points analytically, using the Jacobian developed by [8] to transform from a flat isotropic spin prior to a prior in $\chi_{\rm eff}$.

Figure 1 shows the mean and log standard deviation of the $\chi_{\rm eff}$ truncated Gaussian as a function of m_A, $m_{\rm tot}$, and m_1, from the `chieff`, `chieff_totalmass`, `chieff_mass1` models, respectively. Consistent with the results of [7], we find no evidence of $\chi_{\rm eff}$ varying as a function of mass for any of our models. The $\chi_{\rm eff}$ distribution likely peaks at a small, positive number, also consistent with the results of [2].

Figure 2 illustrates in the same format the results from the `a_A` model. Although we do not find a correlation between μ_{a_A} and m_A, as would be implied by the results of [17], we find that the a_A distribution appears to broaden with mass. We show this broadening explicitly

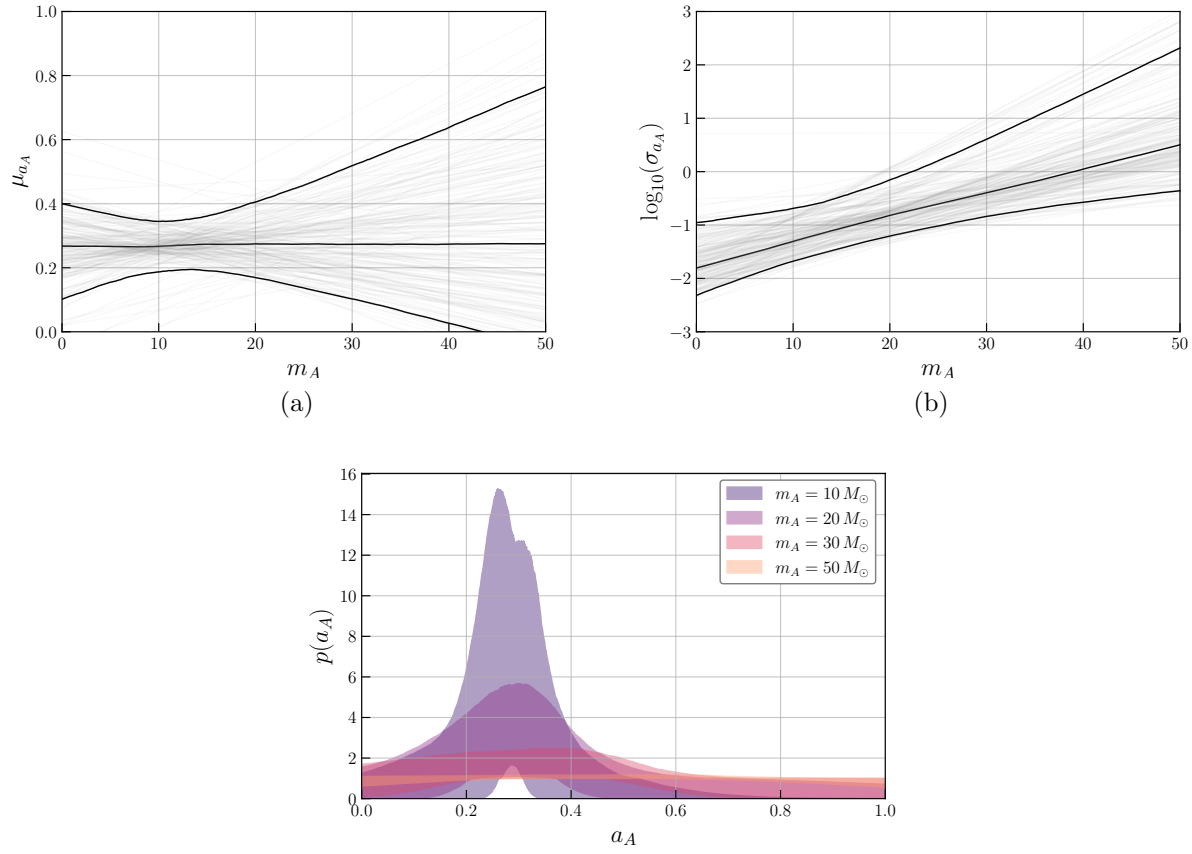

Figure 2: The results of the `a_A` model on GWTC-3 data, showing the mean (a) and log standard deviation (b) of the a_A truncated Gaussian as a function of m_A. Solid lines show the median and 90% symmetric credible intervals, while the faded gray lines show the trace plots of 200 random samples from the hyperposterior. In (c), we have plotted the recovered a_A distribution at different slices of m_A, showing the broadening of the distribution with mass.

in Figure 2(c). Further analysis is needed to show whether or not this result is caused by the enforcement of the condition $N_{\text{eff}} > N$.

5 Results for Injections

In order to validate our results on GWTC-3 data, we perform a simple injection study on the `a_A` and `chieff_mass1` models. We pick true values of the hyperparameters and draw $N = 1000$ events from the resulting distribution to make our mock catalog. We opt to leave out fitting the redshift distribution for simplicity, as it is independent of other components of the model. For the `a_A` model, after m_1 and m_2 are drawn from the `Power Law + Peak` mass model, we assign $m_1 = m_A$ (and $m_2 = m_B$) with a probability of $p = 20\%$, and

correspondingly $m_1 = m_B$ (and $m_2 = m_A$) with a probability of $1 - p = 80\%$. We opt to simulate perfect (i.e. no noise) observations with no selection biases. Therefore, the mock posteriors are composed of a single sample, which is the true value. We perform hierarchical inference with no selection function and a flat PE prior; we do not enforce $N_{\text{eff}} > N$, since there is only one posterior sample per event.

We are encountering separate difficulties for each model. For the `a_A` model injection study, the nested sampler does not converge, with $d \log z = \inf$ at each iteration before eventually stopping. This error began when I began enforcing the a_B Gaussian to be truncated on $[0, a_A]$ rather than $[0, 1]$ in the model. Figures 3(a-b) show the results when the a_B Gaussian was truncated on $[0, 1]$, in which case were able to correctly recover a null correlation between a_A and m_A. On the other hand, Figures 3(c-d) show the results of fitting injected data to the `chieff_m1` model; we were not able to correctly recover the true correlation of $\mu_{\chi_{\text{eff}}}$ with m_1 within a 90% confidence interval.

6 Next Steps

First, we need to resolve the difficulties we are encountering in our injection studies. To investigate the lack of convergence of the `a_A` model with injected data, I will look at the per event likelihood evaluation in the problematic regions of parameter space, as indicated by the trace plots. To investigate the reason behind the incorrect recovery of $\delta\mu$ in the `chieff_m1` model, I will re-run the inference with a lower mass power law index for more events with higher mass to probe whether or not the incorrect recovery is a result of lack of events at high mass.

The other immediate next step is to discuss our results with Jim Fuller, as our models were motivated primarily by his paper.

References

[1] B. P. Abbott et al. "Observation of Gravitational Waves from a Binary Black Hole Merger". In: Phys. Rev. Lett. 116.6, 061102 (Feb. 2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102. arXiv: 1602.03837 [gr-qc].

[2] R. Abbott et al. "Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3". In: *Physical Review X* 13.1, 011048 (Jan. 2023), p. 011048. DOI: 10.1103/PhysRevX.13.011048. arXiv: 2111.03634 [astro-ph.HE].

[3] Gregory Ashton et al. "BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy". In: ApJS 241.2, 27 (Apr. 2019), p. 27. DOI: 10.3847/1538-4365/ab06fc. arXiv: 1811.02042 [astro-ph.IM].

[4] Vishal Baibhav, Zoheyr Doctor, and Vicky Kalogera. "Dropping Anchor: Understanding the Populations of Binary Black Holes with Random and Aligned-spin Orientations". In: ApJ 946.1, 50 (Mar. 2023), p. 50. DOI: 10.3847/1538-4357/acbf4c. arXiv: 2212.12113 [astro-ph.HE].

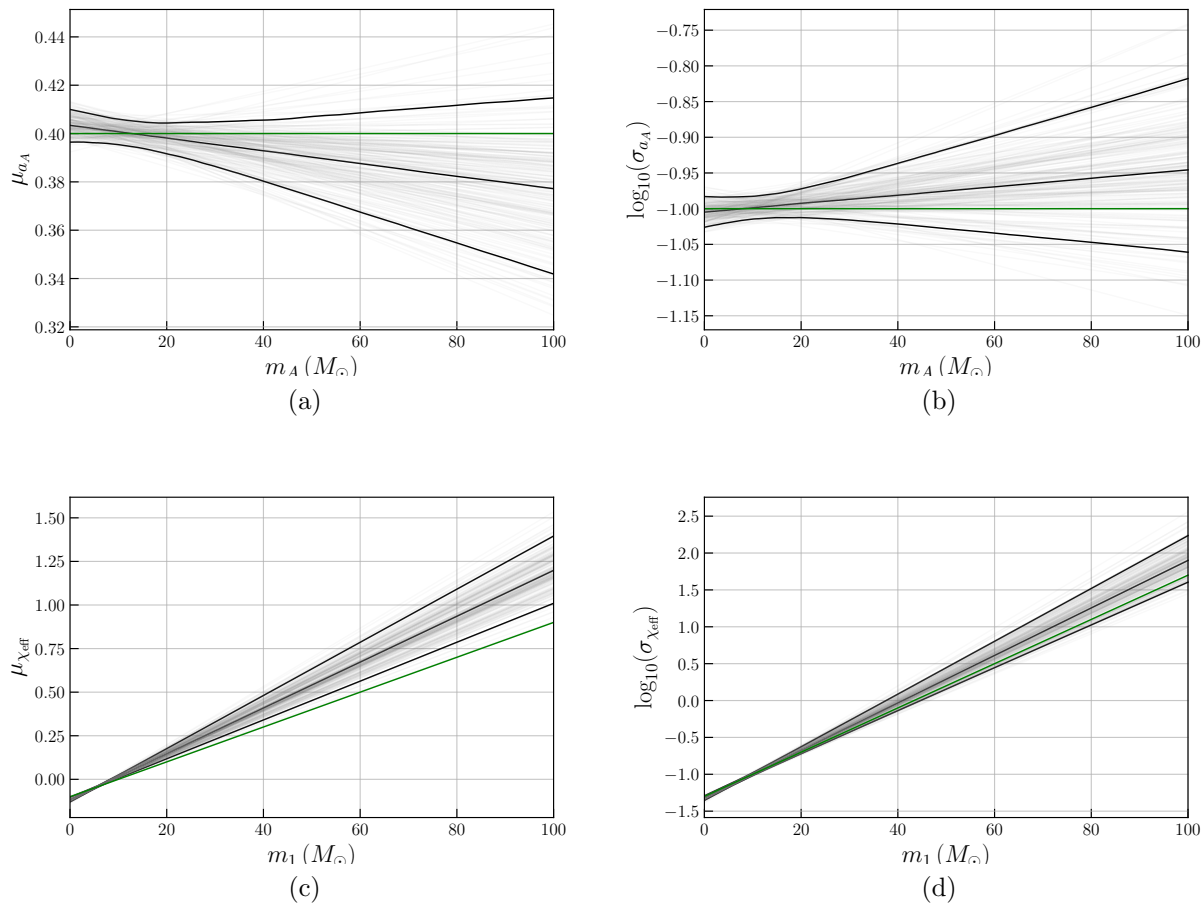

Figure 3: The results of the `a_A` (first row) and `chieff_totalmass` (second row) models on a mock catalog of $N = 1000$ events with no noise or selection biases, showing the mean (a, c) and log standard deviation (b, d) of the a_A and χ_{eff} truncated Gaussians as a function of m_A and m_1, respectively. Solid black lines show the median and 90% symmetric credible intervals, faded gray lines show the trace plots of 200 random samples from the hyperposterior, and the green line shows the true correlation.

[5] Simone S. Bavera, Michael Zevin, and Tassos Fragos. "Approximations of the Spin of Close Black Hole-Wolf-Rayet Binaries". In: *Research Notes of the American Astronomical Society* 5.5, 127 (May 2021), p. 127. DOI: 10.3847/2515-5172/ac053c. arXiv: 2105.09077 [astro-ph.HE].

[6] Sylvia Biscoveanu et al. "New Spin on LIGO-Virgo Binary Black Holes". In: Phys. Rev. Lett. 126.17, 171103 (Apr. 2021), p. 171103. DOI: 10.1103/PhysRevLett.126.171103. arXiv: 2007.09156 [astro-ph.HE].

[7] Sylvia Biscoveanu et al. "The Binary Black Hole Spin Distribution Likely Broadens with Redshift". In: ApJ 932.2, L19 (June 2022), p. L19. DOI: 10.3847/2041-8213/ac71a8. arXiv: 2204.01578 [astro-ph.HE].

[8] T. A. Callister. "A Thesaurus for Common Priors in Gravitational-Wave Astronomy". In: *arXiv e-prints*, arXiv:2104.09508 (Apr. 2021), arXiv:2104.09508. DOI: 10.48550/arXiv.2104.09508. arXiv: 2104.09508 [gr-qc].

[9] LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration. *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*. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) 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. The construction and operation of KAGRA are funded by Ministry of Education, Culture, Sports, Science and Technology (MEXT), and Japan Society for the Promotion of Science (JSPS), National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea, Academia Sinica (AS) and the Ministry of Science and Technology (MoST) in Taiwan. Zenodo, May 2023. DOI: 10.5281/zenodo.7890398. URL: https://doi.org/10.5281/zenodo.7890398.

[10] Will M. Farr. "Accuracy Requirements for Empirically Measured Selection Functions". In: *Research Notes of the American Astronomical Society* 3.5, 66 (May 2019), p. 66. DOI: 10.3847/2515-5172/ab1d5f. arXiv: 1904.10879 [astro-ph.IM].

[11] Maya Fishbach, Daniel E. Holz, and Will M. Farr. "Does the Black Hole Merger Rate Evolve with Redshift?" In: ApJ 863.2, L41 (Aug. 2018), p. L41. DOI: 10.3847/2041-8213/aad800. arXiv: 1805.10270 [astro-ph.HE].

[12] Maya Fishbach, Chase Kimball, and Vicky Kalogera. "Limits on Hierarchical Black Hole Mergers from the Most Negative χ_{eff} Systems". In: ApJ 935.2, L26 (Aug. 2022), p. L26. DOI: 10.3847/2041-8213/ac86c4. arXiv: 2207.02924 [astro-ph.HE].

[13] Gabriele Franciolini and Paolo Pani. "Searching for mass-spin correlations in the population of gravitational-wave events: The GWTC-3 case study". In: Phys. Rev. D 105.12, 123024 (June 2022), p. 123024. DOI: 10.1103/PhysRevD.105.123024. arXiv: 2201.13098 [astro-ph.HE].

[14] Gabriele Franciolini et al. "Searching for a subpopulation of primordial black holes in LIGO-Virgo gravitational-wave data". In: Phys. Rev. D 105.8, 083526 (Apr. 2022), p. 083526. DOI: 10.1103/PhysRevD.105.083526. arXiv: 2105.03349 [gr-qc].

[15] Davide Gerosa and Maya Fishbach. "Hierarchical mergers of stellar-mass black holes and their gravitational-wave signatures". In: *Nature Astronomy* 5 (July 2021), pp. 749–760. DOI: 10.1038/s41550-021-01398-w. arXiv: 2105.03439 [astro-ph.HE].

[16] Jaxen Godfrey, Bruce Edelman, and Ben Farr. "Cosmic Cousins: Identification of a Subpopulation of Binary Black Holes Consistent with Isolated Binary Evolution". In: *arXiv e-prints*, arXiv:2304.01288 (Apr. 2023), arXiv:2304.01288. DOI: 10.48550/arXiv.2304.01288. arXiv: 2304.01288 [astro-ph.HE].

[17] Linhao Ma and Jim Fuller. "Tidal Spin-up of Black Hole Progenitor Stars". In: *arXiv e-prints*, arXiv:2305.08356 (May 2023), arXiv:2305.08356. DOI: 10.48550/arXiv.2305.08356. arXiv: 2305.08356 [astro-ph.HE].

[18] Ethan Payne et al. "Searching for anisotropy in the distribution of binary black hole mergers". In: Phys. Rev. D 102.10, 102004 (Nov. 2020), p. 102004. DOI: 10.1103/PhysRevD.102.102004. arXiv: 2006.11957 [astro-ph.CO].

[19] April Qiu Cheng, Michael Zevin, and Salvatore Vitale. "What You Don't Know Can Hurt You: Use and Abuse of Astrophysical Models in Gravitational-wave Population Analyses". In: *arXiv e-prints*, arXiv:2307.03129 (July 2023), arXiv:2307.03129. DOI: 10.48550/arXiv.2307.03129. arXiv: 2307.03129 [astro-ph.HE].

[20] Joshua S. Speagle. "DYNESTY: a dynamic nested sampling package for estimating Bayesian posteriors and evidences". In: MNRAS 493.3 (Apr. 2020), pp. 3132–3158. DOI: 10.1093/mnras/staa278. arXiv: 1904.02180 [astro-ph.IM].

[21] Colm Talbot. "GWPopulation pipe". In: (Nov. 2021). DOI: 10.5281/zenodo.5654673. URL: https://git.ligo.org/RatesAndPopulations/gwpopulation_pipe.

[22] Colm Talbot and Eric Thrane. "Determining the population properties of spinning black holes". In: Phys. Rev. D 96.2, 023012 (July 2017), p. 023012. DOI: 10.1103/PhysRevD.96.023012. arXiv: 1704.08370 [astro-ph.HE].

[23] Colm Talbot and Eric Thrane. "Measuring the Binary Black Hole Mass Spectrum with an Astrophysically Motivated Parameterization". In: ApJ 856.2, 173 (Apr. 2018), p. 173. DOI: 10.3847/1538-4357/aab34c. arXiv: 1801.02699 [astro-ph.HE].

[24] Colm Talbot et al. "Parallelized inference for gravitational-wave astronomy". In: Phys. Rev. D 100.4, 043030 (Aug. 2019), p. 043030. DOI: 10.1103/PhysRevD.100.043030. arXiv: 1904.02863 [astro-ph.IM].

[25] The LIGO Scientific Collaboration et al. "GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run". In: *arXiv e-prints*, arXiv:2108.01045 (Aug. 2021), arXiv:2108.01045. DOI: 10.48550/arXiv.2108.01045. arXiv: 2108.01045 [gr-qc].

[26] The LIGO Scientific Collaboration et al. "GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run". In: *arXiv e-prints*, arXiv:2111.03606 (Nov. 2021), arXiv:2111.03606. DOI: 10.48550/arXiv.2111.03606. arXiv: 2111.03606 [gr-qc].

[27] Eric Thrane and Colm Talbot. "An introduction to Bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models". In: PASA 36, e010 (Mar. 2019), e010. DOI: 10.1017/pasa.2019.2. arXiv: 1809.02293 [astro-ph.IM].

[28] L. A. C. van Son et al. "No Peaks without Valleys: The Stable Mass Transfer Channel for Gravitational-wave Sources in Light of the Neutron Star-Black Hole Mass Gap". In: ApJ 940.2, 184 (Dec. 2022), p. 184. DOI: 10.3847/1538-4357/ac9b0a. arXiv: 2209.13609 [astro-ph.HE].

[29] Salvatore Vitale, Sylvia Biscoveanu, and Colm Talbot. "Spin it as you like: The (lack of a) measurement of the spin tilt distribution with LIGO-Virgo-KAGRA binary black holes". In: A&A 668, L2 (Dec. 2022), p. L2. DOI: 10.1051/0004-6361/202245084. arXiv: 2209.06978 [astro-ph.HE].

[30] Michael Zevin and Simone S. Bavera. "Suspicious Siblings: The Distribution of Mass and Spin across Component Black Holes in Isolated Binary Evolution". In: ApJ 933.1, 86 (July 2022), p. 86. DOI: 10.3847/1538-4357/ac6f5d. arXiv: 2203.02515 [astro-ph.HE].

[31] Michael Zevin et al. "One Channel to Rule Them All? Constraining the Origins of Binary Black Holes Using Multiple Formation Pathways". In: ApJ 910.2, 152 (Apr. 2021), p. 152. DOI: 10.3847/1538-4357/abe40e. arXiv: 2011.10057 [astro-ph.HE].