




Who Ordered That? Unequal-mass Binary Black Hole Mergers Have Larger Effective Spins

Thomas A. Callister¹ , Carl-Johan Haster^{2,3} , Ken K. Y. Ng^{2,3} , Salvatore Vitale^{2,3} , and Will M. Farr^{1,4} 

¹ Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA; tcallister@flatironinstitute.org, will.farr@stonybrook.edu

² LIGO Laboratory, Massachusetts Institute of Technology, 185 Albany St., Cambridge, MA 02139, USA; haster@mit.edu, kenkyng@mit.edu, salvo@mit.edu,

³ Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139, USA

⁴ Department of Physics and Astronomy, Stony Brook University, Stony Brook NY 11794, USA

Received 2021 June 17; revised 2021 September 28; accepted 2021 October 2; published 2021 November 16

Abstract

Hierarchical analysis of binary black hole (BBH) detections by the Advanced LIGO and Virgo detectors has offered an increasingly clear picture of their mass, spin, and redshift distributions. Fully understanding the formation and evolution of BBH mergers will require not just the characterization of these marginal distributions, but the discovery of any correlations that exist between the properties of BBHs. Here, we hierarchically analyze the ensemble of BBHs discovered by LIGO and Virgo with a model that allows for intrinsic correlations between their mass ratios q and effective inspiral spins χ_{eff} . At 98.7% credibility, we find that the mean of the χ_{eff} distribution varies as a function of q , such that more unequal-mass BBHs exhibit systematically larger χ_{eff} . We find a Bayesian odds ratio of 10.5 in favor of a model that allows for such a correlation over one that does not. Finally, we use simulated signals to verify that our results are robust against degeneracies in the measurements of q and χ_{eff} for individual events. While many proposed astrophysical formation channels predict some degree correlation between spins and mass ratio, these predicted correlations typically act in an opposite sense to the trend we observationally identify in the data.

Unified Astronomy Thesaurus concepts: Compact binary stars (283); Astrostatistics (1882); Gravitational wave astronomy (675); Astrophysical black holes (98); Stellar mass black holes (1611); LIGO (920)

1. Introduction

The growing number of gravitational-wave detections made by the Advanced LIGO (LIGO Scientific Collaboration 2015) and Advanced Virgo (Virgo Collaboration 2015) observatories is enabling exploration of the stellar-mass compact binary population at an ever accelerating pace. With data now available from the first three LIGO–Virgo observing runs (LIGO Scientific Collaboration & Virgo Collaboration 2019a, 2021a), we are beginning to resolve interesting features in the mass, spin, and redshift distributions of binary black holes (BBHs; LIGO Scientific Collaboration & Virgo Collaboration 2021b; Roulet et al. 2020). The BBH primary mass spectrum is characterized by a power law at low masses and a possible “bump” near $40 M_{\odot}$, followed by a steeper decline and a possible secondary feature near $80 M_{\odot}$ (Fishbach & Holz 2017; Wysocki et al. 2019; Kimball et al. 2021; Roulet et al. 2020; Tiwari & Fairhurst 2021; LIGO Scientific Collaboration & Virgo Collaboration 2021b). Black hole spins appear to be small but non-zero, and are oriented neither isotropically nor strictly parallel to the binaries’ orbits but with some spread in spin–orbit tilt angles (Farr et al. 2017, 2018; Tiwari et al. 2018; Roulet & Zaldarriaga 2019; Wysocki et al. 2019; Biscoveanu et al. 2021; Miller et al. 2020). Meanwhile, the BBH merger rate likely increases with redshift at a rate comparable to cosmic star formation (Fishbach et al. 2018; Callister et al. 2020a; LIGO Scientific Collaboration & Virgo Collaboration 2021b; LIGO Scientific Collaboration, Virgo Collaboration & KAGRA Collaboration 2021).

The distributions of black hole parameters encode a valuable range of astrophysical information, and may elucidate processes governing compact binary birth and evolution. In addition to features in these one-dimensional distributions, *correlations* between parameters will be particularly valuable to identify and understand.

Different proposed formation channels, for instance, predict a variety of distinctive correlations that may exist among compact binaries. One might, for example, expect a correlation between BBH mass and spin if black holes experience repeated hierarchical mergers in dense stellar environments (Portegies Zwart & McMillan 2002; McKernan et al. 2012; Antonini & Rasio 2016; Fishbach et al. 2017; Gerosa & Berti 2017; Rodriguez et al. 2019; Doctor et al. 2020; Kimball et al. 2021). And the strength of tidal interactions among field binaries might conceivably regulate BH spins in a way that depends on their mass ratio (Hotokezaka & Piran 2017; Gerosa et al. 2018; Qin et al. 2018; Zaldarriaga et al. 2018; Bavera et al. 2020, 2021).

Here, we report an apparent anti-correlation between the mass ratio and effective inspiral spins of binary black hole mergers: more extreme mass ratios correspond to larger effective spins. At mass ratios near unity, the effective spins of BBHs are consistent with a narrow distribution that is symmetric about zero. At unequal mass ratios, however, the BBH population exhibits preferentially positive effective spins due to an overall shift of the effective spin distribution toward larger values. This behavior is generally inconsistent with predictions of the standard models of compact object formation discussed above.

2. Residual Structure in the Mass Ratio and Effective Spin Plane

The LIGO & Virgo second gravitational-wave transient catalog (GWTC-2) contains 44 BBH candidates with false alarm rates below one per year (LIGO Scientific Collaboration & Virgo Collaboration 2021a). For each of these 44 events, Figure 1(a) shows the joint posterior probability on its mass

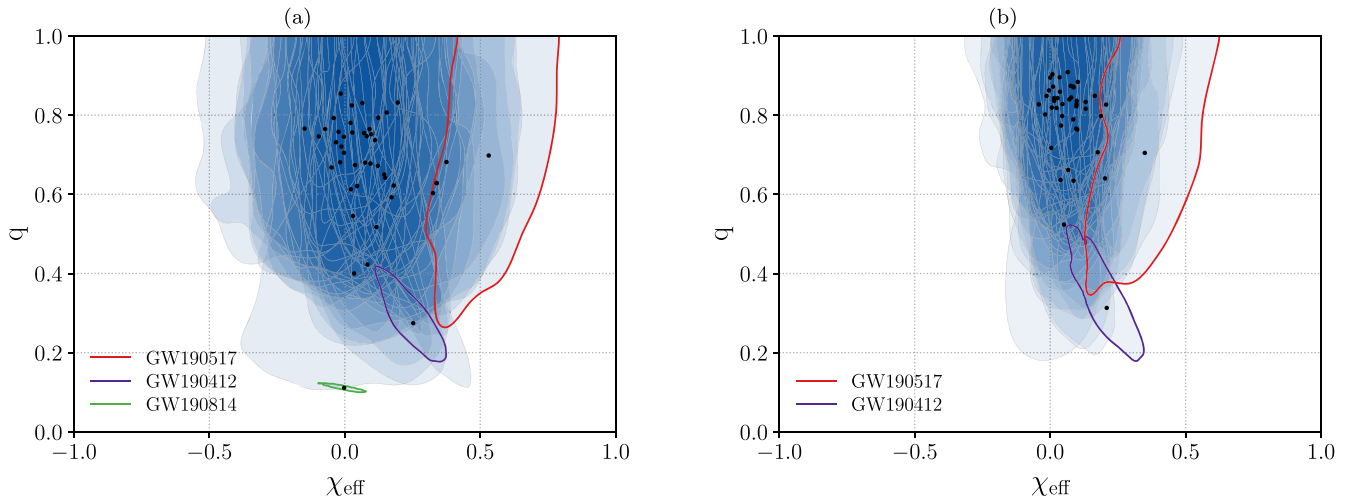


Figure 1. (a): Effective spin and mass ratio posteriors for the 45 BBH candidates in the LIGO/Virgo GWTC-2 catalog (LIGO Scientific Collaboration & Virgo Collaboration 2021a) with false alarm rates below 1 yr^{-1} , as obtained under a default prior. Each shaded region gives the central 90% credible posterior bounds for a given BBH, and black points mark the median χ_{eff} and median q values for each event. Three individually notable events are highlighted. GW190412 has a precisely measured mass ratio constrained well away from unity, while GW190517 likely possesses a large, positive χ_{eff} . Finally, GW190814 exhibits an extreme mass ratio, with a secondary that may either be a massive neutron star or a very light black hole. Due to the unknown nature of GW190814, this event is excluded from our analysis unless otherwise stated. (b): Posteriors for the 44 confident BBHs (excluding GW190814) under a new population-informed prior, obtained by hierarchically fitting the BBH population assuming a Gaussian distribution of effective spins (Equation (4)) and a power-law mass ratio distribution [Equation (3)]. Under a population-informed prior, the ensemble of posteriors has contracted to favor smaller values of χ_{eff} and mass ratios nearer $q \sim 1$.

ratio

$$q = \frac{m_2}{m_1} \quad (1)$$

and its effective spin (Damour 2001)

$$\chi_{\text{eff}} = \frac{\chi_1 \cos t_1 + q \chi_2 \cos t_2}{1 + q}, \quad (2)$$

which quantifies the mass-weighted average of the two component spins when projected parallel to the binary’s orbital angular momentum. Here, m_1 and $m_2 \leq m_1$ are the primary and secondary masses of the component black holes, χ_1 and χ_2 are the dimensionless component spin magnitudes, and t_1 and t_2 are the angles made by each component spin relative to the binary orbital angular momentum. Each contour bounds the central 90% credible region for the given event, under default parameter estimation priors (see Appendix A), and black points mark the one-dimensional median χ_{eff} and q estimates for each event.

Most events in the $\chi_{\text{eff}} - q$ plane are consistent with mass ratios near unity and have effective spins clustered around $\chi_{\text{eff}} \sim 0$. A smaller number, including the events GW190517 and GW190412, appear to have effective spins constrained away from zero. The event GW190412 is also the first BBH to have a confidently unequal mass ratio, with $q \sim 0.3$ (LIGO Scientific Collaboration & Virgo Collaboration 2020a). In Figure 1(a) we have additionally highlighted the event GW190814. Like GW190412, GW190814 has a confidently unequal mass ratio, but its physical nature is unknown; its secondary mass $m_2 \approx 2.5 M_{\odot}$ may be either a very light black hole or an extraordinarily heavy neutron star (LIGO Scientific Collaboration & Virgo Collaboration 2020b). If GW190814 is presumed to be a BBH, it nevertheless remains an outlier relative to the broader BBH population (LIGO Scientific Collaboration & Virgo Collaboration 2021b). We will therefore neglect GW190814 in our analysis below unless stated otherwise (we return to this event in Section 4). Also visible in

Figure 1(a) is the well-known measurement degeneracy between χ_{eff} and q , giving rise to the extended contours that curve down toward low q and large χ_{eff} (Baird et al. 2013; Ohme et al. 2013; Pürrer et al. 2013, 2016; Ng et al. 2018; Tiwari et al. 2018).

In addition to these events comprising GWTC-2, independent reanalyses of Advanced LIGO and Virgo data have identified several additional BBH candidates (Venumadhav et al. 2020; Nitz et al. 2021b). In our analysis below, a critical ingredient is knowledge of the selection effects governing searches for BBH events. The LIGO and Virgo collaborations have made available a set of software injections that we will use to precisely quantify these selection effects (LIGO Scientific Collaboration & Virgo Collaboration 2021b, 2020c); see Appendix A. In order to ensure self-consistency with our sample of BBHs, we will continue to use only those binaries among GWTC-2. While this paper was in preparation, an updated GWTC-2 catalog (GWTC-2.1) was released with the addition of several low-significance BBH candidates, including two consistent with mass ratios $q \lesssim 0.6$ (LIGO Scientific Collaboration & Virgo Collaboration 2021c); we neglect these new candidates.

From Figure 1(a) alone, it is difficult to draw any conclusions about the underlying *population* of BBHs; the data in Figure 1(a) are subject to selection effects and considerable measurement uncertainties, both of which confound any information that might be gleaned by eye about underlying features in this plane. As a first step, we can refine Figure 1(a) by invoking a simple model for the BBH population. The posteriors in Figure 1(a) are obtained via parameter estimation with broad, uninformative priors on χ_{eff} and q . By hierarchically measuring the BBH population, we can reweight each posterior to a new population-informed prior, leveraging the ensemble of events to help us more accurately identify the properties of any one individual system. For the time being, we will assume that mass ratios and effective spins are *uncorrelated*, describing the population distribution of mass ratios via a power law,

$$p(q|m_1, \gamma) \propto q^{\gamma}, \quad (3)$$

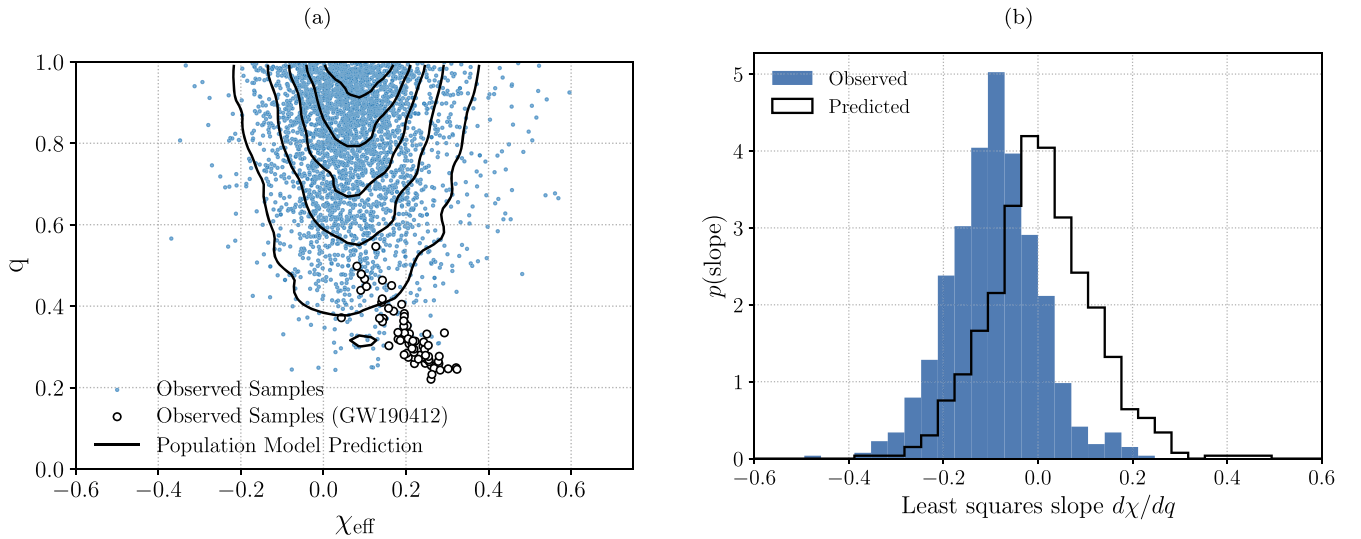


Figure 2. (a): Locations of expected and observed detections in the $\chi_{\text{eff}} - q$ plane. Black contours show the expected density of future BBH detections (bounding 10%, 30%, 50%, 70%, and 90% of expectations), given hierarchical inference on GWTC-2 under the assumption that q and χ_{eff} are uncorrelated. Blue dots, meanwhile, show posterior draws (75 per event) from our set of BBHs in GWTC-2, reweighted to a population-informed prior as in Figure 1(b). Samples drawn from GW190412 are plotted separately as empty circles so that they can be better differentiated from the bulk population. The empty histogram shows the distribution of least-squares slopes $d\chi/dq$ under repeated draws of 44 samples from the *expected* distribution of detections according to our population model. The blue histogram, meanwhile, shows the corresponding distribution of slopes under repeated draws from GWTC-2 posteriors (one sample per event). While our base population model predicts slopes centered about zero, draws from event posteriors yield a distribution of slopes offset noticeably toward negative values. This tension is further explored in Section 3, in which we expand our population model to directly parameterize and measure any correlations between q and χ_{eff} .

with $m_{\text{min}}/m_1 \leq q \leq 1$, and the distribution effective spins as a Gaussian (Roulet & Zaldarriaga 2019; Miller et al. 2020),

$$p(\chi_{\text{eff}}|\mu_\chi, \sigma_\chi) \propto \exp\left[-\frac{(\chi_{\text{eff}} - \mu_\chi)^2}{2\sigma_\chi^2}\right], \quad (4)$$

truncated on the interval $-1 \leq \chi_{\text{eff}} \leq 1$. We obtain posteriors on the parameters governing these distributions using the EMCEE Markov Chain Monte Carlo sampler (Foreman-Mackey et al. 2013). Alongside q and χ_{eff} , we also hierarchically measure the ensemble distribution of primary masses and BBH redshifts. We assume that primary masses are distributed as a power law with a possible Gaussian peak, and a merger rate per comoving volume that evolves as $(1+z)^k$ with redshift (Fishbach et al. 2018; Talbot & Thrane 2018; LIGO Scientific Collaboration & Virgo Collaboration 2019b, 2021b); these models are described in Appendix A, along with details of our hierarchical inference. Our resulting posteriors on μ_χ , σ_χ , and γ are shown in Appendix C.

With this initial fit to the BBH population, we can update our measurements of χ_{eff} and q for each BBH. Figure 1(b) illustrates these reweighted posteriors. The incorporation of a population-informed prior yields two major effects. First, all posteriors have contracted toward small χ_{eff} since we infer both the mean and standard deviation of the effective spin distribution to be small, with $\mu_\chi = 0.06^{+0.05}_{-0.05}$ and $\sigma_\chi = 0.12^{+0.05}_{-0.05}$. Second, since $\gamma = 1.1^{+3.0}_{-1.4}$ is inferred to be positive (favoring mass ratios near unity), many sources have also shifted upward toward $q \sim 1$.

We can now explore whether this baseline population model, with *uncorrelated* spin and mass ratio distributions, is a reasonably good fit to observation. The black contours in Figure 2(a) illustrate the *expected* density of detections in the $\chi_{\text{eff}} - q$ plane, assuming our default population model is correct. These contours are obtained using a set of simulated

BBH signals injected into LIGO and Virgo data, reweighting the successfully recovered injections to our baseline population model (see Appendix A). The expected detections are clustered at $q \sim 1$, with χ_{eff} values spread symmetrically about the population mean at $\chi_{\text{eff}} \approx 0.05$. The collection of points, meanwhile, shows 75 posterior draws from each of our 44 BBHs, also reweighted to a population-informed prior as in Figure 1(b). To better differentiate between samples drawn from the bulk population and those from the low- q event GW190412, GW190412’s samples are shown as empty circles and all other posterior samples as blue dots.

To quantify the degree of tension (if any) that may exist between observed and predicted samples in Figure 2(a), we repeatedly generate and compare catalogs of $\chi_{\text{eff}} - q$ samples consistent with GWTC-2 against predicted catalogs of mock observations drawn from our baseline population model. We begin by choosing a random sample $\Lambda = \{\gamma, \mu_\chi, \sigma_\chi, \dots\}$ drawn from our posterior on the population-level parameters. Given this value of Λ , we reweight each BBH’s posterior to the corresponding population and randomly draw a single posterior sample $\{\chi_{\text{eff}}, q\}$ from every reweighted posterior to yield a catalog of 44 “observed” values consistent with GWTC-2. Under this same proposed population Λ , we similarly draw a “predicted” catalog of mock observations, reweighting and drawing 44 events from the set of successfully found pipeline injections. For both the “predicted” and “observed” catalogs we can then compute a simple least-squares slope $d\chi/dq$ of the 44 samples in the $\chi_{\text{eff}} - q$ plane. The white and blue histograms in Figure 2(b) illustrate the distributions of these least-squares slopes, taken over many random draws of Λ . If our baseline population model, in which χ_{eff} and q are uncorrelated, were a good descriptor of GWTC-2, then draws from our model should predict slopes consistent with observation. The distributions in Figure 2(b), though, exhibit a systematic offset from one another: whereas the baseline population model predicts slopes centered at zero (as it must, if χ_{eff} and q are presumed

independent), GWTC-2 observations preferentially yield negative slopes centered above $\langle d\chi/dq \rangle = -0.1$, with observed slopes lying below predicted ones 74% of the time.

3. Measuring a Correlation between Effective Spin and Mass Ratio

The baseline population model adopted in Section 2, with a power law in mass ratio (Equation (3)) and an uncorrelated Gaussian in effective spins (Equation (4)), is possibly inadequate in capturing the full range of structure exhibited by BBHs in the $\chi_{\text{eff}} - q$ plane. As explored in Figure 2, GWTC-2 is suggestive of a tendency toward larger χ_{eff} with smaller q , a trend that cannot be captured with the simple population model employed so far. Motivated by this tension, in this section we will expand our initial population model to allow for a *correlation* between effective spins and mass ratios and check whether the data are informative or agnostic about the existence of such a correlation.

We continue to describe the mass ratio distribution via Equation (3), but modify our χ_{eff} model such that its mean μ_χ and log standard deviation $\log_{10} \sigma_\chi$ are now allowed to evolve linearly with q :

$$p(\chi_{\text{eff}}|q, \mu_{\chi,0}, \sigma_{\chi,0}, \alpha, \beta) \propto \exp \left[-\frac{(\chi_{\text{eff}} - \mu_\chi(\mu_{\chi,0}, \alpha, q))^2}{2\sigma_\chi^2(\sigma_{\chi,0}, \beta, q)} \right], \quad (5)$$

with

$$\mu_\chi(\mu_{\chi,0}, \alpha, q) = \mu_{\chi,0} + \alpha(q - 0.5) \quad (6)$$

and

$$\log_{10} \sigma_\chi(\sigma_{\chi,0}, \beta, q) = \log_{10} \sigma_{\chi,0} + \beta(q - 0.5). \quad (7)$$

The parameters α and β govern the degree of evolution of χ_{eff} with mass ratio; confidently constraining either parameter to be non-zero would indicate that the BBH effective spin distribution shifts in location or width with increasing q . Note that we choose our linear expansion in the log standard deviation rather than the standard deviation itself, in accordance with our log-uniform prior on σ_χ and $\sigma_{\chi,0}$ in Equations (4) and (7).

We repeat our hierarchical analysis of the BBH events in GWTC-2, now additionally measuring the slope parameters α and β . We again fit simultaneously for the primary mass and redshift distributions of BBHs, as described in Appendix A. Figure 3 shows our resulting posteriors on the subset of parameters governing the spin and mass ratio distributions. At 98.7% credibility, α is constrained to be less than zero, indicating that the effective spin distribution among BBHs in GWTC-2 is *anti-correlated* with mass ratio, shifting toward larger χ_{eff} at lower q .

Whereas we confidently conclude that the mean of the χ_{eff} distribution evolves with q , we can say little about changes in the *width* of the χ_{eff} distribution. Although we rule out very large values of β , this parameter is still permitted to be moderately positive (a broadening χ_{eff} distribution toward large q), very negative (a narrowing distribution), or zero. Note also that the lower boundary on β as shown in Figure 3 is set by our prior and not the data.

As another measure of significance, we compute a Bayes factor between our initial model in Section 2, with independent χ_{eff} and q distributions, and our expanded model that allows for correlations between q and χ_{eff} , with possibly non-zero α

and β . Using the DYNesty nested sampler (Speagle 2020) to compute Bayesian evidences for these two cases, we find a Bayes factor of 10.5 in favor of the expanded model allowing for correlations (see Table 1, discussed further below). We caution, though, that Bayes factors can be difficult to interpret due to their dependence on prior volume and the inability to capture any “trials factors” due to the many exploratory population analyses performed among the community.

Figure 4 illustrates our posterior on $\mu_\chi(q)$ and $\log_{10} \sigma_\chi(q)$ as a function of mass ratio. In each subplot, black lines denote the median and central 90% credible bounds on $\mu_\chi(q)$ and $\log_{10} \sigma_\chi(q)$, while light blue traces show the result of individual draws from our posterior on $\{\mu_{\chi,0}, \log_{10} \sigma_{\chi,0}, \alpha, \beta\}$. Again, μ_χ confidently exhibits evolution with mass ratio, with the 90% credible lower bound on μ_χ at $q=0$ constrained to sit above the 90% credible *upper* bound at $q=1$. Our best measurement of $\mu_\chi(q)$ occurs at $q \approx 0.8$; the fact that we center Equation (6) at $q=0.5$ and not 0.8 is the source of the degeneracy between $\mu_{\chi,0}$ and α seen in Figure 3 above. In contrast, no evolution is evident for $\log_{10} \sigma_\chi(q)$: a horizontal line can be fit inside the 90% credible bounds shown, consistent with the fact that we do not exclude $\beta=0$. Note that, since our posterior on β is bounded by our prior, a non-negligible fraction of the space in Figure 4 is excluded a priori. The shaded gray area shows this region that is artificially excluded by our priors on $\log_{10} \sigma_0$ and β .

Reweighting the posterior of each BBH via an updated population prior from our correlated $\chi_{\text{eff}} - q$ model yields the results shown in Figure 5. Relative to the posteriors obtained via standard parameter estimation (Figure 1(a)) and those obtained by reweighting to a standard uncorrelated population model (Figure 1(b)), the reweighted posteriors in Figure 5 exhibit a significant degree of “shrinkage,” shifting to lie predominantly along the “best-fit” line characterized by $\alpha \approx -0.5$. The event GW190517, in particular, changes character considerably under our updated population prior. Under standard parameter estimation priors, GW190517 is a possible outlier in spin, with the largest χ_{eff} measured among GWTC-2, but otherwise is consistent with a fairly typical mass ratio (LIGO Scientific Collaboration & Virgo Collaboration 2021a). Under a population model that allows for correlated mass ratios and spins, however, the posterior of GW190517 shifts to favor significantly lower χ_{eff} and q , identifying GW190517 as another confidently unequal mass event like GW190412. Moreover, neither GW190517 nor GW190412 are clear outliers relative to the rest of the BBH population; both are quite consistent with the same linear trend favored by the other BBH observations.

Although the BBHs in GWTC-2 favor a χ_{eff} distribution that evolves mass ratio, does this expanded model actually resolve the predictive tension discussed in Figure 2? We repeat the predictive exercise discussed in Section 2, repeatedly drawing a hyperparameter sample $\Lambda = \{\gamma, \mu_{\chi,0}, \sigma_{\chi,0}, \alpha, \beta, \dots\}$, generating “observed” catalogs from our reweighted posteriors and “predicted” catalogs from the population model, and recording their least-squares slopes. The resulting distributions are shown in Figure 6. The expectation values over these distributions are now in far better agreement, with the “observed” and “predicted” slopes having sample means of $\langle d\chi/dq \rangle = -0.46$ and -0.49 , respectively.

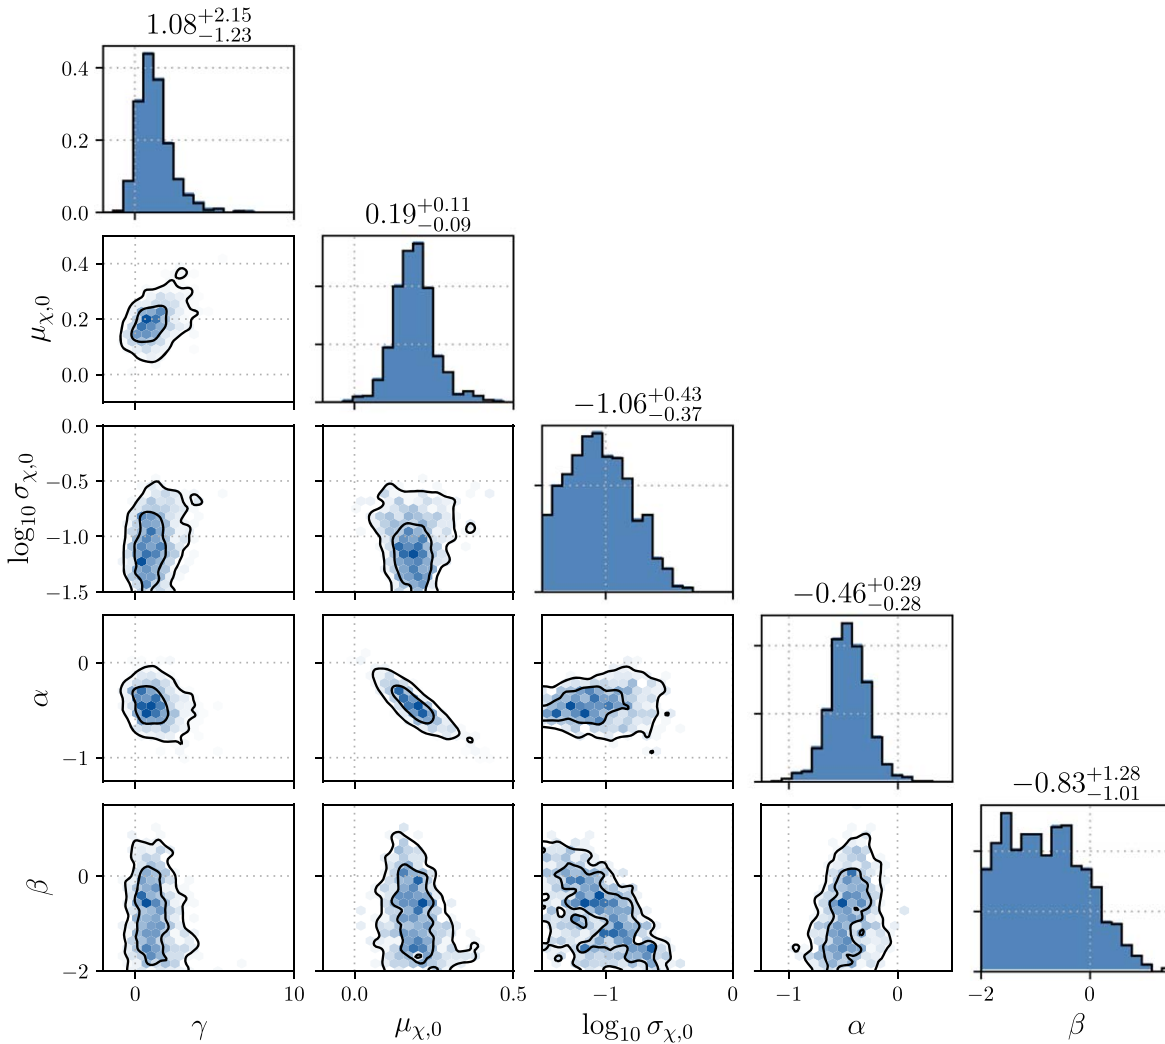


Figure 3. Posterior distributions on the parameters governing the mass ratio and effective spin distributions of BBHs, under a model in which the mean and standard deviation of the χ_{eff} distribution are correlated with q via the parameters α and β [see Equation (5)]. The measurements appearing above each one-dimensional posterior correspond to median estimates and central 90% credible uncertainties. At 98.7% credibility, α is inferred to be negative, indicating a preference for an *anti-correlation* between the effective spins and mass ratios among BBHs in GWTC-2. Posteriors on the remaining hyperparameters used to describe the primary mass and redshift distribution are shown in Appendix C.

Table 1

Relative Bayesian Evidences for the Four Model Variants Discussed in Sections 2, 3, and 5

Model	ln Evidence
Correlated χ_{eff} & q ,	3.65 (± 0.03)
Correlated χ_{eff} & q , with $\chi_{\text{eff}} \geq 0$	3.54 (± 0.03)
Independent χ_{eff} & q ,	1.29 (± 0.04)
Independent χ_{eff} & q , with $\chi_{\text{eff}} \geq 0$	0.00 (± 0.04)

Note. These models collectively encompass the possibilities that q and χ_{eff} are [Equation (5)] or are not [Equation (4)] correlated, together with the existence or non-existence of BBHs with negative effective spins. Evidences are obtained using DYNESTY (Speagle 2020), and the stated uncertainties are obtained by computing each evidence 10 times and recording the standard deviation among runs. We choose a normalization such that the most disfavored model (uncorrelated parameters with purely positive χ_{eff}) has a mean log-evidence of zero.

4. Spurious Sources of Apparent Correlation?

An anti-correlation of BBH effective spins with mass ratio is surprising, given current theories of compact binary formation and evolution. In this section, we therefore detail a number of

checks to bolster our confidence that this anti-correlation is a real observed feature of our data rather than a spurious effect due to uncontrolled systematics.

First, we explore whether our confident measurement of $\alpha < 0$ can be traced to any one event. If so, this would not necessarily indicate that our results are spurious, but would suggest that the event in question is an outlier among the BBH population, or possibly that an error or bias is present in the data associated with this event. We repeat our hierarchical inference of the correlated $\chi_{\text{eff}} - q$ model in Section 3 three additional times, excluding GW190412, excluding GW190517, and excluding both events together. The resulting marginalized posteriors on α are shown in Figure 7. The observed anti-correlation between χ_{eff} and q remains robust against the exclusion of either event. Excluding GW190517 leaves our result effectively unchanged, with $\alpha < 0$ at 99.1% credibility. The same is true of GW190412, whose exclusion gives $\alpha < 0$ at 99.0% credibility. Finally, while excluding *both* GW190412 and GW190517 somewhat weakens our conclusions, in this case, α is still constrained to be less than zero at 95.2% credibility. Therefore, although these two events notably aid in the precision with which α is measured, the




Figure 4. Constraints on the mean $\mu_{\chi}(q)$ and standard deviation $\sigma_{\chi}(q)$ of the χ_{eff} distribution, as a function of BBH mass ratio q . Each blue trace represents a single draw from our posterior on the population-level parameters from Figure 3, while black lines mark the mean and central 90% credible bounds on $\mu_{\chi}(q)$ and $\sigma_{\chi}(q)$ at a given mass ratio. The mean of the χ_{eff} distribution is constrained to decrease with increasing q , while the data are consistent with a non-evolving width. In the lower panel, the shaded gray region denotes the region artificially excluded by our prior bounds on $\log \sigma_0$ and β [see Equation (7)].

preference for negative α is a feature of the broader BBH population. We note that, when GW190412 is excluded, *more negative* values of α are permitted by the data, suggesting that, rather than driving the increase in χ_{eff} with lower q , the tight parameter estimates for GW190412 appear instead to anchor any evolution to somewhat shallower values. This behavior is further explored in Appendix B, where we quantify which events most strongly prefer or most strongly resist an anti-correlation between q and χ_{eff} .

If, on the other hand, we *add in* GW190814, our conclusions do not change qualitatively. Due to its unknown source classification, we have so far excluded GW190814 from the set of BBHs informing our hierarchical inference (LIGO Scientific Collaboration & Virgo Collaboration 2020b, 2021b). If we repeat our analysis but now include GW190814 among our sample, we obtain the reweighted population shown in Figure 8(a), while Figure 8(b) shows the updated posteriors on α and β when GW190814 is present. Given the very precise measurements of GW190814’s mass ratio and effective spin, this event does not shift noticeably under a population-informed prior. Its presence, though, does not override the tendency of the other BBHs to prefer anti-correlated χ_{eff} and q . We still confidently infer $\alpha < 0$, although the strength of this measurement, now at 92.6% credibility, is reduced.

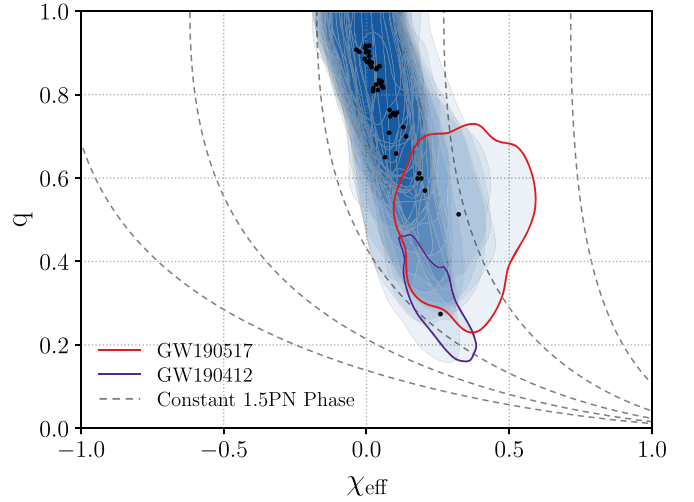


Figure 5. Effective spin and mass ratio posteriors for BBHs among GWTC-2, reweighted to a new population-informed prior given by our hierarchical inference of the (possibly correlated) BBH q and χ_{eff} distributions. Under a model that allows for linear correlations between q and the mean and width of the effective spin distribution, single-event posteriors shift to preferentially lie along a “best-fit line” in the $\chi_{\text{eff}} - q$ plane. Notably, the high-spin event GW190517 has shifted to more unequal mass ratios, and is now constrained to the same neighborhood as the unequal-mass event GW190412. If BBH mass ratio and spins are correlated, it may be the case that GW190412 and GW190517 are two representatives of a shared class of BBH. Also shown are contours (dashed gray lines) of the 1.5PN coefficient governing the phase evolution of a gravitational-wave signal; measurement degeneracies between χ_{eff} and q lie preferentially along these contours.

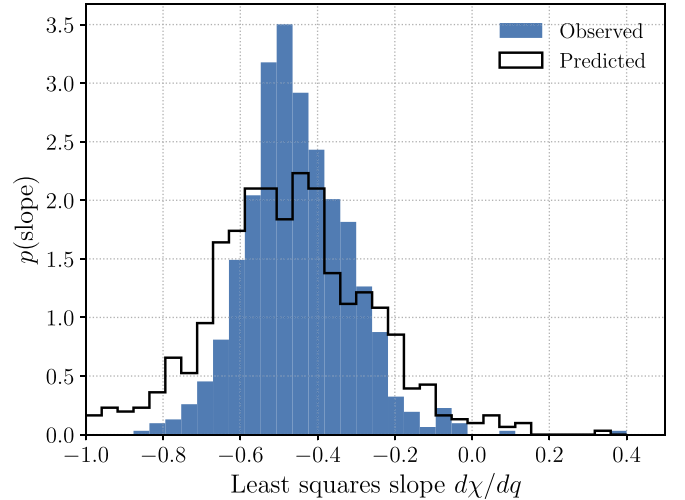


Figure 6. As in Figure 2(b), a comparison of least-squares slopes $d\chi/dq$ of describing draws from the reweighted posteriors of observed BBHs to those slopes predicted by our population model that allows for $\chi_{\text{eff}} - q$ correlations. As shown in Figure 2, a standard population model in which q and χ_{eff} are uncorrelated yields a systematic shift between observed and predicted distributions of slopes. An expanded model that allows for correlations resolves this tension, yielding observed and predicted distributions centered at the same values.

GW190814’s primary effect, however, is to *broaden* the χ_{eff} distribution toward lower q ; we now infer that $\beta < 0$ at 96.8% credibility (although the exact significance will depend on our lower prior bound on β). Qualitatively, since all other BBHs force $\alpha < 0$, the only way for our population model to simultaneously accommodate GW190814 is also to broaden at very low q . This effect can be seen when comparing the reweighted posteriors in Figures 5 and 8(a); whereas the median measurements (black




Figure 7. Posteriors on the parameter α governing the evolution of μ_χ with mass ratio when our hierarchical inference is repeated with GW190412 and/or GW190517 left out of our sample. When one or the other event is excluded, α remains confidently negative. This holds true even when *both* events are excluded, although with reduced statistical significance; in this case, α is constrained below zero at 95.2% credibility.

points) in Figure 5 obey a tight linear correlation, those in Figure 8(a) exhibit visibly increased scatter as we move toward small q .

It could alternatively be the case that biases due to our choice of mass model or systematic uncertainties in the Advanced LIGO & Virgo selection function yield an artificial preference for $\alpha < 0$. The fact that our mass ratio distribution $p(q|m_1)$ is specified as *conditional* on m_1 implies that our conclusions regarding q , and hence also our conclusions regarding correlations between q and χ_{eff} , are possibly impacted by our choice for $p(m_1)$. Our results, though, behave robustly against different choices for the form of $p(m_1)$. So far, we have shown results obtained while assuming $p(m_1)$ is a power law with a Gaussian peak (the POWER LAW+PEAK model of LIGO Scientific Collaboration & Virgo Collaboration 2021b). If we instead adopt a broken power-law form for $p(m_1)$, our conclusions are virtually unchanged. Similarly, if we instead assume that $p(q)$ is itself a Gaussian rather than a power law, we obtain a consistent estimate of α . As discussed in Appendix A, meanwhile, we account for selection effects in O3 by using the results of an actual injection campaign into Advanced LIGO & Virgo data (LIGO Scientific Collaboration & Virgo Collaboration 2021a, 2021b, 2020c). Analogous injections are not available for the O1 and O2 observing runs, however; for these observing runs we instead estimate the LIGO & Virgo selection function using a publicly available set of mock events that pass a semi-analytic signal-to-noise ratio cut (LIGO Scientific Collaboration & Virgo Collaboration 2020d). It is conceivable that this semi-analytic calculation biases us in unexpected ways. To check this possibility, we have verified that our results hold when performing hierarchical analysis using only those events detected in O3a, thereby avoiding any possibility of bias due to the O1 and O2 selection function.

Another possible concern is the well-known fact that measurements of individual events’ χ_{eff} and q exhibit a large degree of correlation (Cutler & Flanagan 1994; Poisson & Will 1995; Baird et al. 2013; Ohme et al. 2013; Pürrer et al. 2013, 2016; Ng et al. 2018; Tiwari et al. 2018). Both the effective spin and mass ratio appear in the 1.5PN coefficient in the post-Newtonian expansion of a BBH’s phase evolution. It is

this coefficient that is most readily extracted from a gravitational-wave observation, and so the resulting individual measurements of χ_{eff} and q are quite degenerate with one another, particularly in the case of low-mass systems for which the inspiral (rather than merger and ringdown) contributes the bulk of the observed signal-to-noise. This degeneracy can be seen in Figure 1; if examined closely, the posteriors for several individual events take the form of extended arcs that curve down and to the right, toward large χ_{eff} and small q . Since this degeneracy acts in the *same* direction as the $\chi_{\text{eff}} - q$ anti-correlation we identify, one might wonder if our observation of an anti-correlation is just a relic of these measurement degeneracies rather than an intrinsic feature of the underlying BBH population.

If the Bayesian parameter estimation of individual gravitational-wave events is unbiased, hierarchical Bayesian inference is robust against measurement degeneracies. Large degeneracies may well inflate our *uncertainties* on the properties of the underlying BBH population, but will not lead to an erroneous identification of some spurious feature or correlation. In practice, though, we cannot be absolutely confident in the complete and unbiased coverage of parameter estimation. The properties of the massive BBHs GW151226 and GW190521, for example, are still under debate due to the possible existence of additional posterior modes identified upon reanalysis of LIGO/Virgo data (LIGO Scientific Collaboration & Virgo Collaboration 2016, 2020e; Nitz & Capano 2021; Estellés et al. 2020; Mateu-Lucena et al. 2021; Chia et al. 2021). Moreover, it is in principle possible for small but systematic errors in our estimates of parameter estimation *priors*, a critical ingredient in hierarchical inference, to have an outsized cumulative impact on our population-level results.

As a purely qualitative check, we can see to what degree our population-level $\chi_{\text{eff}} - q$ anti-correlation and the $\chi_{\text{eff}} - q$ measurement degeneracies lie in the same direction. The 1.5PN phase correction may be written $\Psi_{1.5} = (\pi \mathcal{M}_c f)^{-2/3} \psi_{1.5}$, with the coefficient (Cutler & Flanagan 1994; Poisson & Will 1995; Baird et al. 2013; Ng et al. 2018)

$$\psi_{1.5} = \eta^{-3/5} \left[\frac{(113 - 76\eta)}{128} \chi_{\text{eff}} - \frac{76\eta\delta}{128} \chi_a - \frac{3\pi}{8} \right]. \quad (8)$$

Here, \mathcal{M}_c is the binary chirp mass, η is the symmetric mass ratio, $\delta = (m_1 - m_2)/(m_1 + m_2)$, and $\chi_a = (\chi_1 \cos t_1 - \chi_2 \cos t_2)/2$. The dashed lines in Figure 5 trace contours of equal $\psi_{1.5}$, assuming that $\chi_a = 0$. It is along these contours that measurement degeneracies between q and χ_{eff} preferentially lie. Note that different choices for χ_a will change the exact contours drawn, but that these differences are small since the influence of χ_a in Equation (8) is suppressed by a factor of δ relative to the term involving χ_{eff} . If our hierarchical inference were contaminated by measurement degeneracies, we expect that our analysis would favor a correlation lying parallel to these constant-in-phase contours. The slope $\alpha \approx -0.45$ favored by the BBH population, however, differs from these contours at mass ratios $q \gtrsim 0.6$ where most detections lie.

This check is purely qualitative, however, and may be affected by the fact that different events lie along different 1.5 PN contours as well as the fact that high-mass events have better measured total masses than chirp masses. In Appendix B, we quantify Bayesian evidences for each event in our sample between two fixed populations: one that includes a $q - \chi_{\text{eff}}$ anti-correlation and one that excludes it. Of the four events that

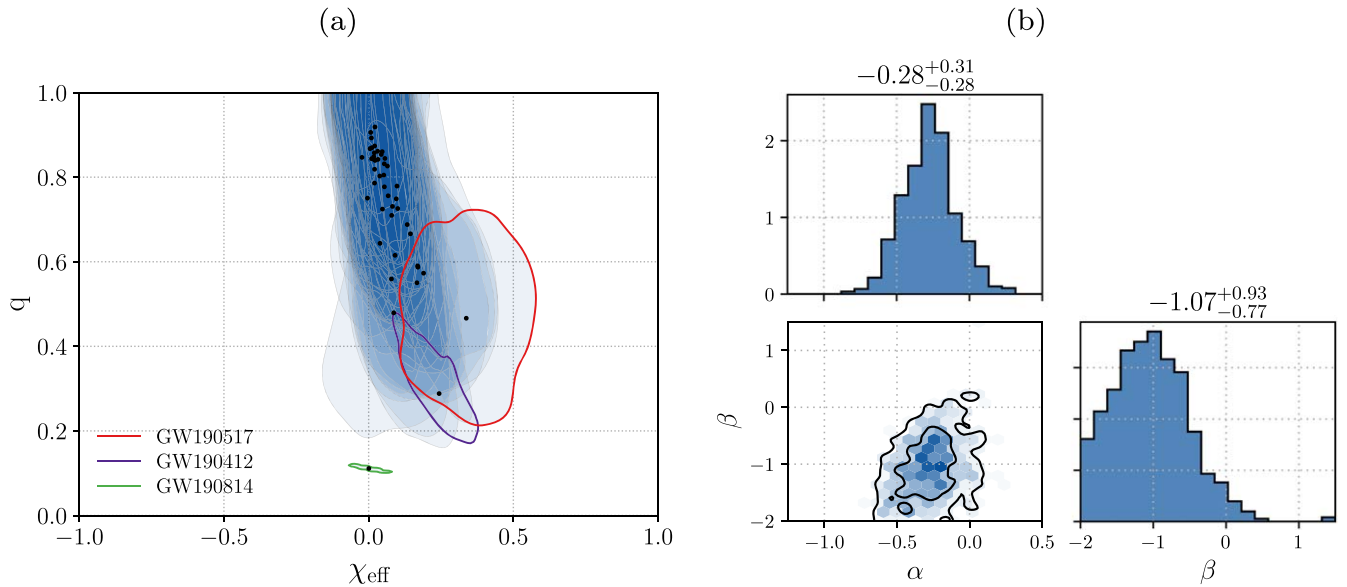


Figure 8. Hierarchical inference results and reweighted single-event posteriors when the outlier event GW190814 is *included* in our BBH sample. The inclusion of GW190814 diminishes our confidence that the mean of the χ_{eff} distribution evolves with mass ratio. Instead, this event leads us to conclude that the *width* of the χ_{eff} distribution likely increases with decreasing q . Comparing the left-hand panel to the reweighted posteriors in Figure 5 (in which GW190814 is excluded from the fit), the bulk of the BBH events still cluster along a best-fit line, but, in order to additionally accommodate the precisely measured properties of GW190814, exhibit with an increasing degree of scatter about this line at smaller mass ratios.

most strongly prefer a correlated population model, none exhibit the sweeping 1.5 PN degeneracies discussed here, offering further evidence that measurement degeneracies are not spuriously driving our measurement of negative α .

As another more concrete test, we verify that measurement degeneracies do not reproducibly confound our analysis by performing an end-to-end injection, recovery, and hierarchical inference of a mock population of BBH events. We randomly draw a large set of BBHs from a broad reference population (see Appendix D for additional details), down-selecting to 5×10^4 “found” events that have a matched-filter network signal-to-noise ratio $\rho \geq 10$ across LIGO-Hanford, LIGO-Livingston, and Virgo. As we are concerned only with χ_{eff} , for simplicity we work only with aligned spins, using the IMRPHENOMD waveform model (Husa et al. 2016; Khan et al. 2016). From our large sample of found injections, we randomly draw $N = 50$ to comprise our “observed” catalog on which we will perform parameter estimation and hierarchical inference. In drawing this catalog, we assign draw probabilities such that our injected population has $\mu_0 = 0.05$, $\sigma_0 = 0.15$, $\alpha = 0$, and $\beta = 0$, with *no* intrinsic correlation between χ_{eff} and q . Additionally, we deliberately introduce a mismatch between the injected distribution of primary masses and the distribution assumed on recovery; while primary masses are drawn from a broken power law, we will fit this population using the same power law and Gaussian peak mixture adopted above.

We perform parameter estimation on each of these 50 events using BILBY (Ashton et al. 2019; Romero-Shaw et al. 2020) in conjunction with the DYNesty (Speagle 2020) nested sampler. Figure 9(a) shows the posteriors we recover from our injections under a default parameter estimation prior. This ensemble of posteriors displays (partially by design) many of the same qualitative features of GWTC-2 seen above in Figure 1(a), with most events clustered near small χ_{eff} and moderately large q , a handful of posteriors whose medians are displaced toward positive χ_{eff} , and several events exhibiting the curving

degeneracy characteristic of joint χ_{eff} and q measurements. Finally, we use the set of resulting posterior samples to hierarchically analyze our injected population, employing the population model described in Section 3 to fit for any correlations between χ_{eff} and q , obtaining the posteriors on α and β shown in Figure 9(b). We recover μ_0 and σ_0 estimates consistent with our injected values, finding $\mu_{\chi,0} = 0.05^{+0.09}_{-0.09}$ and $\log_{10} \sigma_{\chi,0} = 0.13^{+0.09}_{-0.07}$ (medians and central 90% credible uncertainties). More importantly, our results are consistent with *no correlations* between effective spin and mass ratio. Our posteriors allow quite comfortably for the point $(\alpha, \beta) = (0, 0)$ which lies on a contour enclosing 38% of the probability in the $\alpha - \beta$ plane. Hence we correctly conclude that our injected population exhibits no correlations between χ_{eff} and q . As noted above, we deliberately perform our hierarchical inference with an incorrect mass model. For completeness, the dashed histograms in Figure 9(b) show the marginal posteriors obtained on α and β if we instead fit our injection set with the *correct* mass model (a broken power law in primary mass). Our conclusions regarding α are effectively unchanged between these two cases.

This injection case study does not, of course, serve as proof that there are no unknown sources of bias in GWTC-2. It does, however, demonstrate that a false-positive identification of χ_{eff} and q does not readily appear when hierarchically analyzing fully realistic BBH posteriors, particularly those posteriors that exhibit measurement degeneracy between mass ratio and effective spin.

It is possible that our analysis is affected by other systematic biases due to imperfect detector calibration and/or biases in the waveform models used for parameter estimation. Both possibilities, though, are unlikely to give rise to the observed $q - \chi_{\text{eff}}$ anti-correlation. Parameter estimation samples used in this study have been marginalized over a frequency-dependent calibration uncertainty budget (LIGO Scientific Collaboration & Virgo Collaboration 2021a; Cahillane et al. 2017; Sun et al.

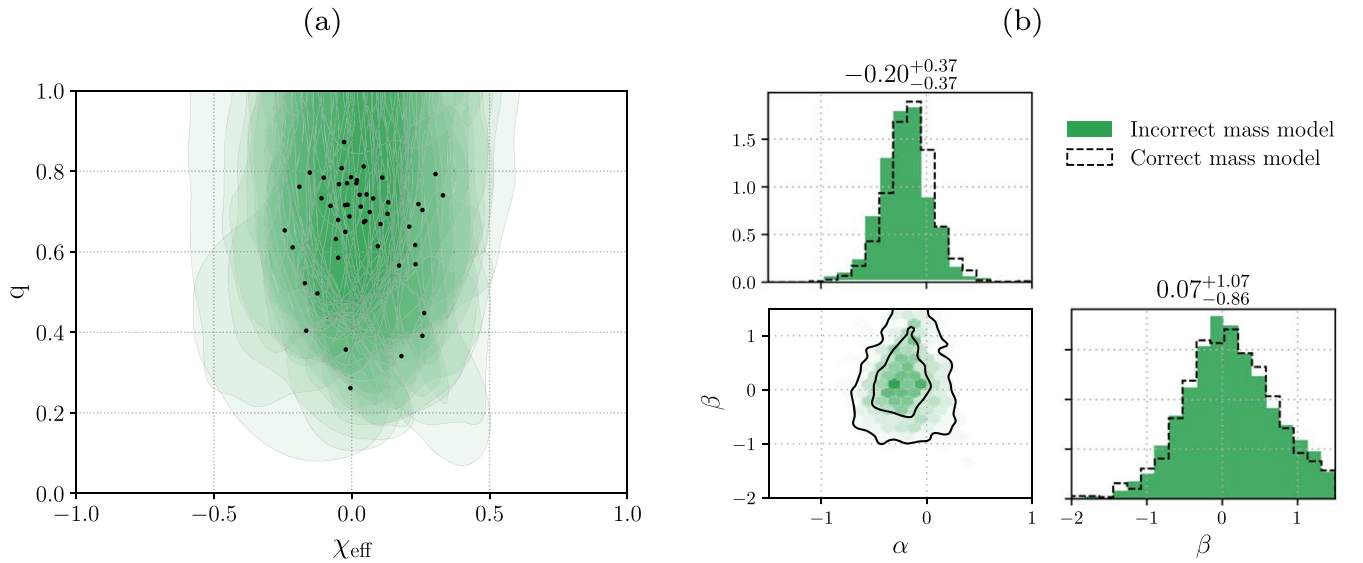


Figure 9. (a): Parameter estimation posteriors for a catalog of mock BBH detections. Each contour encloses the central 90% credible region for a given event, while black points mark one-dimensional medians on q and χ_{eff} . These mock detections are drawn from a population with *no* intrinsic correlations between q and χ_{eff} , but with parameters otherwise consistent with those of GWTC-2. Our resulting posteriors display the same key features as those of real BBH detections in Figure 1(a), including measurement degeneracy between q and χ_{eff} . (b): Inferred posteriors on the parameters governing correlations between q and χ_{eff} [see Equations (6) and (7)] when performing hierarchical analysis on the mock catalog. Our results are consistent with $\alpha = \beta = 0$, correctly reflecting the fact that our underlying injections possess no intrinsic correlations between q and χ_{eff} . As discussed further in the text, we further test the stability of this result by deliberately analyzing this population with an incorrect mass model (power law with a Gaussian peak; green). Results obtained using the *correct* mass model from which injections are drawn (broken power law) are shown via the dashed histograms.

2020). As discussed in Appendix A, meanwhile, we use a union of parameter estimation samples from several distinct waveform families, mitigating potential biases that might be peculiar to any one waveform model.

5. Astrophysical Implications

The physical implications of an anti-correlation between q and χ_{eff} are unclear. There are generally two approaches one might take in attempting to explain this relationship. First, such an anti-correlation could conceivably arise from processes acting within a *single* population of BBHs arising from a common formation channel. Second, a global anti-correlation could arise if observed BBHs originate from some superposition of formation channels, with some favoring high q and low χ_{eff} and others yielding high χ_{eff} with low q . We comment on each of these possibilities in turn.

BBHs arising from isolated stellar binaries have spins that result from a complex interplay of angular momentum transport in stellar cores, tidal torques operating between stars, and episodes of mass transfer (Spruit 2002; Gerosa et al. 2018; Qin et al. 2018; Zaldarriaga et al. 2018; Belczynski et al. 2020; Fuller & Ma 2019; Bavera et al. 2020, 2021). The details of the latter two processes depend on binary mass ratio, and so might impart a relationship between q and χ_{eff} . The generally predicted relationship between q and χ_{eff} , however, differs from the relationship we see here observationally. Bavera et al. (2020), for instance, predict that BBHs originating from common envelope (CE) exhibit an increased *scatter* in χ_{eff} toward positive values with increasing q . Bavera et al. (2021) later identify a similar trend among binaries that undergo Eddington-limited stable mass transfer. In our analysis, this effect would manifest as positive values for α and β ; this possibility is ruled out at high credibility. We do note, though, that the simulations of Bavera et al. (2021) do contain two special cases in which isolated BBHs could exhibit anti-correlated q and χ_{eff} . CE binaries with

very high common envelope efficiencies and stable mass-transfer binaries with super-Eddington accretion (e.g., fourth row in Figure G.2 and third row in Figure G.3 of Bavera et al. 2021, respectively) each predict structure in the $\chi_{\text{eff}} - q$ plane that is at least qualitatively similar to the behavior we find here.

Dynamically assembled binaries in dense stellar clusters may also naturally exhibit a correlation between q and χ_{eff} if they experience repeated, hierarchical mergers. The heavy remnants of BBH mergers are generically rapidly rotating, with spins centered around $\chi \sim 0.7$. If these “second-generation” black holes subsequently undergo additional mergers with “first-generation” black holes, the result will be a population of BBHs observed with very unequal mass ratios and preferentially large spins (Fishbach et al. 2017; Gerosa & Berti 2017; Doctor et al. 2020; Rodriguez et al. 2019; Kimball et al. 2021; Gerosa & Fishbach 2021). In the absence of any preferred directions, though, the result will be a *broadening* of the χ_{eff} distribution toward smaller q (e.g., negative β). This prediction is at odds with the fact that the data prefer larger but *preferentially positive* effective spins with smaller mass ratio.

As noted above, an alternative interpretation is that the BBHs observed in GWTC-2 are a mixture of populations arising from distinct formation channels (e.g., Franciolini et al. 2021; Wong et al. 2021; Zevin et al. 2021); in this picture the measured q versus χ_{eff} anti-correlation arises from the presence of two or more sub-populations that manifest in the $\chi_{\text{eff}} - q$ plane. As discussed in Section 3, both GW190412 and GW190517 lie in the same area of this plane when reweighted to a population-informed prior, each exhibiting small q and confidently positive χ_{eff} . Chia et al. (2021) have also recently argued that GW151226 also exhibits similar characteristics when re-analyzed with a waveform model including spin precessing and higher-order radiation modes (Pratten et al. 2021). These events might together suggest the presence of a secondary sub-population appearing at low q and high χ_{eff} .

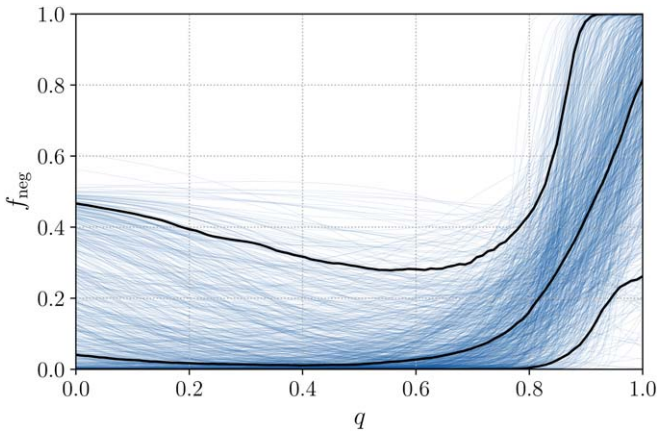


Figure 10. The inferred fraction of BBHs with negative effective spin as a function of mass ratio, assuming that the population is well-described by a Gaussian distribution. Blue traces mark individual draws from our posterior on population-level parameters (see Figure 3) and black traces mark the median and central 90% credible bounds on f_{neg} . When allowing for correlations between q and χ_{eff} , however, the data are fairly agnostic about the *requirement* for negative effective spins; see Table 1 and further discussion in Section 5 below.

If we assume that GWTC-2 comprises a mixture of BBHs arising from isolated stellar evolution, with purely positive χ_{eff} , and systems forming dynamically in stellar clusters, with isotropic spin orientations, we can attempt to characterize the implied mixture fraction between these two populations as a function of mass ratio. Using our hierarchical measurement of $\{\mu_0, \sigma_0, \alpha, \beta\}$, in Figure 10 we show our posterior on the fraction

$$f_{\text{neg}}(q) = \int_{-1}^0 p(\chi_{\text{eff}}|q) d\chi \quad (9)$$

of events with negative effective spins versus q . Individual blue traces show the result of single draws from our posterior on parameters governing the spin–mass ratio distribution, while solid black lines show the mean and central 90% credible bounds on $f_{\text{neg}}(q)$. At $q \sim 1$, the data are consistent with half (or even more) of all systems exhibiting negative effective spins. At low q , meanwhile, our inferred mean on the χ_{eff} distribution has shifted to larger positive values and so the fraction of systems with negative χ_{eff} is likely quite small.

If we very simplistically assume that BBHs formed dynamically in dense clusters exhibit a *symmetric* χ_{eff} distribution, while BBHs arising in the field can possess only positive χ_{eff} , then we obtain crude estimates

$$f_{\text{dyn}}(q) = 2f_{\text{neg}}(q) \quad (10)$$

and

$$f_{\text{field}}(q) = 1 - 2f_{\text{neg}}(q) \quad (11)$$

of the mixture fractions between channels as a function of q . The result is shown in Figure 11. At $q \sim 1$, the fraction of BBHs arising in clusters is at least $f_{\text{dyn}} \approx 0.5$, or as high as one (note that some individual traces rise unphysically to values $f_{\text{dyn}} > 1$). Below $q \approx 0.8$, meanwhile, it is likely that f_{dyn} drops precipitously while the fraction f_{field} of systems formed in the field rises toward unity. The rapid transition at $q \approx 0.8$ corresponds to the mass ratio at which $\mu_\chi(q)$ and $\sigma_\chi(q)$ are each most precisely measured; see Figure 4. This result, with




Figure 11. The implied fractions of BBHs arising from toy “dynamical” (bottom) and “isolated field binary” (top) formation channels, assuming that the former yields a χ_{eff} distribution symmetric about zero, while the latter yields exclusively positive effective spins. Blue traces mark individual draws from the hyperposterior in Figure 3, and black traces denote medians and central 90% credible bounds. Above $q \gtrsim 0.6$, the implied fraction of field events falls sharply with q , while at low mass ratios our uncertainties on f_{field} and f_{dyn} become large due to the small number of confidently unequal-mass BBHs. As in Figure 10, however, we note these results assume that the χ_{eff} distribution at a given q is well-described by a Gaussian that extends to negative χ_{eff} ; this assumption is not necessarily independently supported by observation, as discussed below.

dynamically-formed systems exhibiting preferentially equal mass ratios while isolated BBHs possess unequal mass ratios, is surprising. Theoretical modeling and population synthesis generally predict the opposite, such that BBHs formed in clusters have more unequal mass ratios than those formed in the field, due to the possibility of exchanges during three-body encounters and/or hierarchical mergers (Bouffanais et al. 2019; Di Carlo et al. 2020; Bouffanais et al. 2021).

In constructing and interpreting Figures 10 and 11, it is important to explore whether the constraints on $f_{\text{neg}}(q)$ arise from informative data or whether they are simply extrapolations based on the Gaussian model we adopt for $p(\chi_{\text{eff}}|q)$. When neglecting the possibility of correlations between q and χ_{eff} , the LIGO Scientific Collaboration & Virgo Collaboration (2021b) found evidence for negative effective spins among GWTC-2. In particular, when fitting the χ_{eff} with a Gaussian truncated between $\chi_{\text{eff,min}} \leq \chi_{\text{eff}} \leq 1$, for some unknown $\chi_{\text{eff,min}}$, this lower truncation bound was inferred to be negative at high credibility. This evidence for negative effective spins was corroborated via an alternative model that directly fit for the spin magnitudes and tilt angles of component black holes. Callister et al. (2021b) hierarchically modeled the χ_{eff} distribution not as Gaussian, but as the convolution of an underlying component spin distribution with a distribution of

natal kicks experienced during the core-collapse of black hole progenitors; the results of this analysis too indicated the presence of negative χ_{eff} . More recently, though, Roulet et al. (2021) have argued that present detections *lack* strong evidence for the existence of negative χ_{eff} and/or extreme spin-orbit misalignment; see further discussion below.

These conclusions, though, were generally based on models that did not include correlations between q and χ_{eff} . To understand the extent to which negative effective spins remain required by the data, we use the DYNESTY nested sampler (Speagle 2020) to compute Bayesian evidences for two additional models beyond those introduced in Sections 2 and 3 above:

1. Independent χ_{eff} and q , with purely positive χ_{eff} . In this case, Equation (4) is truncated on $[0, 1]$, rather than $[-1, 1]$, with a restricted prior $0 \leq \mu \leq 1$ on its mean.
2. Correlated χ_{eff} and q , with purely positive χ_{eff} . Equation (5) is truncated on $[0, 1]$, and we adopt a restricted prior enforcing $0 \leq \mu_{\chi,0} \leq 1$.

In each case we use the same modeling assumptions and priors as listed in Appendix A unless otherwise noted.

Table 1 lists the (natural) log-evidences for these two “purely positive spin” models in addition to the evidences for the models of Sections 2 and 3 discussed earlier. The log-evidences are scaled such that the most disfavored model has $\ln \text{Evidence} = 0$. The most favored model is one that allows for both a $\chi_{\text{eff}} - q$ correlation *and* negative effective spins, with a Bayes factor $\ln \mathcal{B} = 3.64$ relative to a model including neither effect. However, allowing for $\chi_{\text{eff}} - q$ correlations *weakens* evidence for negative effective spins; we find a nearly uninformative Bayes factor $\ln \mathcal{B} = 0.10$ between our most favored model and one that allows for correlated χ_{eff} and q but *without* negative χ_{eff} . As the Advanced LIGO and Virgo detectors continue adding to the list of BBH detections, it will be important to revisit this model comparison, refining the evidence (or lack thereof) for negative effective spins and understanding the extent to which conclusions regarding negative χ_{eff} and/or correlations between parameters can confound one another.

We note that, in gauging the evidence for/against negative χ_{eff} , we are still assuming that BBHs are well-described by a *single* population in the $q - \chi_{\text{eff}}$ plane. Roulet et al. (2021) recently demonstrated that conclusions regarding the presence of negative χ_{eff} depend strongly on whether or not one allows for a *second* population of events with near-vanishing spins. When hierarchically measuring the BBH χ_{eff} distribution using a mixture model that included such a vanishing spin sub-population, they argued that evidence for negative effective spins is significantly reduced. It is interesting to note that a χ_{eff} distribution with a vanishing spin sub-population arises naturally from our analysis here. Figure 12 shows the marginal χ_{eff} distribution implied by our results, having integrated the joint mass ratio–spin distribution $p(q, \chi_{\text{eff}})$ over q . The result is an asymmetric distribution consistent with the marginal χ_{eff} model favored by Roulet et al. (2021), exhibiting both a peak at $\chi_{\text{eff}} \approx 0$ (occurring at high q) and a shoulder extending toward positive values (arising from the low- q population).

Finally, it should be emphasized that despite our focus on BBH formation via “canonical” field binary and cluster scenarios in this discussion, these are certainly not the only




Figure 12. The marginal χ_{eff} distribution obtained by integrating our inferred joint distribution $p(q, \chi_{\text{eff}})$ over q . Blue traces give single draws from our hyperposterior in Figure 3, while black traces denote the median and central 90% credible bounds. The marginal distribution exhibits a sharp peak about zero, with an extended tail toward large and positive χ_{eff} . The peak occurs at $q \approx 1$, where the χ_{eff} grows narrow, while the extended tail arises from the shift in χ_{eff} toward more positive values with lower q .

two options that exist. Other possible avenues for BBH formation and mergers include hierarchical triples (Liu & Lai 2017, 2018; Antonini et al. 2018; Rodriguez & Antonini 2018; Martinez et al. 2020) and BBH assembly in the disks of active galactic nuclei (AGNs; McKernan et al. 2018, 2020; Stone et al. 2017; Tagawa et al. 2020); it is not well understood what kinds of q and χ_{eff} correlations might arise under these alternative scenarios. AGN disks, in particular, offer opportunities for hierarchical BBH mergers while maintaining a preferred direction (set by the accretion disk’s angular momentum), and thus might be able to explain the $q - \chi_{\text{eff}}$ anti-correlation we observe here.

6. Conclusions

In this paper, we have presented evidence for an intrinsic anti-correlation between the effective spins and mass ratios of detected BBH mergers. Under a hierarchical analysis of the BBH population using Advanced LIGO and Virgo’s GWTC-2 catalog, we find that equal-mass binaries preferentially exhibit a narrow χ_{eff} distribution centered at zero, while unequal mass events favor a χ_{eff} distribution systematically shifted toward larger, positive values.

This anti-correlation is unexpected. As discussed in Section 5, it is not clear what physical processes might give rise to such an effect; in fact, standard ideas about BBH formation via isolated stellar evolution and dynamical assembly in dense clusters make predictions at odds with the trend we discover here. For this reason, we have attempted to critically examine and test various sources of bias or systematic uncertainty that might lead to spurious conclusions regarding the joint $\chi_{\text{eff}} - q$ distribution in BBH mergers. In Section 4, we found that the observed anti-correlation between q and χ_{eff} could not be clearly attributed to selection effects, measurement degeneracies, or a small number of outlier events among the BBH population.

Looking ahead, the continued detection of additional BBHs by Advanced LIGO and Virgo (to be joined soon by the

KAGRA experiment; Akutsu et al. 2021) will be crucial to better understand the intrinsic relationship between q and χ_{eff} , further bolstering our conclusions or pointing instead to a statistical fluctuation or as-of-yet unknown systematic. Simultaneously, it will be valuable to re-evaluate our understanding of BBH formation in light of the results presented here, exploring which formation channel (or combination of channels) can accommodate the observed trend in q and χ_{eff} . If confirmed to be astrophysical, the anti-correlated mass ratios and effective spins of BBHs may offer a firm observational foothold on our path toward understanding the origin of black hole mergers.

We thank Katerina Chatziioannou for helpful conversation, and our anonymous referees whose suggestions have greatly improved this manuscript. The Flatiron Institute is a division of the Simons Foundation, supported through the generosity of Marilyn and Jim Simons. This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation. We are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. This research has made use of data, software, and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gwopenscience.org/>), a service of LIGO Laboratory, the LIGO Scientific Collaboration, and the Virgo Collaboration. 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, and Spain.

Software: ASTROPY (Astropy Collaboration 2013, 2018), BILBY (Ashton et al. 2019; Romero-Shaw et al. 2020), DYNESTY (Speagle 2020), EMCEE (Foreman-Mackey et al. 2013), H5PY (Collette et al. 2021), MATPLOTLIB (Hunter 2007), NUMPY (Harris et al. 2020), PYCBC (Nitz et al. 2021a), PANDAS (Pandas Development Team, T 2020; McKinney 2010), SCIPY (Virtanen et al. 2020).

Appendix A

Hierarchical Inference of the BBH Population

Here we discuss details of our hierarchical inference of the BBH population with GWTC-2. We model the primary mass distribution following the POWER LAW + PEAK model of Talbot & Thrane (2018) and LIGO Scientific Collaboration & Virgo Collaboration (2021b), in which primary BBH masses are described as a mixture

$$p(m_1|f_p, \lambda, \mu_m, \sigma_m, m_{\text{max}}) = f_p P(m_1|\lambda, m_{\text{max}}) + (1 - f_p)N(m_1|\mu_m, \sigma_m, m_{\text{max}}) \quad (\text{A1})$$

between a power law $P(m_1|\lambda, m_{\text{max}}) \propto m_1^\lambda$ and a Gaussian $N(m_1|\mu_m, \sigma_m, m_{\text{max}}) \propto \exp\left[-\frac{(m_1 - \mu_m)^2}{2\sigma_m^2}\right]$, normalized across the range $5 M_\odot \leq m_1 \leq m_{\text{max}}$. We assume a redshift distribution that is proportional to the differential comoving volume $\frac{dV_c}{dz}$, with a possible evolution in the merger rate toward higher redshift (Fishbach et al. 2018; Callister et al. 2020a),

$$p(z|\kappa) \propto \frac{1}{1+z} \frac{dV_c}{dz} (1+z)^\kappa. \quad (\text{A2})$$

The additional factor of $(1+z)^{-1}$ in Equation (A2) converts a uniform-in-time source-frame distribution to our detector frame. We adopt cosmological parameters consistent with those reported in Ade et al. (2016).

As discussed in the main text, we explore two related models for the mass ratio and effective spin distribution of BBHs. In Section 2, we first consider the case in which these parameters are uncorrelated, describing the conditional distribution $p(q|m_1, \gamma)$ of mass ratios as a power law with index γ and the distribution $p(\chi_{\text{eff}}|\mu_\chi, \sigma_\chi)$ of effective spin parameters as a Gaussian with mean μ_χ and standard deviation σ_χ ; see Equations (3) and (4). In Section 3 and beyond, we subsequently expand this model to allow for population-level correlations between the mass ratios and spins of BBHs. In particular, we preserve the Gaussian form for $p(\chi_{\text{eff}})$, but now allow its mean and standard deviation to vary as a function of mass ratio; see Equations (5), (6), and (7).

Together, we hierarchically fit the collection of hyperparameters that govern the BBH population by considering the $N_{\text{det}} = 44$ BBH candidates among GWTC-2 with false alarm rates below one per year (LIGO Scientific Collaboration & Virgo Collaboration 2021a). We make use of the parameter estimation results described in LIGO Scientific Collaboration & Virgo Collaboration (2019a, 2021a) and made publicly available through the Gravitational-Wave Open Science Center (Vallisneri et al. 2015; LIGO Scientific Collaboration & Virgo Collaboration 2020f, 2020g, 2020h, 2021d). For BBHs first announced in GWTC-1, we use the `Overall_Posterior` samples formed by the union of results under two different waveform families. For BBHs observed in LIGO and Virgo’s O3a observing run, we use the `PrecessingSpinIMRHM` samples generated by waveforms including the effects of spin precession and higher-order modes.

Given posteriors $p(\theta_i|d_i)$ on the individual parameters θ_i (e.g., component masses, spins, etc.) of each event conditioned on its observed data d_i , the corresponding posterior on the population parameters Λ is (Loredo 2004; Taylor & Gerosa 2018; Mandel et al. 2019; Vitale et al. 2020)

$$p(\Lambda | \{d_i\}) \propto p(\Lambda) \xi^{-N_{\text{det}}}(\Lambda) \times \prod_{i=1}^{N_{\text{det}}} \int d\theta_i p(\theta_i|d_i) \frac{p(\theta_i|\Lambda)}{p_{\text{pe}}(\theta_i)}. \quad (\text{A3})$$

This form of the likelihood includes implicit marginalization over the overall *rate* of BBH mergers, using a log-uniform prior on the expected number of detections (Fishbach et al. 2018; Mandel et al. 2019). In Equation (A3), $p_{\text{pe}}(\theta_i)$ is the default prior adopted for purposes of parameter estimation and $\xi(\Lambda)$ is the population-weighted detection efficiency, discussed further below. Parameter estimation for GWTC-2 is achieved using priors that are uniform in *detector* frame masses and

Table 2

Priors Adopted for the Hyperparameters with Which We Describe the Mass, Spin, and Redshift Distributions of BBHs

Parameter	Prior	Defined in
$\mu_{\chi,0}$	$U(-1, 1)$	Equations (4) and (5)
$\log_{10} \sigma_{\chi,0}$	$U(-1.5, 0.5)$	Equations (4) and (5)
α	$U(-2.5, 1)$	Equation (5)
β	$U(-2, 1.5)$	Equation (5)
m_μ	$U(20 M_\odot, 100 M_\odot)$	Equation (A1)
m_σ	$U(1 M_\odot, 10 M_\odot)$	Equation (A1)
f_p	$U(0, 1)$	Equation (A1)
λ	$U(-5, 4)$	Equation (A1)
γ	$U(-2, 10)$	Equation (3)
m_{\max}	$U(60 M_\odot, 100 M_\odot)$	Equation (A1)
κ	$N(0, 6)$	Equation (A2)

Note. Here, $U(\min, \max)$ denotes a uniform prior distribution between the given minimum and maximum values, while $N(a, b)$ is a normal distribution with mean a and standard deviation b . When modeling the χ_{eff} distribution in Section 2 without any correlations with q , our priors on μ_χ and σ_χ are identical to those listed here for $\mu_{\chi,0}$ and $\mu_{\chi,0}$.

Euclidean volume, corresponding to an implicit prior (LIGO Scientific Collaboration & Virgo Collaboration 2019a, 2021a; Callister 2021a)

$$p_{\text{pe}}(m_1, m_2, z) \propto (1+z)^2 D_L^2(z) \frac{dD_L}{dz} \quad (\text{A4})$$

on source-frame masses and redshift, where $D_L(z)$ is the luminosity distance at redshift z . Component spin priors are uniform in magnitude and isotropic in orientation, corresponding to an effective spin prior $p_{\text{pe}}(\chi_{\text{eff}})$ given by Equation (10) of Callister (2021a). Meanwhile, $p(\Lambda)$ is our prior on the population-level parameters. Unless stated otherwise, we use the priors listed in Table 2. We still cannot yet employ Equation (A3) since we do not have direct access to underlying posteriors $p(\theta_i | d_i)$, but instead have a collection of discrete *samples* drawn from each event’s posterior. Instead, we replace integration over $p(\theta_i | d_i)$ with an ensemble average taken over the posterior samples associated with each event:

$$p(\Lambda | \{d_i\}) \propto p(\Lambda) \xi^{-N_{\text{det}}}(\Lambda) \prod_{i=1}^{N_{\text{det}}} \left\langle \frac{p(\theta_i | \Lambda)}{p_{\text{pe}}(\theta_i)} \right\rangle \quad (\text{A5})$$

We sample over Equation (A5) using the EMCEE Markov Chain Monte Carlo sampler (Foreman-Mackey et al. 2013) to obtain the posteriors shown in Figures 3, 14, and 15. When computing and comparing Bayesian evidences in Section 5, we instead implement and integrate over Equation (A5) using the DYNesty nested sampler (Speagle 2020).

In Equations (A3) and (A5), the detection efficiency $\xi(\Lambda)$ quantifies the fraction of events that we expect to pass our detection criteria, given a population described by Λ :

$$\xi(\Lambda) = \int d\theta P_{\text{det}}(\theta) p(\theta | \Lambda). \quad (\text{A6})$$

Here, $P_{\text{det}}(\theta)$ is the detection probability for a particular event with parameters θ . The detection efficiency corrects for search selection effects and so is critical to accurately calculate. We estimate $\xi(\Lambda)$ using the injection campaign reported in LIGO Scientific Collaboration & Virgo Collaboration (2021b, 2020c),

selecting successfully found injections (with recovered false alarm rates below one per year in at least one pipeline) and reweighting to the proposed population Λ :

$$\xi(\Lambda) = \frac{1}{N_{\text{inj}}} \left\langle \frac{p(\theta | \Lambda)}{p_{\text{inj}}(\theta)} \right\rangle_{\text{Found injections}}, \quad (\text{A7})$$

where N_{inj} is the total number of injections (including those that are *not* recovered) and $p_{\text{inj}}(\theta)$ is the reference distribution from which injections were drawn. As discussed in LIGO Scientific Collaboration & Virgo Collaboration (2020c), the injected masses follow $p_{\text{inj}}(m_1) \propto m_1^{-2.35}$ for $2 M_\odot \leq m_1 \leq 100 M_\odot$ and $p(q | m_1) \propto q^2$. The injections additionally have purely aligned component spins ($\theta_1 = \theta_2 = 0$ or π) distributed uniformly between $-1 \leq \chi_z \leq 1$; the corresponding distribution of χ_{eff} is (Callister 2021a)

$$p_{\text{inj}}(\chi_{\text{eff}} | q) = \begin{cases} \frac{(1+q)^2(1+\chi_{\text{eff}})}{4q} & \left(\chi_{\text{eff}} \geq -1 \ \& \ \chi_{\text{eff}} < -\frac{1-q}{1+q} \right) \\ \frac{1+q}{2} & \left(\chi_{\text{eff}} \geq -\frac{1-q}{1+q} \ \& \ \chi_{\text{eff}} \leq \frac{1-q}{1+q} \right) \\ \frac{(1+q)^2(1-\chi_{\text{eff}})}{4q} & \left(\chi_{\text{eff}} > \frac{1-q}{1+q} \ \& \ \chi_{\text{eff}} \leq 1 \right) \end{cases} \quad (\text{A8})$$

One operation frequently performed in this paper is the reweighting of single-event posteriors from the default prior $p_{\text{pe}}(\theta)$ adopted for parameter estimation to some new *population-informed* prior based on our hierarchical inference of Λ . If we denote by $p_{\text{pe}}(\theta | d)$ the posterior on a particular event obtained through standard parameter estimation, then via Bayes’ theorem, the reweighted posterior corresponding to some *fixed* value of Λ is

$$p(\theta | d, \Lambda) \propto p_{\text{pe}}(\theta | d) \left(\frac{p(\theta | \Lambda)}{p_{\text{pe}}(\theta)} \right). \quad (\text{A9})$$

We do not, of course, measure Λ perfectly. What we instead want is a reweighted posterior that has been marginalized over our uncertainty on the nature of the underlying population. Given an original set of samples $\{\theta\}_{\text{pe}}$ drawn from $p_{\text{pe}}(\theta | d)$ and a set of hyperparameter samples $\{\Lambda\}$ generated by our population fit, we can obtain a new set of samples drawn from this reweighted and marginalized posterior by doing the following:

1. Randomly select a hyperparameter sample $\Lambda_i \in \{\Lambda\}$,
2. For each sample $\theta_j \in \{\theta\}_{\text{pe}}$, compute the weights $w_j = p(\theta_j | d, \Lambda_i) / p_{\text{pe}}(\theta_j)$ that correspond to the parenthetical in Equation (A9),
3. Randomly select and store a posterior sample θ_j according to weights w_j , and
4. Repeat.

As is discussed in Callister (2021b) and Farr (2021), this algorithm can be shown to properly avoid the “double counting” of information; the updated prior imposed on any one event is itself informed only by the *other* events in one’s sample, excluding the particular event of interest.

Another operation performed in Sections 2 and 3 is the generation of *predicted* observations, marginalized over the parameters of a particular population model. This proceeds in a

manner similar to posterior reweighting, but instead using the pipeline injection sets that inform our calculation of $\xi(\Lambda)$ above:

1. Randomly select a hyperparameter sample $\Lambda_i \in \{\Lambda\}$,
2. For each found injection, with parameters θ_j , compute weights $w_j = p(\theta_j | \Lambda_i) / p_{\text{inj}}(\theta_j)$ to reweight from the injected reference distribution to the proposed population,
3. Randomly select and store one found injection according to the weights w_j , and
4. Repeat.

Appendix B

Which Events Drive the Spin–Mass Ratio Correlation?

In Section 4, we verified that our measurement of an anti-correlation between q and χ_{eff} is not driven solely by GW190517 and GW190412. Here, we try to understand more generally *which* events are driving (or resisting) our result. For every event in our sample we compute a Bayes factor between two fixed BBH population models: a fixed “correlated” population consistent with our posteriors on the expanded model of Section 3, and a fixed “uncorrelated” population consistent with results from our initial model in Section 2. The events with Bayes factors favoring the correlated population are likely those same events driving our measurement of $\alpha < 0$, while events that favor the uncorrelated population likely act to resist a $q - \chi_{\text{eff}}$ anti-correlation. For both models, we choose hyperparameters consistent with the median values inferred by a complete hierarchical analysis. Specifically, we take $m_\mu = 33 M_\odot$, $m_\sigma = 5 M_\odot$, $f_p = 0.1$, $\lambda = -2.5$, $\gamma = 1.1$, $m_{\text{min}} = 5 M_\odot$, $m_{\text{max}} = 100 M_\odot$, and $\kappa = 2.7$. For the “correlated” population, we take $\mu_{\chi,0} = 0.2$, $\log_{10} \sigma_{\chi,0} = -1$, $\alpha = -0.45$, and $\beta = 0$, while for the “uncorrelated” population we fix $\mu_\chi = 0.05$ and $\log_{10} \sigma_\chi = -1$.

In Figure 13, we show the four events with the largest (top row) and smallest (bottom row) Bayes factors between our fixed correlated and uncorrelated population models. Blue points show each events’ posterior samples (under default parameter estimation priors) and the correlated versus uncorrelated Bayes factor is printed in the lower right corner of each subplot. For reference, the filled blue band shows the range of spins $\mu_\chi \pm \sigma_\chi$ primarily supported by the uncorrelated population, while the sloped red band shows spin range $\mu_\chi(q) \pm \sigma_\chi(q)$ favored in the correlated population.

As anticipated, GW190517 and GW190412 most strongly identify as members of the correlated population. Several other events with posterior morphologies similar to that of GW190517, however, are also picked out as favoring the correlated population model. Most notably, *none of the events that most favor a $q - \chi_{\text{eff}}$ anti-correlation exhibit the long, curving degeneracies which we discussed in Section 4.* This further bolsters our confidence that the $q - \chi_{\text{eff}}$ anti-correlation identified among the BBH population is not due to the combined “leakage” of degenerate measurements for individual events. The common morphologies among GW190517, GW190519, and GW190706 additionally offer intuition as to why excluding GW190412 from our analysis yields an even *stronger* anti-correlation; see Figure 7. If we were to ignore GW190412 and redraw a red band passing through the posteriors of GW190517, GW190519, and GW190706, our band would likely be far shallower than that currently plotted in Figure 13, corresponding to a more negative value of α .

Meanwhile, the four events with the strongest Bayes factors *against* a spin–mass ratio correlation display a variety of morphologies. GW170104 favors low spins and a mass ratio near unity, while GW190929 and GW190512 favor more unequal mass ratios. Interestingly, GW190828 displays the same anti-correlated *measurements* that we considered in




Figure 13. The events with the largest (top row) and smallest (bottom row) Bayes factors, indicated in the lower right corner of each panel, between population models with and without correlations between q and χ_{eff} as described in Appendix B. In each plot, blue points mark posterior samples for each event under default parameter estimation priors, while the red and blue bands show 1σ bounds $\mu_\chi(q) \pm \sigma_\chi(q)$ on the χ_{eff} distribution under the correlated and uncorrelated populations. In Section 4, we worried that large measurement degeneracies between q and χ_{eff} might spuriously cause an apparent correlation between these parameters among the BBH population. Of the events shown here that most favor a population-level correlation, though, none exhibit large curving measurement degeneracies.

Section 4. However, this event resists our measurement of negative α , rather than driving it.

Appendix C Detailed Parameter Estimation Results

Here, we show additional parameter estimation results neglected in the main text. Figure 14 shows posteriors on the parameters governing the BBH effective spin and mass ratio distributions under the base model discussed in Section 2, *without* the possibility of $q - \chi_{\text{eff}}$ correlations. Under this base model, we infer an effective spin distribution centered near $\chi_{\text{eff}} \approx 0.05$ with a standard deviation of $\sigma_\chi \approx 0.1$. Our posteriors on parameters governing the primary mass and redshift distributions are effectively identical to those presented in LIGO Scientific Collaboration & Virgo Collaboration (2021b), and so we do not show their complete posteriors here, although

we include their one-dimensional marginalized distributions in Figure 15 below.

In Figure 3 in the main text, we showed the joint posterior on parameters governing the BBH mass ratio and spin distribution under our expanded model allow for correlations between these parameters. Figure 15 shows the posterior on the remaining parameters when including the effects of the $q - \chi_{\text{eff}}$ correlation. There are no significant correlations between the parameters plotted in Figure 3 and those plotted here. Allowing for a $q - \chi_{\text{eff}}$ correlation minimally impacts our conclusions regarding the primary mass and redshift distributions of BBHs. For comparison, the unfilled gray distributions in Figure 15 show posteriors obtained using the base model from Section 2; these are generally identical to the posteriors obtained from the expanded model (in blue). The largest shift occurs in the posterior for the width σ_m of the Gaussian peak in m_1 . However, σ_m remains effectively unconstrained in both models.




Figure 14. Posteriors on the parameters governing the mass ratio and effective spin distributions of BBHs under the base model in Section 2 with no correlations between q and χ_{eff} . The labels above each one-dimensional posterior give the medians and central 90% credible uncertainties on each parameter, while the contours in each two-dimensional posterior enclose the 50% and 90% credible regions. Posteriors on the remaining parameters governing the BBH primary mass and redshift distributions are effectively identical to those recovered in LIGO Scientific Collaboration & Virgo Collaboration (2021b). For reference, though, the one-dimensional marginalized posteriors on these other parameters are shown as unfilled gray histograms in Figure 15 below.

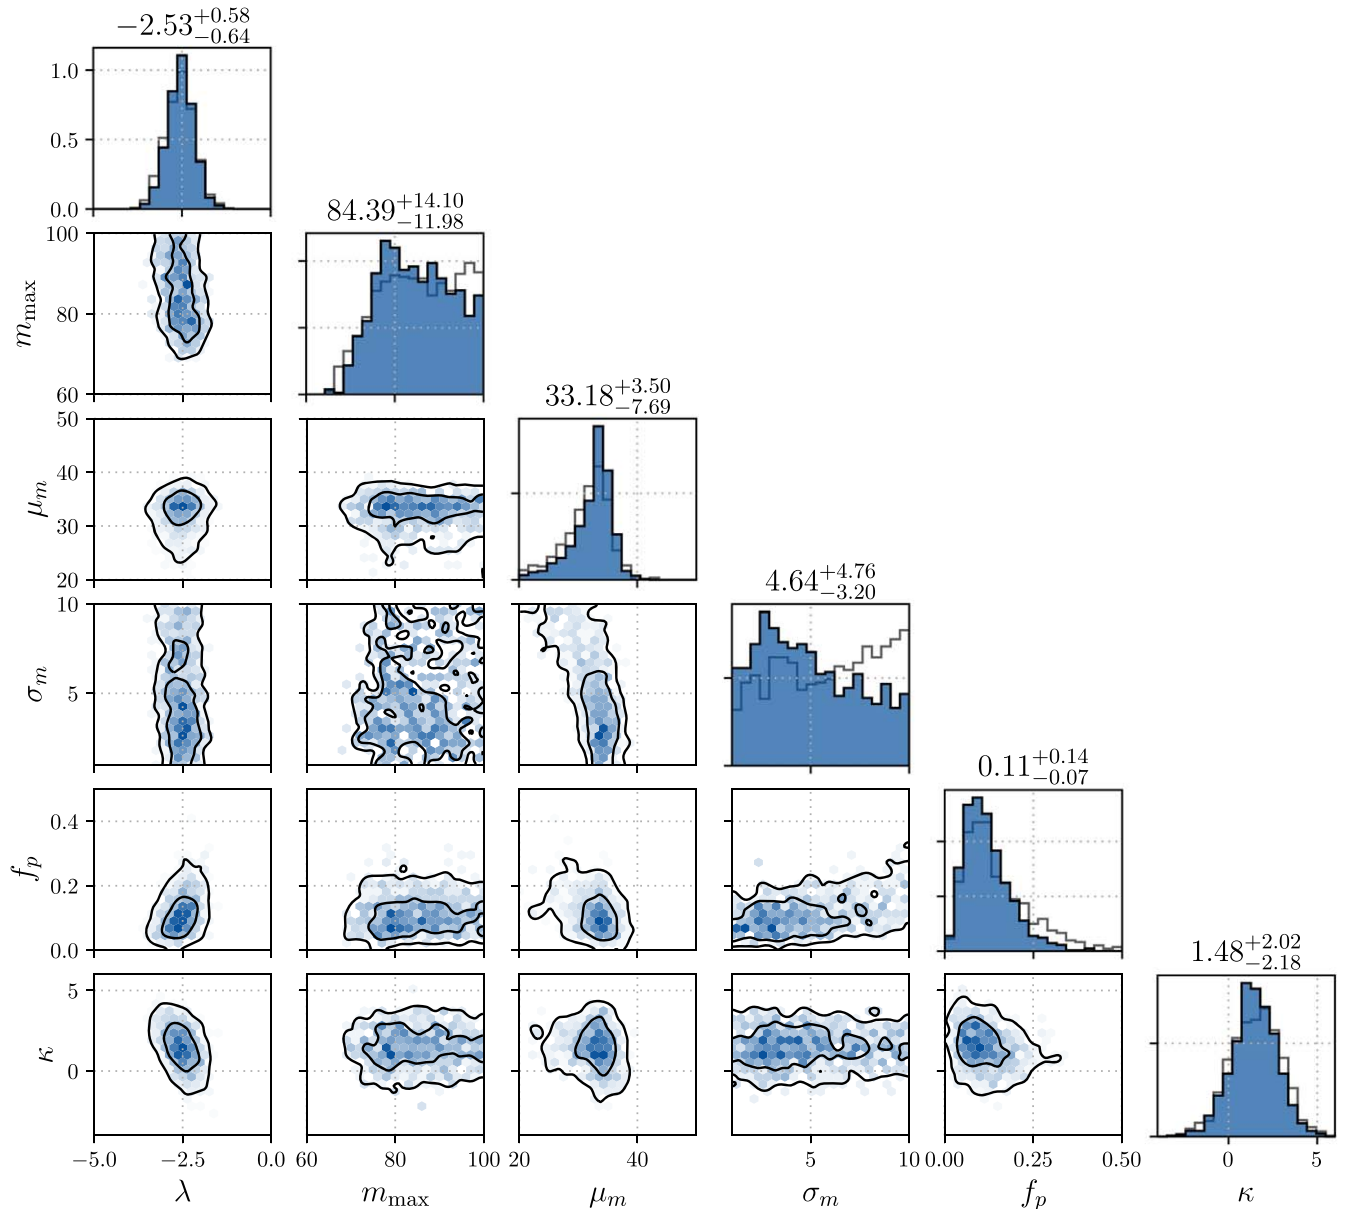


Figure 15. Posteriors on parameters governing the primary mass and redshift distributions of BBHs under the expanded model described in Section 3 that allows for correlations between q and χ_{eff} . Posteriors on the parameters governing the q and χ_{eff} distributions are shown in Figure 3 in the main text. The numbers above each one-dimensional posterior quote medians and central 90% credible uncertainties, and the two contours appearing in each two-dimensional posterior enclose the central 50% and 90% credible regions. For comparison, the unfilled gray histograms show the constraints on these parameters obtained under the base model of Section 2 that neglects any correlations between q and χ_{eff} ; see also Figure 14 above.

Appendix D Recovery of Simulated Events

In Section 4 we described an end-to-end injection study, entailing the parameter estimation and hierarchical analysis of mock signals. This test allowed us to verify that realistic measurement degeneracies between q and χ_{eff} do not seem to generically bias our conclusions regarding underlying correlations between these two parameters. Here, we describe some additional details regarding our generation and analysis of these mock events.

We first randomly draw a large number of BBHs from a mock population with masses between $m_{\text{min}} = 5 M_{\odot}$ and $m_{\text{max}} = 100 M_{\odot}$, redshift evolution described by $\kappa = 2$, and a mass ratio distribution with $\gamma = 0.5$. We assume a broad χ_{eff} distribution

with $\mu_{\chi,0} = 0$, $\sigma_{\chi,0} = 1$, $\alpha = 0$, and $\beta = 0$; note that this is *not* the spin distribution that we ultimately inject and analyze. Since we are concerned only with χ_{eff} and not the individual component spins, we set aligned spin components to $\chi_{1,z} = \chi_{2,z} = \chi_{\text{eff}}$ and in-plane components to zero. We draw primary masses from a broken power law with $p(m_1) \propto m_1^{-2}$ for $m_1 < 35 M_{\odot}$ and $p(m_1) \propto m_1^{-4}$ for $m_1 \geq 35 M_{\odot}$, introducing a deliberate mismatch between the mass distribution of events we inject and the model we eventually assume upon recovery. Events are given random inclinations and sky positions. We use PYCBC (Nitz et al. 2021a) to compute the expected signal-to-noise ratio ρ of each event assuming a network composed of LIGO-Hanford, LIGO-Livingston, and Virgo, using the ‘‘O3 actual’’ noise power spectral densities provided in LIGO Scientific Collaboration &

Virgo Collaboration (2020i) and the IMRPHENOMD aligned spin waveform model (Husa et al. 2016; Khan et al. 2016). We consider events with $\rho \geq 10$ to have been “detected,” and proceed until we have 5×10^4 such events. This set of found events now serves two purposes. First, it provides us a large pool from which to draw a small catalog of mock detections comparable in size to GWTC-2. Second, this pool will allow us to quantify the appropriate selection effects when performing hierarchical inference on our mock catalog, just as the set of pipeline injections described in Appendix A allowed us to correct for selection effects when analyzing GWTC-2.

From our set of 5×10^4 events, we next randomly draw 50 that will comprise our mock injection catalog. In randomly choosing events, each event is assigned a draw weight proportional to $N(\chi_{\text{eff}}|0.05, 0.15)/N(\chi_{\text{eff}}|0, 1)$ so that our catalog corresponds to a much narrower effective spin distribution with $\mu_{\chi,0} = 0.05$ and $\sigma_{\chi,0} = 0.15$ rather than the broad reference distribution used above to generate our initial pool. We use BILBY (Ashton et al. 2019; Romero-Shaw et al. 2020) together with the DYNESTY (Speagle 2020) nested sampler to perform parameter estimation on these 50 events, adding to each signal a random noise realization consistent with the “O3 actual” power spectral densities noted above. Parameter estimation is performed using the IMRPHENOMD aligned spin waveform model (Husa et al. 2016; Khan et al. 2016) with priors that are uniform in detector-frame mass and Euclidean volume [Equation (A4)] and uniform in aligned component spins, corresponding to a χ_{eff} prior given by Equation (A8). This is *not* the same as the χ_{eff} prior implied by a uniform and *isotropic* prior on component spins, the default choice used by the LIGO and Virgo Collaborations, which is much more tightly concentrated about $\chi_{\text{eff}} = 0$. In order to allow for a self-consistent comparison between Figures 1(a) and 9(a), the posterior samples in Figure 9(a) have therefore been reweighted from their original aligned spin prior to the χ_{eff} prior arising from isotropic spins (Equation (10) of Callister 2021a).

ORCID iDs

Thomas A. Callister  <https://orcid.org/0000-0001-9892-177X>
 Carl-Johan Haster  <https://orcid.org/0000-0001-8040-9807>
 Ken K. Y. Ng  <https://orcid.org/0000-0003-3896-2259>
 Salvatore Vitale  <https://orcid.org/0000-0003-2700-0767>
 Will M. Farr  <https://orcid.org/0000-0003-1540-8562>

References

Ade, P. A. R., Aghanim, N., Arnaud, M., et al. 2016, *A&A*, 594, A13
 Akutsu, T., Ando, M., Arai, K., et al. 2021, *PTEP*, 2021, 05A101
 Antonini, F., & Rasio, F. A. 2016, *ApJ*, 831, 187
 Antonini, F., Rodriguez, C. L., Petrovich, C., & Fischer, C. L. 2018, *MNRAS*, 480, L58
 Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, *ApJS*, 241, 27
 Astropy Collaboration, Robitaille, T. P., & Tollerud, E. J. 2013, *A&A*, 558, A33
 Astropy Collaboration, Price-Whelan, A. M., & Sipőcz, B. M. 2018, *AJ*, 156, 123
 Baird, E., Fairhurst, S., Hannam, M., & Murphy, P. 2013, *PhRvD*, 87, 024035
 Bavera, S. S., Fragos, T., Qin, Y., et al. 2020, *A&A*, 635, A97
 Bavera, S. S., Fragos, T., Zevin, M., et al. 2021, *A&A*, 647, A153
 Belczynski, K., Klencki, J., Fields, C. E., et al. 2020, *A&A*, 636, A104
 Biscoveanu, S., Isi, M., Vitale, S., & Varma, V. 2021, *PhRvL*, 126, 171103
 Bouffanais, Y., Mapelli, M., Gerosa, D., et al. 2019, *ApJ*, 886, 25
 Bouffanais, Y., Mapelli, M., Santoliquido, F., et al. 2021, *MNRAS*, 507, 5224

Cahillane, C., Betzwieser, J., Brown, D. A., et al. 2017, *PhRvD*, 96, 102001
 Callister, T., Fishbach, M., Holz, D. E., & Farr, W. M. 2020a, *ApJL*, 896, L32
 Callister, T. A. 2021a, arXiv:2104.09508
 Callister, T. A. 2021b, LIGO Document Control Center, T2100301, <https://dcc.ligo.org/LIGO-T2100301/public>
 Callister, T. A., Farr, W. M., & Renzo, M. 2021b, *ApJ*, 920, 157
 Chia, H. S., Olsen, S., Roulet, J., et al. 2021, arXiv:2105.06486
 Collette, A., Kluyver, T., Caswell, T. A., et al. 2021, h5py/h5py: 3.2.1, v3.2.1, Zenodo, doi:10.5281/zenodo.4584676
 Cutler, C., & Flanagan, E. E. 1994, *PhRvD*, 49, 2658
 Damour, T. 2001, *PhRvD*, 64, 124013
 Di Carlo, U. N., Mapelli, M., Giacobbo, N., et al. 2020, *MNRAS*, 498, 495
 Doctor, Z., Wysocki, D., O’Shaughnessy, R., Holz, D. E., & Farr, B. 2020, *ApJ*, 893, 35
 Estellés, H., Husa, S., Colleoni, M., et al. 2020, arXiv:2105.06360
 Farr, B., Holz, D. E., & Farr, W. M. 2018, *ApJL*, 854, L9
 Farr, W. M. 2021, Reweighting Samples to Account for a Population Distributions, GitHub, <https://github.com/farr/Reweighting>
 Farr, W. M., Stevenson, S., Miller, M. C., et al. 2017, *Natur*, 548, 426
 Fishbach, M., & Holz, D. E. 2017, *ApJL*, 851, L25
 Fishbach, M., Holz, D. E., & Farr, B. 2017, *ApJL*, 840, L24
 Fishbach, M., Holz, D. E., & Farr, W. M. 2018, *ApJL*, 863, L41
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
 Franciolini, G., Baibhav, V., De Luca, V., et al. 2021, arXiv:2105.03349
 Fuller, J., & Ma, L. 2019, *ApJL*, 881, L1
 Gerosa, D., & Berti, E. 2017, *PhRvD*, 95, 124046
 Gerosa, D., Berti, E., O’Shaughnessy, R., et al. 2018, *PhRvD*, 98, 084036
 Gerosa, D., & Fishbach, M. 2021, *NatAs*, 5, 749
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, 585, 357
 Hotokezaka, K., & Piran, T. 2017, arXiv:1707.08978
 Hunter, J. D. 2007, *CSE*, 9, 90
 Husa, S., Khan, S., Hannam, M., et al. 2016, *PhRvD*, 93, 044006
 Khan, S., Husa, S., Hannam, M., et al. 2016, *PhRvD*, 93, 044007
 Kimball, C., Talbot, C., Berry, C. P. L., et al. 2021, *ApJL*, 915, L35
 LIGO Scientific Collaboration, Aasi, J., & Abbott, B. P. 2015, *CQG*, 32, 074001
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, B. P., & Abbott, T. D. 2016, *PhRvL*, 116, 241103
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, B. P., & Abbott, R. 2019a, *PhRvX*, 9, 031040
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, B. P., & Abbott, R. 2019b, *ApJL*, 882, L24
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, R., & Abbott, T. D. 2020a, *PhRvD*, 102, 043015
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, R., & Abbott, T. D. 2020b, *ApJL*, 896, L44
 LIGO Scientific Collaboration, & Virgo Collaboration 2020c, LIGO Document Control Center, P2000217-v2, <https://dcc.ligo.org/LIGO-P2000217/public>
 LIGO Scientific Collaboration, & Virgo Collaboration 2020d, LIGO Document Control Center, P2000434, <https://dcc.ligo.org/LIGO-P2000434/public>
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, R., & Abbott, T. D. 2020e, *PhRvL*, 125, 101102
 LIGO Scientific Collaboration, & Virgo Collaboration 2020f, LIGO Document Control Center, P1800370, <https://dcc.ligo.org/LIGO-P1800370/public>
 LIGO Scientific Collaboration, & Virgo Collaboration 2020g, LIGO Document Control Center, P2000223, <https://dcc.ligo.org/LIGO-P2000223/public>
 LIGO Scientific Collaboration, & Virgo Collaboration 2020h, Gravitational Wave Open Science Center, <https://www.gw-openscience.org/>
 LIGO Scientific Collaboration, & Virgo Collaboration 2020i, LIGO Document Control Center, T2000012, <https://dcc.ligo.org/LIGO-T2000012/public>
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, R., & Abbott, T. D. 2021a, *PhRvX*, 11, 021053
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, R., & Abbott, T. D. 2021b, *ApJL*, 913, L7
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, R., & Abbott, T. D. 2021c, arXiv:2108.01045
 LIGO Scientific Collaboration, & Virgo Collaboration, Abbott, R., & Abbott, T. D. 2021d, *SoftX*, 13, 100658
 LIGO Scientific Collaboration, Virgo Collaboration, & KAGRA Collaboration, Abbott, R., & Abbott, T. D. 2021, *PhRvD*, 104, 022004
 Liu, B., & Lai, D. 2017, *ApJL*, 846, L11
 Liu, B., & Lai, D. 2018, *ApJ*, 863, 68

- Loredo, T. J. 2004, in AIP Conf. Ser., 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. R. Fischer, R. Preuss, & U. V. Toussaint (Melville, NY: AIP), 195
- Mandel, I., Farr, W. M., & Gair, J. R. 2019, *MNRAS*, 486, 1086
- Martinez, M. A. S., Fragione, G., Kremer, K., et al. 2020, *ApJ*, 903, 67
- Mateu-Lucena, M., Husa, S., Colleoni, M., et al. 2021, arXiv:2105.05960
- McKernan, B., Ford, K. E. S., Lyra, W., & Perets, H. B. 2012, *MNRAS*, 425, 460
- McKernan, B., Ford, K. E. S., O’Shaughnessy, R., & Wysocki, D. 2020, *MNRAS*, 494, 1203
- McKernan, B., Saavik Ford, K. E., Bellovary, J., et al. 2018, *ApJ*, 866, 66
- McKinney, W. 2010, in Proc. 9th Python in Science Conf., ed. S. van der Walt & J. Millman, 56
- Miller, S., Callister, T. A., & Farr, W. M. 2020, *ApJ*, 895, 128
- Ng, K. K., Vitale, S., Zimmerman, A., et al. 2018, *PhRvD*, 98, 083007
- Nitz, A., Harry, I., Brown, D., et al. 2021a, gwastro/pycbc: 1.18.0 release of PyCBC, v1.18.0, Zenodo, doi:10.5281/zenodo.4556907
- Nitz, A. H., & Capano, C. D. 2021, *ApJL*, 907, L9
- Nitz, A. H., Capano, C. D., Kumar, S., et al. 2021b, arXiv:2105.09151
- Ohme, F., Nielsen, A. B., Keppel, D., & Lundgren, A. 2013, *PhRvD*, 88, 042002
- Pandas Development Team, T, Reback, J., & Jbrockmendl 2020, pandas-dev/pandas: Pandas, v:latest, Zenodo, doi:10.5281/zenodo.3509134
- Poisson, E., & Will, C. M. 1995, *PhRvD*, 52, 848
- Portegies Zwart, S. F., & McMillan, S. L. W. 2002, *ApJ*, 576, 899
- Pratten, G., García-Quirós, C., Colleoni, M., et al. 2021, *PhRvD*, 103, 104056
- Pürrer, M., Hannam, M., Ajith, P., & Husa, S. 2013, *PhRvD*, 88, 064007
- Pürrer, M., Hannam, M., & Ohme, F. 2016, *PhRvD*, 93, 084042
- Qin, Y., Fragos, T., Meynet, G., et al. 2018, *A&A*, 616, A28
- Rodriguez, C. L., & Antonini, F. 2018, *ApJ*, 863, 7
- Rodriguez, C. L., Zevin, M., Amaro-Seoane, P., et al. 2019, *PhRvD*, 100, 043027
- Romero-Shaw, I. M., Talbot, C., Biscoveanu, S., et al. 2020, *MNRAS*, 499, 3295
- Roulet, J., Chia, H. S., Olsen, S., et al. 2021, *PhRvD*, 104, 083010
- Roulet, J., Venumadhav, T., Zackay, B., Dai, L., & Zaldarriaga, M. 2020, *PhRvD*, 102, 123022
- Roulet, J., & Zaldarriaga, M. 2019, *MNRAS*, 484, 4216
- Speagle, J. S. 2020, *MNRAS*, 493, 3132
- Spruit, H. C. 2002, *A&A*, 381, 923
- Stone, N. C., Metzger, B. D., & Haiman, Z. 2017, *MNRAS*, 464, 946
- Sun, L., Goetz, E., Kissel, J. S., et al. 2020, *CQGrA*, 37, 225008
- Tagawa, H., Haiman, Z., & Kocsis, B. 2020, *ApJ*, 898, 25
- Talbot, C., & Thrane, E. 2018, *ApJ*, 856, 173
- Taylor, S. R., & Gerosa, D. 2018, *PhRvD*, 98, 083017
- Tiwari, V., & Fairhurst, S. 2021, *ApJL*, 913, L19
- Tiwari, V., Fairhurst, S., & Hannam, M. 2018, *ApJ*, 868, 140
- Vallisneri, M., Kanner, J., Williams, R., Weinstein, A., & Stephens, B. 2015, *JPhCS*, 610, 012021
- Venumadhav, T., Zackay, B., Roulet, J., Dai, L., & Zaldarriaga, M. 2020, *PhRvD*, 101, 083030
- Virgo Collaboration, Acernese, F., & Agathos, M. 2015, *CQGrA*, 32, 024001
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Methods*, 17, 261
- Vitale, S., Gerosa, D., Farr, W. M., & Taylor, S. R. 2020, arXiv:2007.05579
- Wong, K. W. K., Breivik, K., Kremer, K., & Callister, T. 2021, *PhRvD*, 103, 083021
- Wysocki, D., Lange, J., & O’Shaughnessy, R. 2019, *PhRvD*, 100, 043012
- Zaldarriaga, M., Kushnir, D., & Kollmeier, J. A. 2018, *MNRAS*, 473, 4174
- Zevin, M., Bavera, S. S., Berry, C. P. L., et al. 2021, *ApJ*, 910, 152