The Binary Black Hole Spin Distribution Likely Broadens with Redshift

Sylvia Biscoveanu[1,2] , Thomas A. Callister[3] , Carl-Johan Haster[1,2] , Ken K. Y. Ng[1,2] , Salvatore Vitale[1,2] , and
Will M. Farr[3,4]
[1] Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA
02139, USA; sbisco@mit.edu
[2] LIGO Laboratory, Massachusetts Institute of Technology, 185 Albany Street, Cambridge, MA 02139, USA
[3] Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA
[4] Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794, USA
Received 2022 April 4; revised 2022 May 19; accepted 2022 May 19; published 2022 June 17

Abstract

The population-level distributions of the masses, spins, and redshifts of binary black holes (BBHs) observed using gravitational waves can shed light on how these systems form and evolve. Because of the complex astrophysical processes shaping the inferred BBH population, models allowing for correlations among these parameters will be necessary to fully characterize these sources. We hierarchically analyze the BBH population detected by LIGO and Virgo with a model allowing for correlations between the effective aligned spin and the primary mass and redshift. We find that the width of the effective spin distribution grows with redshift at 98.6% credibility. We determine this trend to be robust under the application of several alternative models and additionally verify that such a correlation is unlikely to be spuriously introduced using a simulated population. We discuss the possibility that this correlation could be due to a change in the natal black hole spin distribution with redshift.

Unified Astronomy Thesaurus concepts: Astrophysical black holes (98); Gravitational wave astronomy (675); Gravitational wave sources (677)

1. Introduction

The growing catalog of compact binary mergers detected with gravitational waves (Abbott et al. 2019a, 2021a, 2021b) has allowed for increasingly precise characterization of the population properties of binary black holes (BBHs) and neutron stars (Abbott et al. 2021c), which can shed light on how these systems form and evolve (Wong et al. 2020; Bouffanais et al. 2021; Zevin et al. 2021). The latest data from the Advanced LIGO (Aasi et al. 2015) and Virgo (Acernese et al. 2015) observatories have revealed that the BBH mass distribution has substructure (Li et al. 2021; Tiwari & Fairhurst 2021; Veske et al. 2021; Edelman et al. 2022) beyond a smooth power law and single peak or break at $\sim 40\,M_\odot$ (Fishbach & Holz 2017; Talbot & Thrane 2018; Abbott et al. 2019b). BBH spins are found to be small but nonzero (Roulet & Zaldarriaga 2019; Wysocki et al. 2019; Miller et al. 2020; Biscoveanu et al. 2021; García-Bellido et al. 2021), with some of their tilts misaligned to the orbital angular momentum (Talbot & Thrane 2017; Abbott et al. 2019b, 2021c), although these conclusions have been challenged when applying a different population model (Galaudage et al. 2021; Roulet et al. 2021). The merger rate is found to evolve with redshift at a rate consistent with the star formation rate (Madau & Dickinson 2014; Fishbach et al. 2018; Abbott et al. 2021c).

In addition to fitting the mass, spin, and redshift distributions independently, previous works have looked for correlations among these parameters (Safarzadeh et al. 2020; Abbott et al. 2021c; Callister et al. 2021; Franciolini & Pani 2022; Tiwari 2022). Such correlations can be imprinted via evolutionary processes within a single formation channel or can be

caused by the presence of multiple populations arising from distinct formation channels. For example, systems formed dynamically in dense environments via repeated mergers are expected to be more massive and have higher spins (Portegies Zwart & McMillan 2002; Miller & Lauburg 2009; McKernan et al. 2012; Benacquista & Downing 2013; Rodriguez et al. 2015; Antonini & Rasio 2016; Fishbach et al. 2017; Gerosa & Berti 2017; Rodriguez et al. 2019; Kimball et al. 2020; Gerosa & Fishbach 2021). Callister et al. (2021) and Abbott et al. (2021c) found statistically significant evidence for a correlation between the distribution of the mass-weighted spin aligned with the orbital angular momentum, $\chi_{\rm eff}$, and the binary mass ratio, q, where the mean of the $\chi_{\rm eff}$ distribution increases for more extreme mass ratios. This sort of correlation is unexpected for most individual BBH formation models, with the possible exception of formation in the disks of active galactic nuclei (McKernan et al. 2012; Stone et al. 2017; Mckernan et al. 2018; McKernan et al. 2020; Tagawa et al. 2020; Callister et al. 2021) and super-Eddington accretion during stable mass transfer for systems formed via isolated binary evolution (Bavera et al. 2021; Zevin & Bavera 2022), so may hint at the superposition of multiple populations formed via independent channels with unique q versus $\chi_{\rm eff}$ signatures.

Safarzadeh et al. (2020) examined whether the effective spin distribution correlates with various mass parameters of the binary. They found a possible negative correlation between the mean effective spin and each of the chirp mass ($\mathcal{M} = (m_1 m_2)^{3/5}/(m_1 + m_2)^{1/5}$), total mass, and primary mass parameters and a possible positive correlation between the dispersion of the effective spin and mass; neither finding reached high statistical significance ($\sim 80\%$ credibility depending on the mass and correlation parameters).

Abbott et al. (2021c) and Tiwari (2022) also found that the spread in the component spins aligned with the orbital angular momentum may increase with the binary chirp mass. This is

explained as an effect of the paucity of events at large chirp masses, which manifests as a degradation of the constraint on the spin distribution, rather than a firm measurement of an increase in the width of the distribution at larger chirp masses. Franciolini & Pani (2022) also find initial evidence of an evolution of the χ_{eff} distribution with increasing total mass, which they interpret as evidence for a subpopulation of dynamically formed or primordial black hole binaries (De Luca et al. 2020a, 2020b, 2020c; Franciolini et al. 2022), as both models predict correlations between mass and spin.

Previous works have also explored potential correlations between the BBH mass and redshift distributions (Abbott et al. 2021c; Fishbach et al. 2021). Heavier black holes are predicted to form from lower-metallicity stellar progenitors at higher redshifts, which could lead to such a correlation (Mapelli et al. 2019; Neijssel et al. 2019; Belczynski et al. 2020). For systems formed via isolated binary evolution, those that undergo a common envelope event tend to be both less massive and have shorter delay times, merging at higher redshifts (van Son et al. 2022). Motivated by these theoretical predictions, Fishbach et al. (2021) find that redshift dependence of the maximum black hole mass is required if the primary mass distribution has a sharp cutoff, but a gradual tapering of the primary mass distribution is consistent with no evolution with redshift, a finding that was confirmed with the latest catalog of events by Abbott et al. (2021c).

In this work, we search for correlations between the χ_{eff} distribution and the primary mass and redshift distributions. We find robust evidence of a correlation between χ_{eff} and redshift, where the width of the χ_{eff} distribution increases with redshift. We also find a weaker correlation between the width of the χ_{eff} distribution and primary mass. When allowing the χ_{eff} distribution to correlate with both redshift and primary mass, we find a preference in the data for a correlation with one or the other, although we cannot distinguish which. In Section 2, we describe the Bayesian methods and models employed in our analysis. The results on data from the latest GWTC-3 catalog (Abbott et al. 2021b) of compact binaries observed by LIGO-Virgo are presented in Section 3. Section 4 includes a validation of our results using simulated populations and alternative models applied to the real data. We conclude and comment on the potential astrophysical implications of our finding in Section 5.

2. Methods

We employ the framework of hierarchical Bayesian inference to constrain the hyperparameters governing the population-level distributions of the masses, spins, and redshifts of the binary black hole systems detected by LIGO and Virgo. We assume the distribution of primary masses, m_1, is described by the sum of a truncated power law with low-mass smoothing and a Gaussian component (Talbot & Thrane 2018), dubbed the POWER LAW + PEAK model in Abbott et al. (2021c, 2021d), and that the mass ratio distribution is also a power law, bounded such that the minimum and maximum masses of the secondary component are the same as those of the primary (Fishbach & Holz 2020). The merger rate is allowed to evolve with redshift as a power law, following the model in Fishbach et al. (2018) and Abbott et al. (2021c). The hyperparameters governing these mass and redshift distributions are described in Table 1.

Our goal is to determine if there is a correlation between the BBH spin distribution and the distributions of masses or

Table 1
Mass and Redshift Hyperparameter Priors for the Standard POWER-LAW + PEAK and POWER-LAW REDSHIFT Distributions

Parameter	Description	Prior
α	m_1 power-law index	$U(-4, 12)$
β	q power-law index	$U(-4, 7)$
m_{max}	Maximum BH mass	$U(30\,M_\odot, 100\,M_\odot)$
m_{min}	Minimum BH mass	$U(2\,M_\odot, 10\,M_\odot)$
δ_m	Low-mass smoothing parameter	$U(0\,M_\odot, 10\,M_\odot)$
μ_m	PISN peak location	$U(20\,M_\odot, 50\,M_\odot)$
σ_m	PISN peak width	$U(1\,M_\odot, 10\,M_\odot)$
λ	Fraction of systems in PISN peak	$U(0, 1)$
λ_z	z power-law index	$U(-2, 10)$

Table 2
Spin Hyperparameter Priors and Descriptions for the Model in Equation (2)

Parameter	Description	Prior
μ_0	Independent χ_{eff} mean	$U(-1, 1)$
$\log\sigma_0$	Log of independent χ_{eff} width	$U(-1.5, 0.5)$
$\delta\mu_q$	q-dependent χ_{eff} mean	$U(-2.5, 1)$
$\delta\log\sigma_q$	Log of q-dependent χ_{eff} width	$U(-2, 1.5)$
$\delta\mu_z$	z-dependent χ_{eff} mean	$U(-2.5, 1)$
$\delta\log\sigma_z$	Log of z-dependent χ_{eff} width	$U(-0.5, 1.5)$
$\delta\mu_{m_1}$	m_1-dependent χ_{eff} mean	$U(-2.5, 1)$
$\delta\log\sigma_{m_1}$	Log of m_1-dependent χ_{eff} width	$U(-2, 1.5)$

redshifts. To this end, we modify the correlated spin model introduced in Callister et al. (2021) so that the effective aligned spin (Damour 2001; Santamaria et al. 2010; Ajith et al. 2011; Ajith 2011),

$$\chi_{eff} = \frac{\chi_1 \cos\theta_1 + q\chi_2 \cos\theta_2}{1 + q}, \quad (1)$$

is modeled as a truncated Gaussian on $[-1, 1]$ whose mean and variance can each evolve linearly with primary mass and redshift,

$$\pi_{pop}(\chi_{eff}|\Lambda_{\chi_{eff}}, m_1, z) = \mathcal{N}(\chi_{eff}; \mu_\chi, \sigma_\chi), \quad (2)$$

$$\mu_\chi(z, m_1) = \mu_0 + \delta\mu_z(z - 0.5) + \delta\mu_{m_1}\left(\frac{m_1}{10\,M_\odot} - 1\right) \quad (3)$$

$$\log\sigma_\chi(z, m_1) = \log\sigma_0 + \delta\log\sigma_z(z - 0.5) + \delta\log\sigma_{m_1}\left(\frac{m_1}{10\,M_\odot} - 1\right), \quad (4)$$

where we use log to indicate log base 10 everywhere and ln to indicate the natural logarithm. The pivot points of $z = 0.5$ and $m_1 = 10\,M_\odot$ are chosen to be near the peak of the population distributions of those parameters, but we have verified that the exact values do not affect our results. To avoid unphysically narrow distributions using the model above, we impose a cut such that

$$\pi_{pop}(\chi_{eff}|\Lambda_{\chi_{eff}}) = \begin{cases} 0, & \text{if } \log_{10}\sigma_\chi(z, m_1) < -3 \\ \mathcal{N}(\chi_{eff}; \mu_\chi, \sigma_\chi), & \text{elsewhere.} \end{cases} \quad (5)$$

The full set of $\Lambda_{\chi_{eff}}$ parameters are described in Table 2. We note that unlike the model of Callister et al. (2021), our model in Equation (2) does not allow for correlations between χ_{eff} and
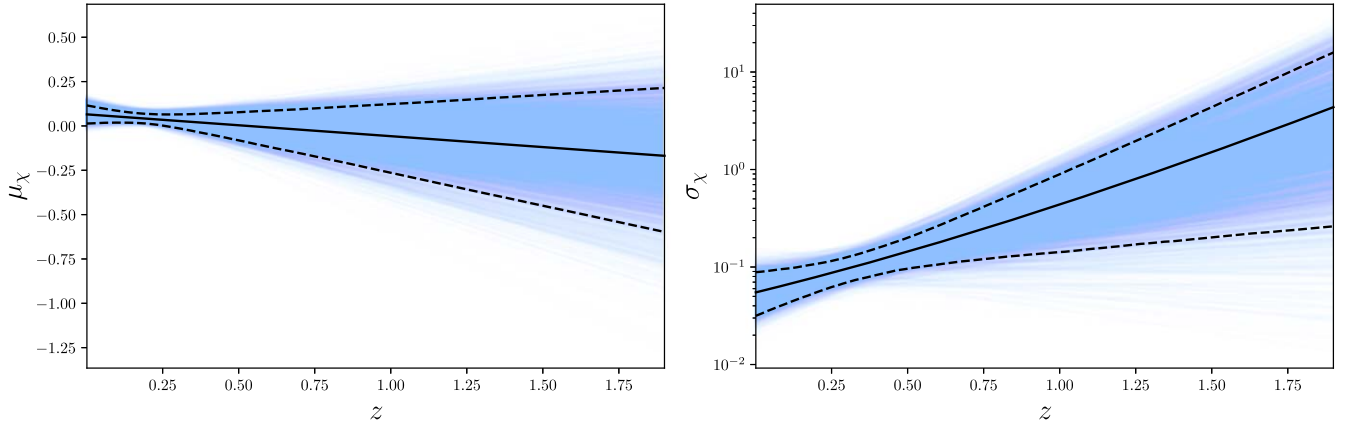
Figure 1. Posteriors for the mean, μ_χ, and standard deviation, σ_χ, of the binary black hole effective spin distribution as a function of redshift obtained for GWTC-3 events. The solid black line shows the mean, while the dashed black lines show the 90% credible region.

mass ratio. However, we amend the model to incorporate mass ratio correlations later in Section 3.

Given a set of posterior samples for the binary parameters θ of N individual BBH events, they can be combined to obtain posteriors on the hyperparameters governing the population-level distributions:

$$p(\Lambda|\{d\}) \propto \mathcal{L}(\{d\}|\Lambda)\pi(\Lambda), \qquad (6)$$

$$\mathcal{L}(\{d\}|\Lambda) \propto \frac{1}{\alpha(\Lambda)^N} \prod_i^N \sum_j \frac{\pi_{\mathrm{pop}}(\theta_{i,j}|\Lambda)}{\pi_{\mathrm{PE}}(\theta_{i,j})}. \qquad (7)$$

The factor of $\pi(\Lambda)$ in Equation (6) represents the prior probability for the hyperparameters, given in Tables 1–2. The likelihood in Equation (7) consists of a Monte Carlo integral over the posterior samples j for each individual event i, where $\pi_{\mathrm{pop}}(\theta_{i,j}|\Lambda)$ is the product of the population distributions for the masses, redshift, and χ_{eff} in Equation (2); $\pi_{\mathrm{PE}}(\theta_{i,j})$ is the original prior that was applied during the parameter estimation of the binary parameters for individual events. Finally, the factor of $\alpha(\Lambda)^N$ accounts for the fact that the observed BBH sources are a biased sample of the underlying astrophysical distribution (Loredo 2004; Mandel et al. 2019; Thrane & Talbot 2019; Vitale 2020). This selection effect arises from the dependence of the sensitivity of the detector network on the intrinsic parameters of the source.

We use the GWPOPULATION package (Talbot et al. 2019) and the DYNESTY nested sampler (Speagle 2020) to obtain samples from the hyperparameter posterior in Equation (6). We include the 69 BBH events reported in GWTC-3 with a false alarm rate (FAR) less than 1 per year (Abbott et al. 2021b), as was done for the BBH analyses in Abbott et al. (2021c). We use the individual-event posterior samples publicly released by LIGO and Virgo (Abbott et al. 2018, 2020a, 2021e, 2021f) obtained with the IMRPhenomPv2 waveform model for events first published in GWTC-1 (Abbott et al. 2019a), and using a combination of different waveform models including the effects of spin precession and higher-order modes for events reported in later catalogs[5] (Abbott et al. 2021a, 2021b, 2021g). We calculate $\alpha(\Lambda)$ following the method described in Farr (2019), using the sensitivity estimates for BBH systems

released at the end of the most recent LIGO-Virgo observing run (O3) obtained via a simulated injection campaign (Abbott et al. 2021h).

3. Results for GWTC-3

We first allow the χ_{eff} distribution to be correlated with redshift and primary mass individually. We recover mass and redshift hyperparameter posteriors consistent with those reported in Abbott et al. (2021c). The evolution of μ_χ and σ_χ as a function of redshift for individual hyperparameter posterior samples obtained with the model only allowing for redshift correlations is shown in Figure 1. In this case the $\delta\mu_{m_1}$ and $\delta\log\sigma_{m_1}$ parameters are fixed to zero. While μ_χ does not exhibit significant evolution with redshift, σ_χ increases with redshift. The corresponding posteriors on $\Lambda_{\chi_{\mathrm{eff}}}$ are shown in Figure 2. The $\delta\mu_z$ posterior is consistent with 0, indicating no evolution of the mean of the χ_{eff} distribution as a function of redshift, but we find that $\delta\log\sigma_z = 0$ is disfavored at 98.6% credibility. We recover $\delta\log\sigma_z = 0.93^{+0.54}_{-0.54}$, (maximum posterior value and 90% credible interval calculated with the maximum posterior density method) indicating that the spin distribution broadens with increasing redshift. This correlation can also be visualized by examining slices of the χ_{eff} distribution at fixed redshift, as shown in Figure 3. The thickness of the 90% credible region for each slice is comparable, but the distribution becomes noticeably broader with increasing redshift. This indicates that the apparent correlation between the width of the spin distribution and the redshift is not dominated by increased uncertainty in the constraint on the distribution at higher redshifts. The location of the peak of the χ_{eff} distribution does not vary significantly between the different redshift slices, consistent with the lack of observed evolution in μ_χ.

While these results present compelling evidence for a correlation between χ_{eff} and redshift, it is possible that the BBH mass and redshift distributions are themselves correlated, so that our recovered redshift correlation is actually a manifestation of a mass versus χ_{eff} correlation viewed through the wrong model. To test this, we now allow the χ_{eff} distribution to be correlated only with primary mass, fixing $\delta\mu_z = 0$, $\delta\log\sigma_z = 0$. The posteriors for the spin hyperparameters for the primary mass correlation model are shown in Figure 4, and the distributions for χ_{eff} along different slices in primary mass are shown in Figure 5. We find a similar but less significant trend between the width of the χ_{eff} distribution and

---

[5] These correspond to the `PublicationSamples`, `PrecessingSpinIMRHM`, and `C01:Mixed` data sets for GWTC-2, GWTC-2.1, and GWTC-3 events, respectively.
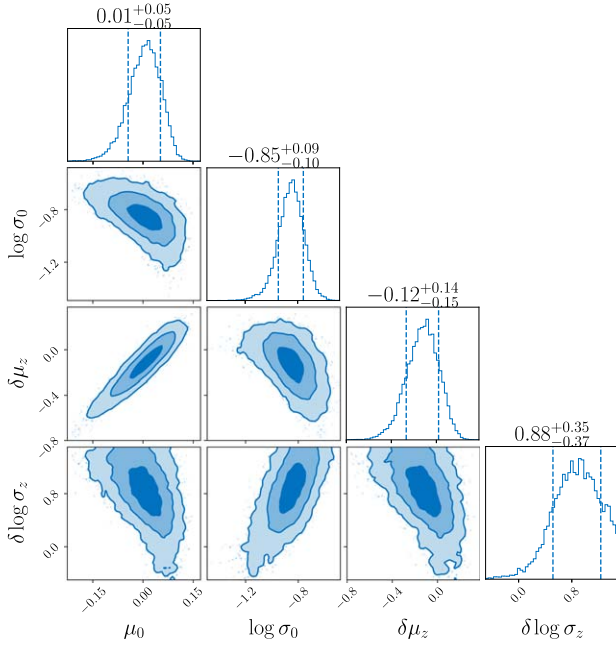
Figure 2. Spin hyperparameter posteriors obtained for GWTC-3 events when allowing for redshift correlations only. The colors indicate the 1, 2, and 3σ 2D credible regions, the dashed lines on the individual histograms show the 1D 1σ credible interval, and the median and 1σ credible interval are printed above each histogram.
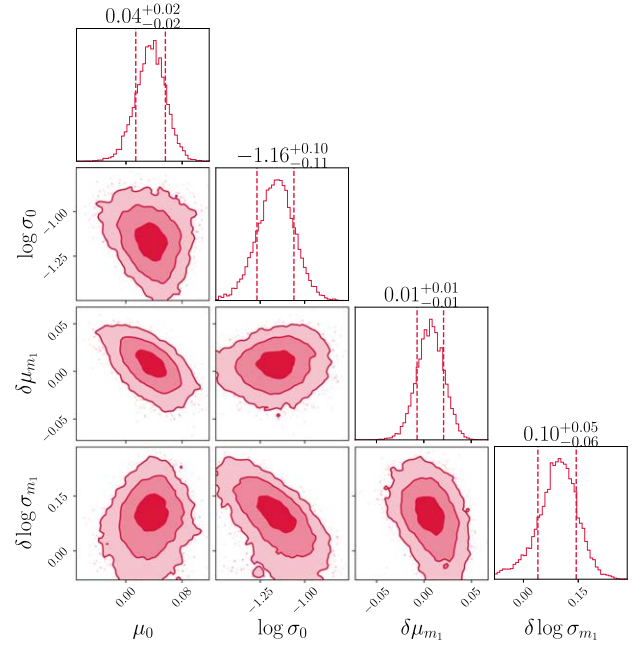

Figure 4. Spin hyperparameter posteriors obtained for GWTC-3 events when allowing for a correlation between $\chi_{\rm eff}$ and primary mass only. The colors indicate the 1, 2, and 3σ 2D credible regions, the dashed lines on the individual histograms show the 1D 1σ credible interval, and the median and 1σ credible interval are printed above each histogram.
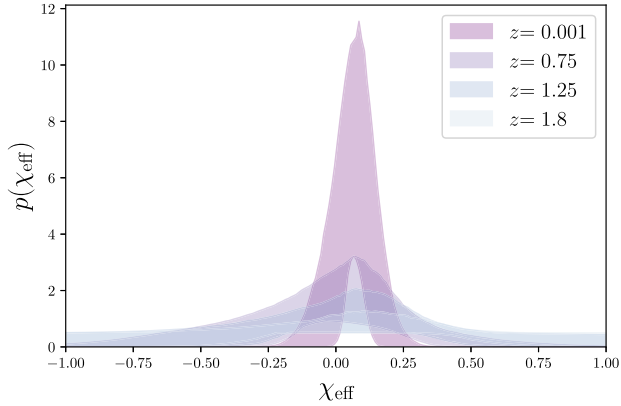

Figure 3. 90% credible region for slices of the $\chi_{\rm eff}$ distribution at four different values of redshift.
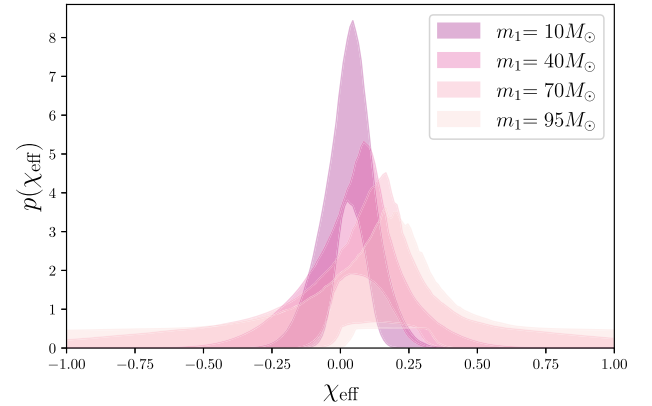

Figure 5. 90% credible region for slices of the $\chi_{\rm eff}$ distribution at four different values of primary mass.

the primary mass, where $\delta\log\sigma_{m_1} = 0.09^{+0.10}_{-0.08}$, and $\delta\log\sigma_{m_1} = 0$ is excluded at 93.5% credibility. While the $\chi_{\rm eff}$ distributions in Figure 5 appear visually to increase both in mean and width with increasing primary mass, the posterior on $\delta\mu_{m_1}$ is still consistent with 0 at 31.1% credibility. In Section 4, we comment on the possibility of a correlation between primary mass and $\chi_{\rm eff}$ being spuriously introduced due to mismodeling a true correlation between redshift and $\chi_{\rm eff}$.

In order to determine if there is a preference in whether the redshift or the primary mass drives the evolution of the width of the $\chi_{\rm eff}$ distribution, we now analyze the data with a model that allows for correlations with both parameters. The posteriors on the spin hyperparameters are shown in Figure 6. We obtain much weaker constraints on $\delta\log\sigma_z$ and $\delta\log\sigma_{m_1}$ individually, but both posteriors have more support for positive than negative values, and the point $\delta\log\sigma_z$, $\delta\log\sigma_{m_1} = (0, 0)$ is disfavored, lying at the 96.2% credibility contour.

Based on these posteriors alone, we cannot confidently identify whether the correlation is dominated by the redshift or the primary mass, but comparing the Bayes factors between the various models that we have considered can provide an indication of which correlation is statistically preferred. Table 3 shows the natural log Bayes factors between the three correlated models we have presented so far and an uncorrelated model, where only μ_0 and $\log\sigma_0$ in Equation (2) are left as free parameters. We repeat the analysis allowing only for redshift correlations with fixed $\delta\mu_z = 0$ in order to gauge the effect of the Occam penalty on the Bayes factor, since this subhypothesis is consistent with our initial redshift correlation result. While none of the values are particularly statistically significant, we find that the models including a correlation with primary mass are disfavored relative to the redshift-only models. We emphasize that while we have ensured consistent priors between all the models, which are subhypotheses of each
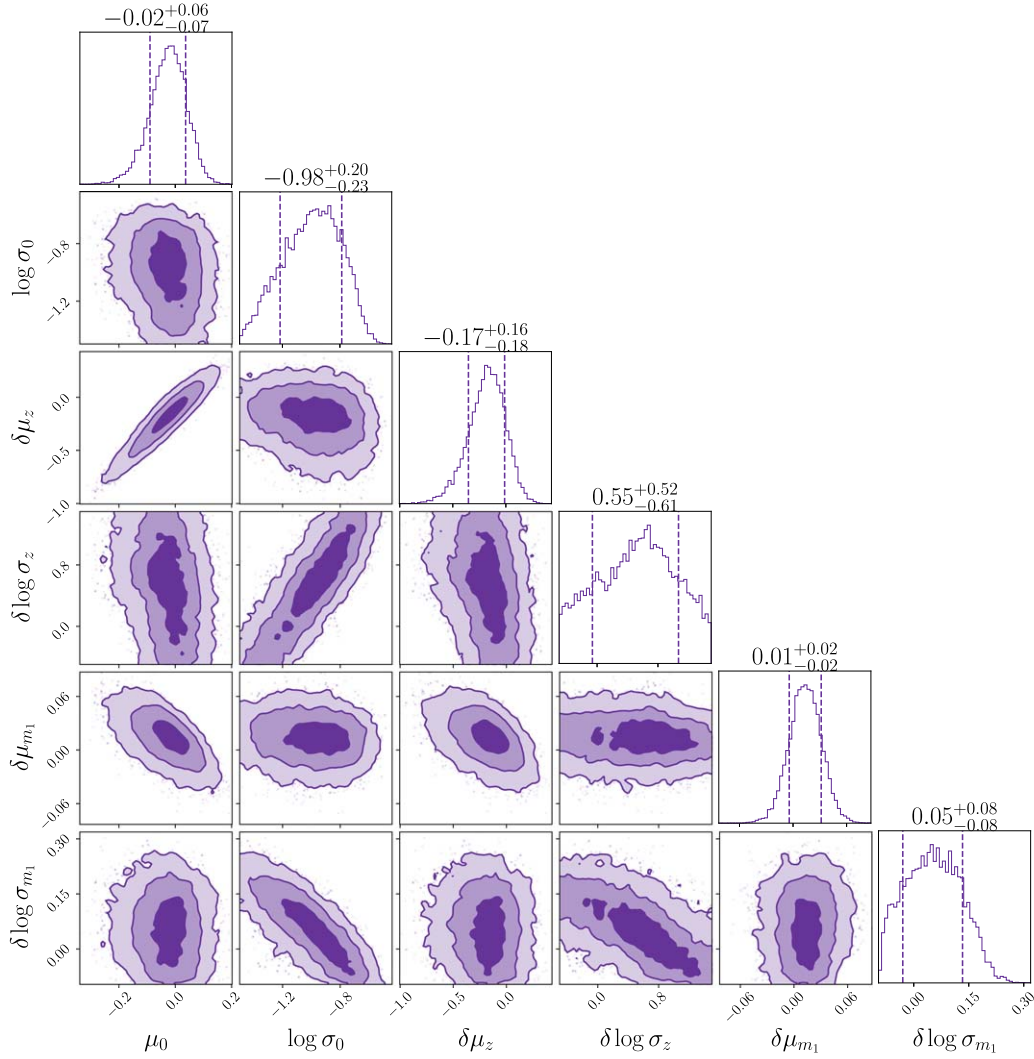
Figure 6. Spin hyperparameter posteriors obtained for GWTC-3 events when allowing for a correlation between $\chi_{\rm eff}$ and both primary mass and redshift. The colors indicate the 1, 2, and 3σ 2D credible regions, the dashed lines on the individual histograms show the 1D 1σ credible interval, and the median and 1σ credible interval are printed above each histogram.

other, the choice of priors for individual hyperparameters is necessarily somewhat arbitrary, complicating the interpretation of the Bayes factors.[6]

Finally, we want to ensure that the apparent correlation between redshift and $\chi_{\rm eff}$ is not falsely introduced by the previously reported correlation between the mean of the $\chi_{\rm eff}$ distribution and mass ratio (Abbott et al. 2021c; Callister et al. 2021). We modify the model in Equation (2) to allow for correlations with both redshift and mass ratio rather than redshift and primary mass:

$$\mu_\chi(z, q) = \mu_0 + \delta\mu_z(z - 0.5) + \delta\mu_q(q - 0.7) \qquad (8)$$

$$\log_{10}\sigma_\chi(z, q) = \log\sigma_0 + \delta\log\sigma_z(z - 0.5)$$
$$+ \delta\log\sigma_q(q - 0.7). \qquad (9)$$

Table 3
Natural Log Bayes Factors Comparing the Models Allowing for Correlations between Redshift, Primary Mass, Mass Ratio, and $\chi_{\rm eff}$ and the Base Model without any Correlations

Correlation	ln(BF)
Redshift, $\delta\mu_z = 0$	1.43
Redshift	0.24
Primary mass	−2.09
Mass ratio	3.75
Redshift & primary mass	−3.63
Redshift & mass ratio	3.20

Note. The redshift-only models include both $\delta\mu_z$ and $\delta\log\sigma_z$ as free parameters and only $\delta\log\sigma_z$ as a free parameter with $\delta\mu_z$ fixed to zero. The models including a correlation with primary mass are disfavored by the data.

The priors on $\delta\mu_q$ and $\delta\log\sigma_q$ are given in Table 2, and the Bayes factor between this model and the uncorrelated model is also shown in Table 3. The results we obtain with this model are consistent with both the previously reported mass ratio correlation and the redshift correlation presented earlier in this work. The $\delta\mu_q$ posterior peaks at negative values, indicating

---

[6] We can use a simple back-of-the-envelope calculation to determine the effect of the narrower width of the prior on $\delta\log\sigma_z$ on the Bayes factor for the redshift correlation models. Since the prior volume for $\delta\log\sigma_z$ is 2/3.5 times the prior volume for $\delta\log\sigma_q$ and $\delta\log\sigma_{m_1}$, the Bayes factor would be increased for the redshift correlation model by $\sim -\ln 2/3.5 \sim 0.56$. This is comparable to the uncertainty in the Bayes factor due to the Monte Carlo integration in the likelihood in Equation (7).

Table 4
True and Recovered Values of the Hyperparameters Describing the Mass, Spin, and Redshift Distributions from which the Simulated Population with No Correlation was Drawn

Parameter	Value	Recovery
α	4	$4.74^{+1.33}_{-0.98}$
β	1.5	$1.02^{+2.10}_{-1.44}$
$m_{\max}$	$50\ M_\odot$	$46.63^{+46.35}_{-7.87}\ M_\odot$
$m_{\min}$	$5\ M_\odot$	$4.99^{+0.87}_{-1.10}\ M_\odot$
δ_m	$5\ M_\odot$	$4.88^{+4.12}_{-2.12}\ M_\odot$
μ_m	$35\ M_\odot$	$31.40^{+1.59}_{-1.59}\ M_\odot$
σ_m	$4\ M_\odot$	$3.09^{+0.99}_{-2.09}\ M_\odot$
λ	0.04	$0.012^{+0.019}_{-0.012}$
λ_z	3	$5.15^{+1.83}_{-1.83}$
μ_0	0.05	$0.025^{+0.076}_{-0.088}$
$\log \sigma_0$	-0.85	$-0.82^{+0.22}_{-0.20}$
$\delta\mu_z$	0	$-0.17^{+0.21}_{-0.11}$
$\delta\log\sigma_z$	0	$-0.24^{+0.62}_{-0.47}$
$\delta\mu_{m_1}$	0	$0.023^{+0.035}_{-0.035}$
$\delta\log\sigma_{m_1}$	0	$-0.14^{+0.07}_{-0.08}$

Note. The correlated population is described by the same hyperparameters with the exception of $\delta\log\sigma_z = 0.85$. The recovered values are represented by the maximum posterior value and 90% credible interval calculated with the maximum posterior density method obtained with the model that allows for both primary mass and redshift correlations

that the mean of the χ_{eff} distribution shifts toward smaller values as the mass ratio becomes more equal. Simultaneously, the $\delta\log\sigma_z$ posterior peaks at positive values, similar to the posteriors shown in Figures 2 and 6. Thus, we conclude that the increase in the width of the χ_{eff} distribution that we report here is independent of the previously established correlation between χ_{eff} and mass ratio.

4. Validation of Results

The results presented so far suggest that for the 69 BBH events we analyze with FAR $< 1\ \mathrm{yr}^{-1}$ detected up to the end of O3, the χ_{eff} distribution broadens with increasing redshift. As noted above, though, it is difficult to disentangle a correlation between spin and redshift from a correlation between spin and mass and to distinguish the extent to which the broadening of the χ_{eff} distribution is due to increased uncertainty in spin measurements at high redshifts. We test the robustness of the observed redshift correlation first with a series of simulated populations, then by analyzing the GWTC-3 data with alternative models allowing for a correlation between χ_{eff} and redshift.

4.1. Simulated Populations

One potential concern is that such a spin-redshift correlation could be introduced due to the degradation of the constraint on χ_{eff} for individual events at farther redshifts, as these sources are generally detected with smaller signal-to-noise ratios (SNRs). To verify whether such a spurious correlation could be introduced, we simulate a population of 69 BBH events detected by a Hanford–Livingston detector network operating at O3 sensitivity (Abbott et al. 2020b) drawn from an intrinsically uncorrelated population. The true values of the hyperparameters describing the population are given in
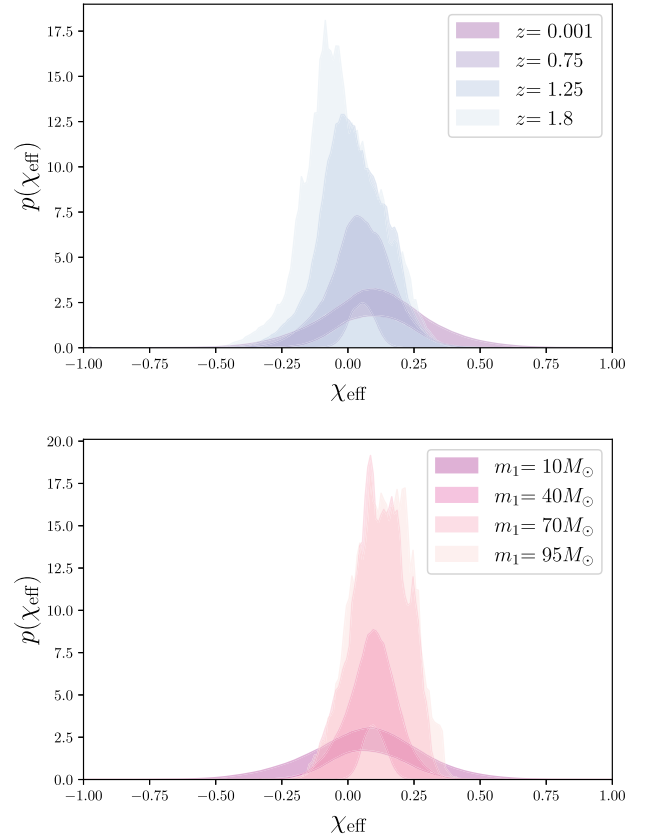

Figure 7. 90% credible region for slices of the χ_{eff} distribution at four different values of redshift (top) and primary mass (bottom) for the simulated population with no correlations.

Table 4. We consider an event to be "detected" if it has a network optimal SNR $\geqslant 9$. We perform individual-event parameter estimation using the reduced order quadrature implementation (Smith et al. 2016) of the IMRPhenomPv2 waveform (Hannam et al. 2014; Husa et al. 2016; Khan et al. 2016) and the DYNESTY nested sampler (Speagle 2020) through the BILBY package (Ashton et al. 2019; Romero-Shaw et al. 2020). We then recover posteriors on the mass, redshift, and spin hyperparameters using the three population models applied to the real data in the previous section.

The true values of all hyperparameters describing this simulated population are recovered within at least the 3σ level, although the posteriors for both $\delta\log\sigma_z$ and $\delta\log\sigma_{m_1}$ tend to prefer negative values. This suggests that it is easier to rule out a distribution that broadens with redshift or mass rather than one that becomes narrower. A broadening distribution implies the presence of highly spinning sources at high mass/redshift, which are easier to detect. Therefore, their absence from the observed population indicates that they are also absent from the underlying population. The converse is not true. If the χ_{eff} distribution instead became narrower, there would be an increasing number of sources with nearly zero spin at high mass/redshift. These sources are harder to detect, so their absence from the observed distribution does not necessarily imply their absence from the underlying distribution.

The χ_{eff} distributions along different slices in redshift and primary mass are shown in Figure 7. These distributions are morphologically distinct from those obtained for the real data in Figures 3 and 5. The distributions obtained for this uncorrelated population appear to initially get narrower

Table 5
Natural Log Bayes Factors Comparing the Models Allowing for Correlations between Redshift, Primary Mass, and $\chi_{\rm eff}$ and the Base Model without any Correlations for the Simulated Population with No Correlations

Correlation	ln(BF)
Redshift	−1.93
Primary mass	−1.59
Both	−3.35

between the lower two values of redshift and primary mass but then broaden again. This indicates that there is no strong evidence for either an increase or a decrease in the width of the $\chi_{\rm eff}$ distribution with either parameter. There is more uncertainty in the distributions at higher primary mass and redshift, a feature that is not observed for the real data, further increasing our confidence in our measurement of a broadening in the spin distribution with mass and/or redshift.

The natural log Bayes factors between each of the models allowing for correlations and the model without any correlations for this simulated population are given in Table 5. In this case, the redshift correlation model is disfavored at a level comparable to the primary-mass-correlation model, as expected since the true population does not contain any correlations. This simulation suggests that a positive correlation between the width of the $\chi_{\rm eff}$ distribution and either the redshift or primary mass is unlikely to be spuriously introduced by the worsening constraint on $\chi_{\rm eff}$ for individual events at higher redshift or primary mass.

We next seek to verify if we are able to successfully recover a correlation in a simulated population similar to the one we find in real data, and whether a redshift correlation can manifest as a primary mass correlation when analyzed with the wrong model. We again generate 69 sources detected at O3 sensitivity now drawn from a distribution with $\delta\log\sigma_z = 0.85$. All the other true hyperparameter values are the same as for the uncorrelated population given in Table 4. We initially analyze this simulated population allowing for mass and redshift correlations in turn. As before, all the hyperparameter values are recovered within 3σ credibility, although the posterior on $\delta\log\sigma_z$ peaks below the true value. The distributions for $\chi_{\rm eff}$ along different slices in mass and redshift shown in Figure 8 are similar to those recovered in the real data. They become consistently wider between increasing values of redshift and primary mass, although they also become more uncertain, which is not observed in the real data. This is because the posteriors for individual events in the simulated data set only extend up to $z = 1.14$, $m_1 = 104\ M_\odot$, while those for real events include values up to $z = 1.90$, $m_1 = 296 M_\odot$, meaning there is more resolving power for the mass and redshift distributions at large masses and redshifts in the real data.

While the posterior on $\delta\log\sigma_{m_1}$ includes 0 within the 81.4% credible level, this simulation indicates that a true correlation between redshift and $\chi_{\rm eff}$ could be perceived as a correlation between primary mass and $\chi_{\rm eff}$ if analyzed with the wrong model. It also reinforces the significance of our finding a preference for positive values of $\delta\log\sigma_z$ in the real data, since both simulations tend to recover smaller values of $\delta\log\sigma_z$ compared to the true value. This is further supported by the Bayes factors we recover for the correlated simulation, given in Table 6. In this case, the recovered correlation with redshift-only is not significant enough to overcome the Occam penalty
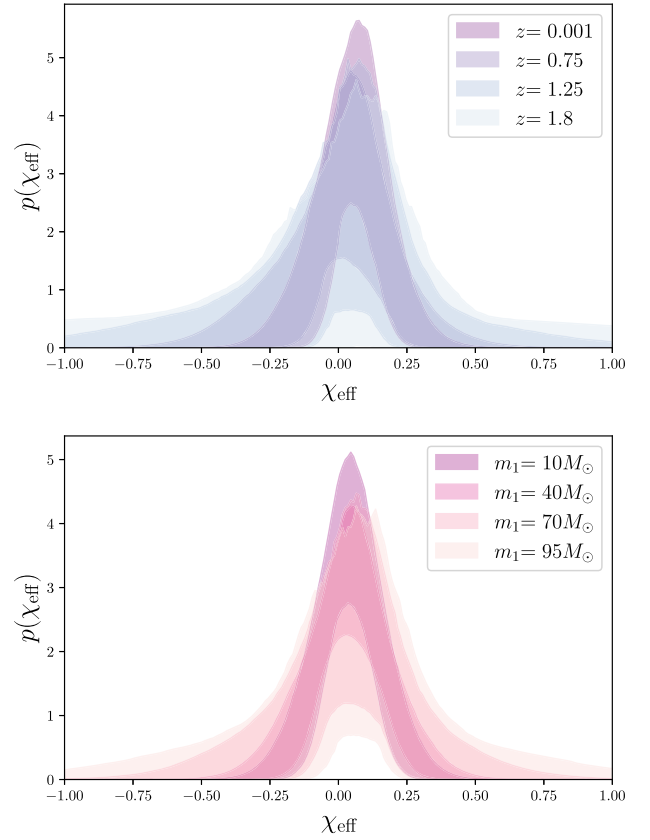

Figure 8. 90% credible region for slices of the $\chi_{\rm eff}$ distribution at four different values of redshift (top) and primary mass (bottom) for the simulated population with a correlation between the width of the $\chi_{\rm eff}$ distribution and redshift.

Table 6
Natural Log Bayes Factors Comparing the Models Allowing for Correlations between Redshift, Primary Mass, and $\chi_{\rm eff}$ and the Base Model without any Correlations for the Simulated Population with $\delta\log\sigma_z = 0.85$

Correlation	ln(BF)
Redshift	−2.45
Primary mass	−4.55
Both	−6.97

for adding parameters relative to the uncorrelated model. However, the relative Bayes factors between the three models considered are similar to those obtained for the real data. This indicates that the redshift-only correlation found in the real data is preferred over the other correlated models we explore at a similar level of statistical significance to the simulated population with a known correlation in $\delta\log\sigma_z$. On the other hand, the relative Bayes factors for the uncorrelated simulation span a much narrower range, revealing that none of the correlated models are strongly preferred over the others in the case of no correlation.

Conversely, it is also possible that a true correlation between primary mass and $\chi_{\rm eff}$ could manifest as a correlation between redshift and $\chi_{\rm eff}$ if analyzed with the wrong model. To verify whether this is a potential false source of the redshift-$\chi_{\rm eff}$ correlation we observe, we simulate a third population of 69 events detected at O3 sensitivity with hyperparameters identical to the previous two, except now $\delta\log\sigma_z = 0$ and $\delta\log\sigma_{m_1} = 0.15$. We find that when this population is analyzed with the model allowing for a redshift correlation only,
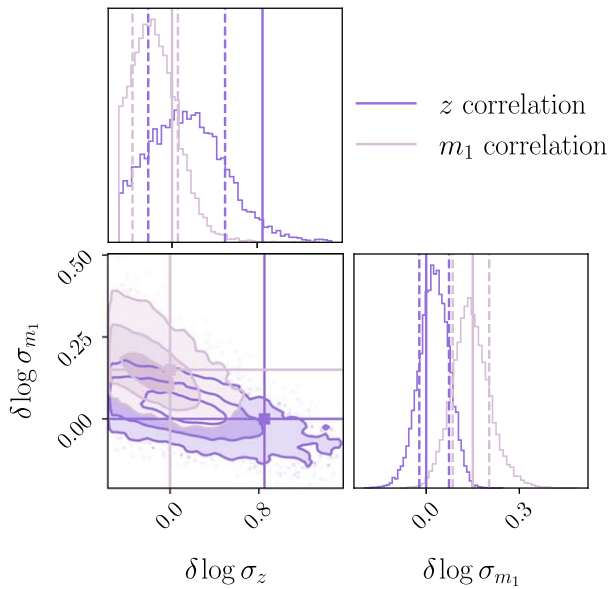
Figure 9. Posteriors for $\delta\log\sigma_{m_1}$, $\delta\log\sigma_z$ obtained under the model allowing for both redshift and primary mass correlations with $\chi_{\rm eff}$, for two different simulated populations. The dark purple shows the results for the population with a simulated redshift correlation, $\delta\log\sigma_z = 0.85$, while the light purple shows the results for the population with a simulated primary mass correlation, $\delta\log\sigma_{m_1} = 0.15$. The shading indicates the 1, 2, and 3σ 2D credible regions and the dashed lines on the individual histograms show the 1D 1σ credible interval. While the posterior for the redshift-correlated population is consistent with primary mass being the sole source of correlation, the opposite is not true.

$\delta\log\sigma_z = 0$ is incorrectly ruled out at 99.0% credibility, indicating that a primary-mass-$\chi_{\rm eff}$ correlation can indeed be confused for a redshift-$\chi_{\rm eff}$ correlation. This is similar to the case explored in the previous two paragraphs where $\delta\log\sigma_{m_1} = 0$ was disfavored at the 81.4% credible level for a true correlation between redshift and spin analyzed with the wrong model.

However, we find in general that a correlation between spin and primary mass is easier to confidently identify with 69 events detected at O3 sensitivity. The posterior on $\delta\log\sigma_{m_1}$ disfavors 0 at 99.7% credibility under the model allowing only for a primary mass correlation and at 99.5% credibility under the model allowing for both primary mass and redshift correlations. This is illustrated in the corner plot of the posteriors on $\delta\log\sigma_{m_1}$, $\delta\log\sigma_z$ in Figure 9 for both the mass-correlated population and the redshift-correlated population previously presented in Figure 8. While it is difficult to experimentally distinguish between mass and redshift evolution of the spin distribution in the case of a redshift-correlated population, the same ambiguity is not present for the mass-correlated population. This further supports the redshift-$\chi_{\rm eff}$ correlation we find in the real data, since a correlation between primary mass and $\chi_{\rm eff}$ would likely have been unambiguously identified.

4.2. Alternative Models

While the results of the two simulated populations lend credence to the finding that the width of the $\chi_{\rm eff}$ distribution increases with redshift in the real data, we want to verify whether this result is model-driven. To this end, we now analyze the GWTC-3 events with a different model—a mixture of truncated Gaussians with a redshift-dependent mixing fraction,

$$\pi(\chi_{\rm eff}|\Lambda_{\chi_{\rm eff}}) = \mathcal{N}(\chi_{\rm eff}; \mu_b, 10^{\sigma_b})(1 - f(z)) \tag{10}$$

$$+ \mathcal{N}(\chi_{\rm eff}; \mu_a, 10^{\sigma_a})f(z), \tag{11}$$

where

$$f(z) = \frac{1}{1 + \exp\left[f_A + f_B(z - 0.5)\right]}. \tag{12}$$

We choose three different priors on the μ and σ parameters for each Gaussian, outlined in Table 7, as proxies for different astrophysical scenarios.

When modeling the effective spin distribution with a single, redshift-independent Gaussian, it is found that the mean effective spin is $\chi_{\rm eff} = 0.06$ (Miller et al. 2020; Abbott et al. 2021c). For our first prior choice, we assume that this result encapsulates the bulk of the population, fixing $\mu_b = 0.06$, and let the second mixture component model the departure from this assumption as we look to higher redshifts. We allow this secondary subpopulation to peak at either positive or negative values of $\chi_{\rm eff}$, but the absolute value of the peak must be $\geqslant 0.1$, so as to minimize the degeneracy between the bulk and this second subpopulation. The evolution of the mixture fraction with redshift for individual hyperparameter posterior samples is shown in Figure 10. The posterior on the mixture fraction differs substantially from the prior, which is symmetric about $f = 0.5$ for all redshifts. This result indicates that the contribution of the Gaussian defined by the parameters μ_a and σ_a becomes more significant as redshift increases. The posterior for μ_a is 3.4 times more likely to be positive than negative, and the positive part of the posterior rails against the boundary at $\mu_a = 0.1$, indicating that it may be trying to replicate the bulk distribution rather than to extract an independent component; σ_a is only weakly constrained to be > -1.02 at 90% credibility. Nonetheless, the $\chi_{\rm eff}$ distributions at different redshift slices shown in the top of Figure 11 for this model and prior choice paint a similar picture to the linear evolution model presented in Section 3; although skewed toward positive values of $\chi_{\rm eff}$, the distributions become wider with increasing redshift.

The second prior choice targets two populations with $\chi_{\rm eff}$ distributions characterized by different widths. The peak of the narrower bulk distribution is now fixed to $\mu_b = 0$ to capture the BBH population formed dynamically with small spin and the majority of field binaries predicted to have negligible spin (Fuller & Ma 2019). Meanwhile the posterior on the peak of the broader Gaussian, μ_a, is largely uninformative. The posteriors for the width parameters and the mixture fraction are very similar to those obtained with the first prior, with $\sigma_a > -0.25$ at 90% credibility and σ_b more strongly constrained to peak at $-1.23^{+0.17}_{-0.17}$. These results do not provide significant evidence for two distinct populations, but the $\chi_{\rm eff}$ distributions at different redshift slices shown in the middle panel of Figure 11 still exhibit some broadening with increasing redshift.

The final prior choice is designed to look for a redshift-dependent excess of preferentially aligned BBH systems. The mean of the bulk distribution is still fixed to $\mu_b = 0$, as expected for a population of dynamically formed systems with isotropic spin tilts (Portegies Zwart & McMillan 2002; Rodriguez et al. 2015; Antonini & Rasio 2016; Gerosa & Berti 2017; Rodriguez et al. 2019; Kimball et al. 2020; Gerosa & Fishbach 2021). The mean of the aligned population is restricted to $0.1 < \mu_a < 1$, in

Table 7
Spin Hyperparameter Priors and Descriptions for the Model in Equation (12)

Parameter	Description	Prior 1	Prior 2	Prior 3
μ_a	Mean of the secondary Gaussian	$U(-1, -0.1) \cup U(0.1, 1)$	$U(-1, 1)$	$U(0.1, 1)$
μ_b	Mean of the bulk Gaussian	0.06	0	0
σ_a	Log-width of secondary Gaussian	$U(-1.5, 0.5)$	$U(-0.7, 1.5)$	$U(-1.5, 0.5)$
σ_b	Log-width of the bulk Gaussian	$U(-1.5, 0.5)$	$U(-2, -0.7)$	$U(-1.5, 0.5)$
$\chi_{a,\mathrm{min}}$	Lower bound of the secondary Gaussian	-1	-1	0
f_A	Independent offset in mixing fraction	$\mathcal{N}(0, 1.5)$	$\mathcal{N}(0, 1.5)$	$\mathcal{N}(0, 1.5)$
f_B	z-dependent offset in mixing fraction	$\mathcal{N}(0, 1.5)$	$\mathcal{N}(0, 1.5)$	$\mathcal{N}(0, 1.5)$
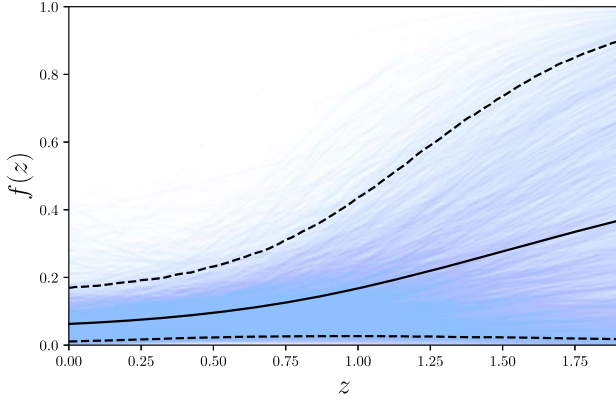
Figure 10. Posteriors for the mixture fraction in Equation (12) obtained for GWTC-3 events under the first prior choice in Table 7. The solid black line shows the mean, while the dashed black lines show the 90% credible region.

order to represent systems formed via isolated binary evolution in the field (Gerosa et al. 2018; Qin et al. 2018; Zaldarriaga et al. 2018; Bavera et al. 2020; Belczynski et al. 2020; Bavera et al. 2021, 2022), and the corresponding Gaussian is truncated on [0,1] instead of [−1, 1]. The χ_{eff} distributions at different redshift slices for this prior choice are shown in the bottom panel of Figure 11. The population model is asymmetric by construction, so the mean of the distribution also increases with redshift along with the width.

The posterior on the mixture fraction is again very similar to the one obtained with the first prior choice in Figure 10. Since the mixture fraction increases with redshift for all three prior choices, the second Gaussian added on top of the bulk distribution always contributes more at high redshifts, consistent with the conclusion that the χ_{eff} distribution broadens with increasing redshift. These results indicate that the apparent correlation between the width of the χ_{eff} distribution and the redshift is not driven by our choice of linear evolution model in Section 2, but can be observed under a variety of population models and priors.

5. Discussion and Conclusion

In this work, we have found weak but robust evidence for an increase in the width of the χ_{eff} distribution of BBH systems with increasing redshift. We have verified that this correlation is unlikely to be spuriously introduced by the increased uncertainty in the χ_{eff} posteriors for individual high-redshift sources using a simulated population and that this trend remains present using alternative population models. We also observe a less significant correlation between the width of the χ_{eff} distribution and the primary mass, although we find that such a correlation can be
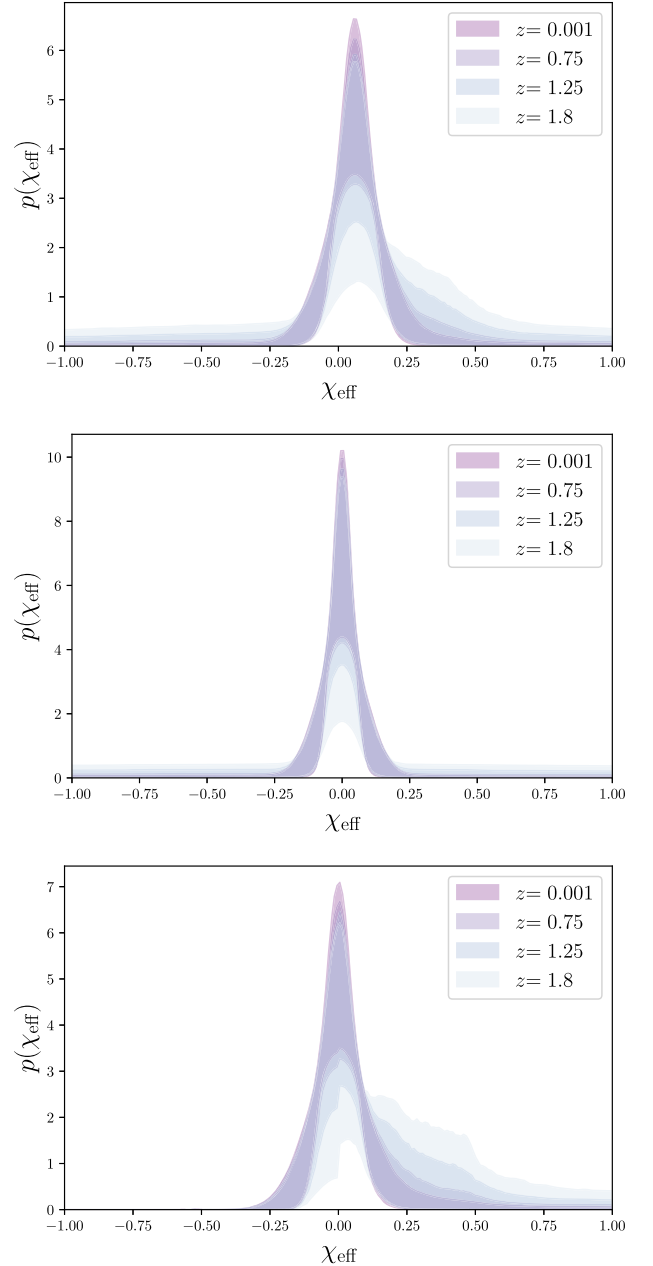

Figure 11. 90% credible region for slices of the χ_{eff} distribution at four different values of redshift for the three different prior choices in Table 7 (1–3 from top to bottom).

falsely recovered if a population with an χ_{eff}-redshift correlation is instead analyzed assuming only an χ_{eff}-primary mass correlation. When allowing for correlations with both redshift and primary

mass, we find evidence for an increase in the width of the χ_{eff} distribution with one or the other but cannot distinguish which.

A correlation between χ_{eff} and redshift might be explained by one of two broad scenarios. First, the correlation could be indicative of multiple subpopulations arising from distinct formation channels, each of which occupy a different region in the (χ_{eff}, z) plane. The third model defined in Section 4.2, for example, serves as a proxy for this scenario. In this case, the symmetric χ_{eff} component centered at zero might serve to capture binaries assembled dynamically in dense stellar environments (Portegies Zwart & McMillan 2002; Rodriguez et al. 2015; Antonini & Rasio 2016; Gerosa & Berti 2017; Rodriguez et al. 2019; Kimball et al. 2020; Gerosa & Fishbach 2021), while the preferentially positive χ_{eff} component corresponds to a subpopulation of systems formed via isolated binary evolution in the field (Gerosa et al. 2018; Qin et al. 2018; Zaldarriaga et al. 2018; Bavera et al. 2020; Belczynski et al. 2020; Bavera et al. 2021, 2022). However, such a mixture between two subpopulations generally leads to a shift in the *mean* of the χ_{eff} distribution with redshift. Under the more generic model explored in Section 3, we much more strongly prefer evolution of the *width* with redshift, although evolution of μ_χ is not strictly ruled out.

The second possibility, broadly, is that a spin-redshift correlation arises due to evolutionary processes operating within a single population. Within isolated binary evolution, for example, tidal interactions may be responsible for correlating black hole spins and redshifts. Some authors predict that isolated black holes are naturally born slowly rotating due to efficient angular momentum transport from stellar cores (Spruit 2002; Fuller et al. 2019; Fuller & Ma 2019). Spin can nevertheless be introduced via tidal torques exerted by the first-born black hole on its companion, prior to the companion's core collapse (Qin et al. 2018; Zaldarriaga et al. 2018; Bavera et al. 2020, 2021; Fuller & Lu 2022). In this scenario, the shortest-period binaries both acquire the largest spins and merge most promptly, correlating the spins of binary black holes and the redshifts at which they merge (Qin et al. 2018; Fuller & Ma 2019). This effect may be enhanced by the lower metallicities present at high redshifts, which diminish the strength of stellar winds and hence prevent binary orbits from widening and avoiding tidal spin-up. Bavera et al. (2022) corroborate this prediction using population synthesis simulations, finding that an increasing fraction of systems undergo tidal spin-up in close orbits at high redshifts. Because there are systems that do not meet the criteria for efficient tidal spin-up at all redshifts, they find that the spin distribution both broadens and increases in mean.

Another possible explanation comes from the effect of metallicity on the efficiency of angular momentum transport in the envelope of the stellar progenitors of the BBH system. Because stellar winds are weaker at low metallicity (higher redshift), this leads to a less efficient removal of the angular momentum stored in the stellar envelope and hence a possibly higher spin of the resulting black hole (Qin et al. 2018; Fuller & Ma 2019). However, this naive picture is complicated by the interplay between the strength of stellar winds, the extent of the hydrogen-burning region inside the star, and the efficiency of elemental mixing within the star. As such, Belczynski et al. 2020 find a nonmonotonic relationship between the metallicity and the black hole spin for a given carbon-oxygen-core mass of

the progenitor. This relationship further depends on the stellar evolution model employed in the binary evolution simulation. Because of the uncertainty in the processes that impart natal spin to the black hole, we cannot place meaningful constraints on these models.

While both of the possibilities described above lead to systems with larger spin magnitudes at increasing redshift, we must also explain the increased number of systems with negative values of χ_{eff} due to the symmetric broadening of the distribution. One potential explanation within the framework of isolated binary evolution is that natal kicks are stronger at higher redshift, leading to more significant misalignment of the BH spin relative to the orbital angular momentum upon birth. It is not clear how such an effect would arise, however. Furthermore, even with moderately strong ($\sim$100 km s^{-1}) natal kicks, the median of the χ_{eff} distribution for systems formed in the field is still much higher than what we recover (Gerosa et al. 2018). Alternatively, the mixing fraction between systems formed via isolated evolution and dynamical assembly may remain constant with redshift, but the spin *magnitudes* of systems formed via both channels can increase due to one or both of the scenarios previously discussed. In this case, the orientations of the spins in systems formed dynamically would remain isotropic, but their magnitudes would be larger at higher redshifts, leading to a broadening of the distribution but not a shift in the mean.

One caveat with our analysis is that we restrict any correlations to be with χ_{eff}, rather than also allowing for mass-redshift correlations. It is likely that the complex processes leading to the formation of BBH systems would produce correlations between all of these parameters, rather than just with spin (Qin et al. 2018; Bavera et al. 2021; Fuller & Lu 2022; Zevin & Bavera 2022). However, it is difficult to encapsulate all possible correlations in a simple phenomenological model without dramatically increasing the dimensionality of the parameter space. We leave the exploration of simultaneous correlations between all parameters to future work.

ORCID iDs

Sylvia Biscoveanu ● https://orcid.org/0000-0001-7616-7366
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

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, CQGra, 32, 074001
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, Parameter estimation sample release for GWTC-1, doi:10.7935/KSX7-QQ51
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, PhRvX, 9, 031040
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, ApJL, 882, L24
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020b, Noise curves used for Simulations in the update of the Observing Scenarios Paper, https://dcc.ligo.org/LIGO-T2000012/public
Abbott, R., Abbott, T. D., Abraham, S., et al. 2020a, GWTC-2 Data Release: Parameter Estimation Samples and Skymaps, https://dcc.ligo.org/LIGO-P2000223/public/
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021d, ApJL, 913, L7
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021g, PhRvX, 11, 021053
Abbott, R., Abbott, T. D., Acernese, F., et al. 2021a, arXiv:2108.01045
Abbott, R., Abbott, T. D., Acernese, F., et al. 2021b, arXiv:2111.03606
Abbott, R., Abbott, T. D., Acernese, F., et al. 2021c, arXiv:2111.03634
Abbott, R., Abbott, T. D., Acernese, F., et al. 2021e, GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run—Parameter Estimation Data Release, vv1, Zenodo, doi:10.5281/zenodo.5117703
Abbott, R., Abbott, T. D., Acernese, F., et al. 2021f, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run—Parameter estimation data release, Zenodo, doi:10.5281/zenodo.5546663
Abbott, R., Abbott, T. D., Acernese, F., et al. 2021h, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run—O1.O2.O3 Search Sensitivity Estimates, Zenodo, doi:10.5281/zenodo.5636816
Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, CQGra, 32, 024001
Ajith, P. 2011, PhRvD, 84, 084037
Ajith, P., Hannam, M., Husa, S., et al. 2011, PhRvL, 106, 241101
Antonini, F., & Rasio, F. A. 2016, ApJ, 831, 187
Ashton, G., Hübner, M., Lasky, P., et al. 2019, ApJS, 241, 27
Bavera, S. S., Fishbach, M., Zevin, M., Zapartas, E., & Fragos, T. 2022, arXiv:2204.02619
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
Benacquista, M. J., & Downing, J. M. 2013, LRR, 16, 4
Biscoveanu, S., Isi, M., Vitale, S., & Varma, V. 2021, PhRvL, 126, 171103
Bouffanais, Y., Mapelli, M., Santoliquido, F., et al. 2021, MNRAS, 507, 5224
Callister, T. A., Haster, C.-J., Ng, K. K. Y., Vitale, S., & Farr, W. M. 2021, ApJL, 922, L5
Damour, T. 2001, PhRvD, 64, 124013
De Luca, V., Franciolini, G., Pani, P., & Riotto, A. 2020a, JCAP, 04, 052
De Luca, V., Franciolini, G., Pani, P., & Riotto, A. 2020b, PhRvD, 102, 043505
De Luca, V., Franciolini, G., Pani, P., & Riotto, A. 2020c, JCAP, 06, 044
Edelman, B., Doctor, Z., Godfrey, J., & Farr, B. 2022, ApJ, 924, 101
Farr, W. M. 2019, RNAAS, 3, 66
Fishbach, M., Doctor, Z., Callister, T., et al. 2021, ApJ, 912, 98
Fishbach, M., & Holz, D. E. 2017, ApJL, 851, L25
Fishbach, M., & Holz, D. E. 2020, ApJL, 891, L27
Fishbach, M., Holz, D. E., & Farr, B. 2017, ApJL, 840, L24

Fishbach, M., Holz, D. E., & Farr, W. M. 2018, ApJL, 863, L41
Franciolini, G., Cotesta, R., Loutrel, N., et al. 2022, PhRvD, 105, 063510
Franciolini, G., & Pani, P. 2022, arXiv:2201.13098
Fuller, J., & Lu, W. 2022, MNRAS, 511, 3951
Fuller, J., & Ma, L. 2019, ApJL, 881, L1
Fuller, J., Piro, A. L., & Jermyn, A. S. 2019, MNRAS, 485, 3661
Galaudage, S., Talbot, C., Nagar, T., et al. 2021, ApJL, 921, L15
García-Bellido, J., Nuño Siles, J. F., & Ruiz Morales, E. 2021, PDU, 31, 100791
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, 8
Hannam, M., Schmidt, P., Bohé, A., et al. 2014, PhRvL, 113, 151101
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., L. Berry, C. P., et al. 2020, ApJ, 900, 177
Li, Y.-J., Wang, Y.-Z., Han, M.-Z., et al. 2021, ApJ, 917, 33
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
Madau, P., & Dickinson, M. 2014, ARA&A, 52, 415
Mandel, I., Farr, W. M., & Gair, J. R. 2019, MNRAS, 486, 1086
Mapelli, M., Giacobbo, N., Santoliquido, F., & Artale, M. C. 2019, MNRAS, 487, 2
McKernan, B., Ford, K. E. S., Bellovary, J., et al. 2018, ApJ, 866, 66
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
Miller, M., & Lauburg, V. M. 2009, ApJ, 692, 917
Miller, S., Callister, T. A., & Farr, W. 2020, ApJ, 895, 128
Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., et al. 2019, MNRAS, 490, 3740
Portegies Zwart, S. F., & McMillan, S. L. W. 2002, ApJ, 576, 899
Qin, Y., Fragos, T., Meynet, G., et al. 2018, A&A, 616, A28
Rodriguez, C. L., Morscher, M., Pattabiraman, B., et al. 2015, PhRvL, 115, 051101
Rodriguez, C. L., Zevin, M., Amaro-Seoane, P., et al. 2019, PhRvD, 100, 043027
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., & Zaldarriaga, M. 2019, MNRAS, 484, 4216
Safarzadeh, M., Farr, W. M., & Ramirez-Ruiz, E. 2020, ApJ, 894, 129
Santamaría, L., Ohme, F., Ajith, P., et al. 2010, PhRvD, 82, 064016
Smith, R., Field, S. E., Blackburn, K., et al. 2016, PhRvD, 94, 044031
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
Tagawa, H., Haiman, Z., & Kocsis, B. 2020, ApJ, 898, 25
Talbot, C., Smith, R., Thrane, E., & Poole, G. B. 2019, PhRvD, 100, 043030
Talbot, C., & Thrane, E. 2017, PhRvD, 96, 023012
Talbot, C., & Thrane, E. 2018, ApJ, 856, 173
Thrane, E., & Talbot, C. 2019, PASA, 36, e010
Tiwari, V. 2022, ApJ, 928, 155
Tiwari, V., & Fairhurst, S. 2021, ApJL, 913, L19
van Son, L. A. C., de Mink, S. E., Callister, T., et al. 2022, ApJ, 931, 17
Veske, D., Bartos, I., Márka, Z., & Márka, S. 2021, ApJ, 922, 258
Vitale, S., Gerosa, D., Farr, W. M., & Taylor, S. R. 2020, Inferring the Properties of a Population of Compact Binaries in Presence of Selection Effects, Handbook of Gravitational Wave Astronomy (Singapore: Springer), 1
Wong, K. W. K., Contardo, G., & Ho, S. 2020, PhRvD, 101, 123005
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. 2022, arXiv:2203.02515
Zevin, M., Bavera, S. S., Berry, C. P. L., et al. 2021, ApJ, 910, 152