



Using Spin to Understand the Formation of LIGO and Virgo’s Black Holes

Ben Farr^{1,2,3,4} , Daniel E. Holz^{2,3,4,5,6} , and Will M. Farr^{4,7}

¹ Department of Physics, University of Oregon, Eugene, OR 97403, USA; bfarr@uoregon.edu

² Enrico Fermi Institute, University of Chicago, Chicago, IL 60637, USA

³ Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

⁴ DARK, Niels Bohr Institute, University of Copenhagen, Blegdamsvej 17, DK-2100 Copenhagen, Denmark

⁵ Department of Physics, University of Chicago, Chicago, IL 60637, USA

⁶ Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA

⁷ Birmingham Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK

Received 2017 November 9; revised 2018 January 22; accepted 2018 January 24; published 2018 February 8

Abstract

With the growing number of binary black hole (BBH) mergers detected by the Advanced LIGO and Virgo detectors, it is becoming possible to constrain the properties of the underlying population and better understand the formation of these systems. Black hole (BH) spin orientations are one of the cleanest discriminators of formation history, with BHs in dynamically formed binaries in dense stellar environments expected to have spins distributed isotropically, in contrast to isolated populations where stellar evolution is expected to induce spins preferentially aligned with the orbital angular momentum. In this work, we propose a simple, model-agnostic approach to characterizing the spin properties of LIGO/Virgo’s BBH population. Using measurements of the effective spin of the binaries, we introduce a simple parameter to quantify the fraction of the population that is isotropically distributed, regardless of the spin magnitude distribution of the population. Once the orientation characteristics of the population have been determined, we show how measurements of effective spin can be used to directly constrain the BH spin magnitude distribution. We find that most effective spin measurements are too small to be informative, with the first four events showing a slight preference for a population with alignment, with an odds ratio of 1.2. We argue that it will be possible to distinguish symmetric and anti-symmetric populations at high confidence with tens of additional detections, although mixed populations may take significantly longer to disentangle. We also derive BH spin magnitude distributions from LIGO’s first four BBHs under the assumption of aligned or isotropic populations.

Key words: binaries: general – gravitational waves – methods: data analysis – stars: black holes

1. Introduction

The direct detection of gravitational waves (GWs) from binary black holes (BBHs) has become almost routine, with three confident events (GW150914, Abbott et al. 2016a; GW151226, Abbott et al. 2016b; GW170104, Abbott et al. 2017) and one candidate event LVT151012 (Abbott et al. 2016c) identified by the Advanced LIGO detectors. The GW signatures of these binaries encode properties of the binary (Veitch et al. 2015) and, in particular, can be used to measure the spin properties of the binary’s black holes (BHs). The astrophysical processes by which these systems form remain uncertain; two generic formation channels include the evolution of an isolated pair of stars that were born together (i.e., “in the field”; Tutukov & Yungelson 1993; Belczynski et al. 2016; Mandel & de Mink 2016; Marchant et al. 2016; Stevenson et al. 2017b) and dynamical interactions in a dense stellar environment (i.e., globular clusters; Sigurdsson & Hernquist 1993; Portegies Zwart & McMillan 2000; Rodriguez et al. 2015, 2016).

Computational models of these formation channels provide predictions for the rates and distributions of binary masses. However, these predictions are highly dependent upon assumptions about poorly understood processes (e.g., common envelope evolution). More robust are the predictions for binary spin properties, particularly the *orientation* of BH spins with respect to the orbital angular momentum of the binary. A generic characteristic of dynamical formation is that the spin orientations of the component BHs are isotropic with respect to the orbital

angular momentum (Sigurdsson & Hernquist 1993; Portegies Zwart & McMillan 2000; Schnittman 2004; Bogdanović et al. 2007; Rodriguez et al. 2015; Rodriguez et al. 2016; Stone et al. 2017). Isolated binaries, on the other hand, are generally expected to be preferentially aligned with the orbital angular momentum (Tutukov & Yungelson 1993; Belczynski et al. 2016; Mandel & de Mink 2016; Marchant et al. 2016; Stevenson et al. 2017b; Zaldarriaga et al. 2018). This conclusion can be weakened through the impact of natal kicks (O’Shaughnessy et al. 2017), but it is difficult for the isolated channel to produce a significant fraction of binary mergers with large misalignment ($\hat{S} \cdot \hat{L} < 0$, where \hat{S} and \hat{L} point along the stellar spin and orbital angular momentum vectors, respectively). In the absence of firm predictions for the full spin distributions (orientation and magnitude) from various formation channels, we provide a model-agnostic approach specially suited for looking for spin isotropy in the BBH population.

LIGO/Virgo’s spin measurements come with large, but well-quantified, uncertainties. Thus, it is difficult to utilize spin to constrain the formation channel of any particular binary, but with many events the properties of the population can start to be inferred. In principle, there are two astrophysically interesting spin quantities for each of the component BHs: the spin magnitude and its misalignment angle with the orbital angular momentum. However, LIGO and Virgo typically provide only one well-constrained spin quantity: the effective spin, χ_{eff} , which is the mass-weighted combination of the aligned components of BH spins

(Abbott et al. 2016d):

$$\chi_{\text{eff}} \equiv \frac{1}{M}(m_1\chi_1 + m_2\chi_2), \quad (1)$$

where $M = m_1 + m_2$ is the total mass of the system, and $-1 < \chi_{1,2} < 1$ are the projections of the spin vectors of each component BH along the orbital axis. As pointed out in Farr et al. (2017), by measuring the effective spin distribution of the population we can infer misalignment characteristics and, consequently, formation channels.

Where Farr et al. (2017) focused on assessing the plausibility of specific model populations through χ_{eff} measurements, we generalize this approach to make model-agnostic statements about the misalignment characteristics of the population. Our approach is based on the basic principle that any isotropic population, regardless of the population’s spin magnitude distribution, must have a symmetric χ_{eff} distribution. Thus, we can use a hierarchical analysis (Hogg et al. 2010; Mandel 2010) of LIGO/Virgo’s detected events to compare the relative rates of binaries with $\chi_{\text{eff}} < 0$ and $\chi_{\text{eff}} > 0$ to assess whether or not the detected population has spins that are isotropically distributed. If ρ is the fraction of systems with $\chi_{\text{eff}} > 0$, then if we consider the possibility of a population that is isotropically distributed, preferentially aligned, or a mixture thereof, the isotropic (or symmetric) fraction is $\sim 2(1 - \rho)$ and the aligned fraction is $\sim 2\rho - 1$. A single definitive measurement of $\chi_{\text{eff}} < 0$ would imply that $\rho < 1$; to date, no such measurement exists. Section 2 presents our hierarchical analysis framework, our proposed models for the χ_{eff} distribution, and a prescription for determining the underlying BH spin magnitude distribution once the orientation distribution has been characterized. Section 3 presents the results from applying these techniques to LIGO’s first four BBH candidates.

2. Methods

2.1. Simulated Populations

To test the sensitivity of the techniques outlined in this Letter we consider several fiducial distributions. Our approach is similar to Farr et al. (2017). We will simulate isotropic and aligned populations with four different component spin magnitude distributions: a “high-spin” distribution with $p(a) = 2a$, a “low-spin” distribution with $p(a) = 2(1 - a)$, a “very-low-spin” population with $a \sim |\mathcal{N}(0, 0.1)|$, and a “very-very-low-spin” population with $a \sim |\mathcal{N}(0, 0.01)|$, where $\mathcal{N}(\text{mean}, \text{standarddeviation})$ represents a normal distribution. For the purposes of simulating aligned populations we assume exact alignment of component spins and the orbital angular momentum, but the techniques described here are general and can be applied to populations with preferential (but not exact) alignment. In particular, unless the processes that misalign an isolated (field) population can produce anti-aligned spin vectors ($\hat{L} \cdot \hat{S} < 0$) the population will still have $\chi_{\text{eff}} > 0$. Posterior constraints on χ_{eff} for each event are simulated to be Gaussian with widths assigned randomly from LIGO’s first four BBH constraints (Abbott et al. 2016c, 2017).

2.2. Hierarchical Population Inference

Ultimately, we would like to calculate $p(\lambda|\{d^i\})$, the posterior density function for the parameters λ describing our population given the data measured around each event $\{d^i\}$.

Since the effective spin cannot be determined precisely from any given observation, we want to marginalize over it, leading to the following expression for the posterior density function for population parameters λ :

$$p(\lambda|\{d^i\}) \propto \left[\prod_{i=1}^{N_{\text{obs}}} \int d\chi_{\text{eff}}^i p(d^i|\chi_{\text{eff}}^i) p(\chi_{\text{eff}}^i|\lambda) \right] p(\lambda), \quad (2)$$

where $p(d^i|\chi_{\text{eff}}^i)$ is the likelihood of measuring data d^i given χ_{eff}^i , $p(\chi_{\text{eff}}^i|\lambda)$ is the probability of observing χ_{eff}^i given population parameters λ , and $p(\lambda)$ is the prior probability of observing population parameters λ . In principle, there is an additional factor that accounts for selection effects. Systems with $\chi_{\text{eff}} > 0$ accumulate more inspiral cycles than $\chi_{\text{eff}} < 0$ systems (Blanchet et al. 1995; Reisswig et al. 2009), making them more likely to be detected. However, this effect only becomes appreciable for large ranges in χ_{eff} ; for GW151226-like binaries ($m_1 \simeq 7.5 M_{\odot}$, $m_2 \simeq 15 M_{\odot}$) the sensitive Volume \times Time, or VT, for systems with $\chi_{\text{eff}} = 0.5$ is $\sim 30\%$ larger than for average systems, and $\sim 70\%$ larger than that of $\chi_{\text{eff}} = -0.5$ systems. We ignore this effect for the purposes of this work, though it should be accounted for as constraints improve, especially if the range of measured χ_{eff} values increases.

Given N_i samples $\{\chi_{\text{eff}}^{ij}\}$ from the marginal posteriors for event i generated using a prior $p(\chi_{\text{eff}})$ we can avoid reanalysis of the data by reweighing the individual event posteriors given a particular population model, leading to a more practical form of the posterior for model parameters

$$p(\lambda|\{d^i\}) \propto \left[\prod_{i=1}^{N_{\text{obs}}} \frac{1}{N_i} \sum_{j=1}^{N_i} \frac{p(\chi_{\text{eff}}^{ij}|\lambda)}{p(\chi_{\text{eff}}^{ij})} \right] p(\lambda). \quad (3)$$

We will make use of this hierarchical posterior density function throughout this work using various population models.

2.3. Two-bin χ_{eff} Model

Since we do not have concrete predictions for the expected distribution of spin parameters for various formation channels, we want a simple yet general approach to characterize the binary population that does not depend on unknown properties, such as the spin magnitude distribution. Binaries that form dynamically, regardless of their spin magnitude distribution, are expected to have a χ_{eff} distribution that is symmetric about 0. A simple way to measure this is to use a two-bin, one-parameter model to describe the χ_{eff} distribution:

$$p(\chi_{\text{eff}}) = \begin{cases} 1 - \rho & -1 \leq \chi_{\text{eff}} < 0 \\ \rho & 0 \leq \chi_{\text{eff}} \leq 1 \end{cases} \quad (4)$$

where ρ can be interpreted as the fraction of systems with $\chi_{\text{eff}} > 0$. Under this model, purely isotropic populations have $\rho = 1/2$, aligned populations have $\rho = 1$, and purely anti-aligned populations ($\chi_{\text{eff}} < 0$) have $\rho = 0$. In this work, we simulate populations with perfect alignment ($\rho = 1$), and while natal kicks are expected to introduce some misalignment in the isolated population, systems with $\chi_{\text{eff}} < 0$ are expected to be rare since this requires torquing the black hole all the way over to retrograde spin. Thus, ρ is expected to be very close to 1 for isolated populations (O’Shaughnessy et al. 2017).




Figure 1. Evolution of the odds ratio in favor of an isotropic distribution (where $\mathcal{O}_{\text{iso}} > 1$ indicates an isotropic distribution is preferred over aligned) for simulated populations using the two-bin model, applied to aligned populations (top) and isotropic populations (bottom). Dashed lines indicate an infinite odds ratio, occurring when a definitive $\chi_{\text{eff}} < 0$ measurement is made. If the population’s spin magnitudes are *very low* ($a \lesssim 0.01$), the two-bin model breaks down and can wrongly conclude that a population is isotropically distributed. This limitation is addressed with the three-bin model introduced in the next section.

If one believes that the BBHs detected by LIGO and Virgo have to either come from a perfectly isotropic population *or* a perfectly aligned population, then the odds ratio in favor of isotropy (assuming equal prior odds) goes as $\mathcal{O}_{\text{iso}} = p(\rho = 1/2)/p(\rho = 1)$. If we consider the possibility of a mixed population of binaries, then the posterior density function for ρ provides a direct measure of the mixing fraction, with $f_{\text{aligned}} = 2\rho - 1$ (Stevenson et al. 2017a; Talbot & Thrane 2017; Vitale et al. 2017).

To test the effectiveness of the two-bin model, we simulate up to 250 events from isotropic and aligned spin populations with observational uncertainties drawn randomly from LIGO’s first four BBHs for the various spin magnitude distributions outlined in Section 2.1. To summarize the performance of this model, Figure 1 shows how the odds ratio \mathcal{O}_{iso} in favor of an isotropic population estimated using the two-bin model evolves as an increasing number of events are detected from each population. For isotropic populations the odds ratio behaves as expected, increasing steadily with the number of detections. For aligned populations with very-very-low-spin, however, the odds ratio in fact favors the wrong conclusion (i.e., the pink curve in the top panel of Figure 1 goes up, not down). Figure 2 shows the posterior constraints on ρ for the very-very-low-spin aligned and isotropic populations, where for the aligned population the simulated value ($\rho = 1$) is incorrectly excluded with high confidence. This shortcoming, due to the mismatch




Figure 2. Posterior distributions for the population parameter ρ , the fraction of systems with $\chi_{\text{eff}} > 0$, after 250 simulated events from the isotropic and aligned, very-very-low-spin ($a \sim |\mathcal{N}(0, 0.01)|$) populations. If the detected population has such low spins, the two-bin model fails (see Section 2.3), motivating the use of the three-bin model (see Section 2.4).

between the population model and the simulated population, is addressed in the following subsection.

For populations with small spins, we have $w \ll \sigma_{\text{obs}} \ll 1$, where w is the characteristic width of the population about $\chi_{\text{eff}} = 0$ (0.01 for the very-very-low-spin population) and σ_{obs} is the typical observational uncertainty of χ_{eff} (~ 0.1 for the events used in this work). In this case, the log posterior for ρ becomes (cf. Equation (2))

$$\begin{aligned} \log p(\rho|\{d^i\}) = \text{const} \\ + \sum_i 2 \left[\left(A_i - \frac{1}{2} \right) (2\rho - 1) - \left(A_i - \frac{1}{2} \right)^2 \right. \\ \left. \times (2\rho - 1)^2 + \mathcal{O} \left(A_i - \frac{1}{2} \right)^3 \right], \end{aligned} \quad (5)$$

where

$$A_i = \frac{\int_0^1 d\chi p(d^i|\chi)}{\int_{-1}^1 d\chi p(d^i|\chi)} \quad (6)$$

is the fraction of the likelihood that supports $\chi_{\text{eff}} > 0$. Thus,

$$\begin{aligned} \log p(\rho|\{d^i\}) \simeq \text{const} \\ + 2N(\mu(2\rho - 1) - \sigma^2(2\rho - 1)^2), \end{aligned} \quad (7)$$

where

$$\mu \equiv \left\langle A - \frac{1}{2} \right\rangle \quad (8)$$

and

$$\sigma^2 \equiv \left\langle \left(A - \frac{1}{2} \right)^2 \right\rangle. \quad (9)$$

The posterior in Equation (7) peaks at

$$\hat{\rho} = \frac{\mu + 2\sigma^2}{4\sigma^2}. \quad (10)$$

For the aligned, very-very-low-spin population considered in this work, $\mu \simeq 0.022$ and $\sigma \simeq 0.29$, leading to $\hat{\rho} \simeq 0.57$, as

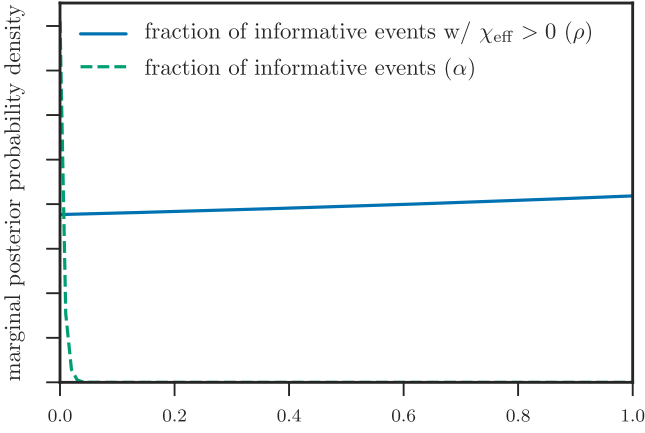


Figure 3. Marginal posterior distributions for the fraction of informative systems α and the fraction of informative systems with $\chi_{\text{eff}} > 0$, ρ , after 250 events for the aligned, very-very-low-spin population. In contrast to Figure 2, almost all systems are uninformative and leave ρ only marginally constrained.

observed in Figure 2, and explaining why the correct (aligned) model is not preferred even as the number of measurements increases in Figure 1. Any population where χ_{eff} is typically much smaller than the observational uncertainty will be subject to the same effect in the two-bin model, motivating our introduction of a three-bin model in the next subsection.

2.4. Three-bin χ_{eff} Model

The two-bin approach outlined in Section 2.3 is a simple approach to answering the question of isotropy. However, systems with $\chi_{\text{eff}} \approx 0$ will fall on the boundary between the two bins and are therefore uninformative about ρ (the fraction of the population with $\chi_{\text{eff}} > 0$). When the typical width of the population is smaller than the observational uncertainty, the two-bin model breaks down, as explained in Section 2.3. To be more robust to the case of low- χ_{eff} systems, we introduce a third bin to the model described in Section 2.3. This additional bin is narrow and centered on $\chi_{\text{eff}} = 0$, leaving us with a two-parameter model:

$$p(\chi_{\text{eff}}) = \begin{cases} 2\alpha(1 - \rho)/(2 - \Delta) & -1 \leq \chi_{\text{eff}} < -\frac{\Delta}{2} \\ (1 - \alpha)/\Delta & -\frac{\Delta}{2} \leq \chi_{\text{eff}} < \frac{\Delta}{2} \\ 2\alpha\rho/(2 - \Delta) & \frac{\Delta}{2} \leq \chi_{\text{eff}} \leq 1 \end{cases}, \quad (11)$$

where Δ is the width of the central bin, α is the fraction of “informative” systems that lie outside of the central (low- χ_{eff}) bin, and ρ is the fraction of those systems with $\chi_{\text{eff}} > 0$. This third, narrow bin accounts for systems with χ_{eff} too small to be informative. By marginalizing over the number of informative systems, we obtain a marginal posterior distribution function $p(\rho)$ that can be interpreted in the same way as the two-bin model. We have experimented with several choices for the central bin width Δ . Since our results do not depend sensitively on this choice, we have used $\Delta = 0.05$ throughout this work; this width is smaller than, but comparable to, the typical observational uncertainty for existing LIGO/Virgo events.

Figure 3 shows the results from applying the three-bin model to the same very-very-low-spin aligned population that failed

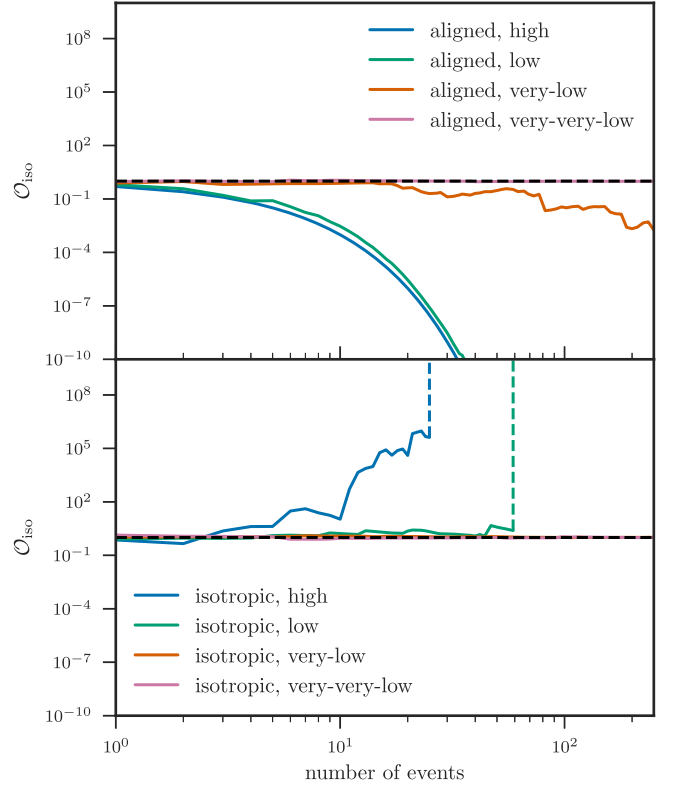


Figure 4. Evolution of the odds ratio in favor of an isotropic population over an aligned population for the various simulated populations using the three-bin model, applied to aligned populations (top) and isotropic populations (bottom). Dashed lines indicate an infinite odds ratio, occurring when a definitive $\chi_{\text{eff}} < 0$ measurement is made. Unlike the two-bin model presented in Figure 1, this model is robust against very-very-low-spin populations.

with the two-bin model (see Figure 1). The marginal posterior distribution for α (the fraction of systems lying outside of the low- χ_{eff} bin) shows a vanishing number of events to be informative, while the marginal posterior distribution for ρ (the fraction of informative systems with $\chi_{\text{eff}} > 0$) remains unconstrained. Figure 4 shows how the odds ratio in favor of an isotropic distribution, calculated from the marginal posterior distribution of ρ , performs on isotropic and aligned populations. While this approach proves to be robust against populations with very low spin magnitudes, it requires more events to be constrained to the same precision as the two-bin model. We find that tens of events will likely be necessary to characterize the underlying population at high confidence, with the precise number depending on the fraction of the population with effective spins indistinguishable from zero. However, for populations predominately composed of systems with very low spins, where the statistical uncertainty of each spin measurement is much greater than the spin distribution of the population, distinguishing between population models may not be possible even with hundreds of events.

2.5. Constraining Component Spins

In general, it is difficult to infer the spin magnitude distribution for the component BHs from the measured distribution of χ_{eff} , since the relation between the two distributions depends on mass ratios and spin alignments. However, using our simple method to identify the alignments of the population can dramatically reduce the complexity of this task. If the population has been determined to be either completely




Figure 5. Mean and 90% credible bounds on the spin magnitude distribution of the low-spin isotropic population after 250 detections, under the assumption that the population is isotropic.

isotropic or aligned, we can then use χ_{eff} measurements to constrain the underlying BH spin magnitude distribution. We will use a three-bin model for the spin magnitude distribution

$$p(a) = \begin{cases} A_1/3 & 0 \leq a < 1/3 \\ A_2/3 & 1/3 \leq a < 2/3, \\ (1 - (A_1 + A_2))/3 & 2/3 \leq a \leq 1 \end{cases} \quad (12)$$

where A_1 and A_2 are the heights of the low ($0 \leq a < 1/3$) and moderate ($1/3 \leq a < 2/3$) spin bins, respectively. We choose three bins here because it seems to be meaningfully constrained with the current data. As we detect more BBHs, we will be able to constrain higher-resolution models. With this model in hand, assuming either an isotropic or aligned distribution of orientations leads to a unique prediction for the distribution of χ_{eff} that can be used in place of the two- and three-bin models described above. As an example, Figure 5 shows the posterior constraints on the spin magnitude distribution after detecting 250 events from a low-spin isotropic population.

3. Results from LIGO’s First Four BBHs

We now make inferences about the spin distributions of the BBH population being detected by LIGO/Virgo. In lieu of true posterior samples from the LIGO/Virgo analyses, which are not publicly available, we approximate the posterior estimates of χ_{eff} for the four likely GW events detected so far (GW150914, GW151226, GW170104, and LVT151012) following a similar prescription to that in Farr et al. (2017). We approximate the posterior as a Gaussian whose central 90% credible interval matches the stated interval for each event (Abbott et al. 2016c, 2017).

Figure 6 shows the marginal posterior density functions from a three-bin analysis of LIGO’s first four likely BBH detections. We find that the information for distinguishing between symmetric and anti-symmetric spin distributions is dominated by GW151226; the rest of the events are consistent with the low- χ_{eff} central bin. Roughly speaking, in the symmetric case a nonzero χ_{eff} measurement would have a 50%–50% change of being aligned ($\chi_{\text{eff}} > 0$) or anti-aligned ($\chi_{\text{eff}} < 0$). In the aligned case, all of these systems would have $\chi_{\text{eff}} > 0$. Thus, finding GW151226 with $\chi_{\text{eff}} > 0$, and no definitive systems with $\chi_{\text{eff}} < 0$, weakly favors aligned versus anti-aligned, if




Figure 6. Marginal posterior distributions for the fraction of informative systems α and the fraction of informative systems with $\chi_{\text{eff}} > 0$, ρ , for the first four likely GW events detected by LIGO, indicating an odds ratio of $\mathcal{O}_{\text{ali}} = 1.2$, i.e., marginal support for an aligned population.




Figure 7. Mean and 90% credible bounds on the spin magnitude distribution inferred from LIGO’s first four likely BBH detections, assuming the population is isotropically distributed.

those are the only two possibilities. Assuming that the current results are “representative,” roughly speaking we expect every sample of ~ 4 events to contain one with $|\chi_{\text{eff}}| \neq 0$ at high confidence. It will rapidly become evident whether these are evenly distributed between positive and negative χ_{eff} values, or whether they are preferentially positive. A single confident measurement of $\chi_{\text{eff}} < 0$ will invalidate the asymmetric population assumption (i.e., will imply $\rho \neq 1$).

If we now assume a particular formation scenario for LIGO/Virgo’s BBHs, we can infer the BH spin magnitude distribution following the approach in Section 2.5. Figure 7 shows the posterior constraints on the spin magnitude distribution assuming LIGO’s first four BBHs were dynamically formed. Figure 8 shows the posterior constraints on the spin magnitude distribution assuming LIGO’s first four BBHs are aligned, the currently preferred scenario. If the population is indeed aligned, then even with just the four candidate BBH mergers detected to date we can already say that there are likely more low-spin systems ($a \lesssim 1/3$) than moderate ($1/3 \lesssim a \lesssim 2/3$) or high ($a \gtrsim 2/3$) spin systems.

A maximal spin ($a \equiv 1$) aligned model is strongly disfavored, with an odds ratio compared to our three-bin model of $\sim 10^{-39}$, while under an isotropic assumption the maximal spin model is disfavored by a modest odds ratio of




Figure 8. Mean and 90% credible bounds on the spin magnitude distribution inferred from LIGO’s first four likely BBH detections, assuming the population is aligned.

0.33; the striking difference between the aligned and isotropic assumptions reflects the difficulty in distinguishing different spin magnitude models under the isotropic assumption (Farr et al. 2017). Our approach considers only χ_{eff} and neglects information about the precessing component of the spins, χ_p . In the case of maximal spins, however, the typically large precessing components of the spins are likely to be measurable, and thus the odds ratio against a maximal, isotropic spin population from the full GW data set is probably larger than the 3:1 value quoted above.

4. Conclusions

We have described a simple, model-agnostic approach for distinguishing dynamically formed and isolated populations from spin measurements (χ_{eff}) of LIGO/Virgo’s events. Our model asks whether the χ_{eff} distribution is symmetric (suggesting a dynamical formation channel), or whether there is evidence for asymmetry (i.e., a preference for aligned versus anti-aligned systems, suggesting isolated formation). We introduce a mixing fraction, ρ , where $\rho = 0.5$ is isotropic, $\rho = 1$ is 100% aligned, and a value in between represents a mixed population. Furthermore, we show that once the alignment of the population has been characterized, the spin magnitude distribution can be directly inferred from the χ_{eff} distribution.

With the first four likely BBH systems observed by LIGO, we find that an aligned formation scenario (i.e., isolated or field formation) is slightly preferred over an isotropic scenario (i.e., dynamical), with an odds ratio in favor of alignment of 1.2. Similarly to Farr et al. (2017), but with a more general model, we find that if all of LIGO’s first four BBHs are assumed to come from this aligned population, then most BH spins must be low ($a \lesssim 1/3$).

Figure 4 shows that ~ 10 additional detections will be sufficient to distinguish between a pure aligned or isotropic population, unless the intrinsic spin magnitude distribution is very low (and GW151226 turns out to be an outlier in spin). LIGO and Virgo are on the cusp of providing important constraints on the formation mechanisms of its BBHs.

We thank Jonathan Gair for valuable discussions. B.F. and D.E.H. were partially supported by NSF CAREER grant PHY-1151836 and NSF grant PHYS-1708081. They were also supported by the Kavli Institute for Cosmological Physics at the University of Chicago through NSF grant PHY-1125897 and an endowment from the Kavli Foundation. The authors thank the Niels Bohr Institute for its hospitality while part of this work was completed and acknowledge the Kavli Foundation and the DNRF for supporting the 2017 Kavli Summer Program.

ORCID iDs

Ben Farr <https://orcid.org/0000-0002-2916-9200>
 Daniel E. Holz <https://orcid.org/0000-0002-0175-5064>
 Will M. Farr <https://orcid.org/0000-0003-1540-8562>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *PhRvL*, **116**, 061102
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *PhRvL*, **116**, 241103
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016c, *PhRvX*, **6**, 041015
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016d, *PhRvL*, **116**, 241102
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvL*, **118**, 221101
- Belczynski, K., Holz, D. E., Bulik, T., & O’Shaughnessy, R. 2016, *Natur*, **534**, 512
- Blanchet, L., Damour, T., Iyer, B. R., Will, C. M., & Wiseman, A. G. 1995, *PhRvL*, **74**, 3515
- Bogdanović, T., Reynolds, C. S., & Miller, M. C. 2007, *ApJL*, **661**, L147
- Farr, W. M., Stevenson, S., Miller, M. C., et al. 2017, arXiv:1706.01385
- Hogg, D. W., Myers, A. D., & Bovy, J. 2010, *ApJ*, **725**, 2166
- Mandel, I. 2010, *PhRvD*, **81**, 084029
- Mandel, I., & de Mink, S. E. 2016, *MNRAS*, **458**, 2634
- Marchant, P., Langer, N., Podsiadlowski, P., Tauris, T. M., & Moriya, T. J. 2016, *A&A*, **588**, A50
- O’Shaughnessy, R., Gerosa, D., & Wysocki, D. 2017, *PhRvL*, **119**, 011101
- Portegies Zwart, S. F., & McMillan, S. L. W. 2000, *ApJL*, **528**, L17
- Reisswig, C., Husa, S., Rezzolla, L., et al. 2009, *PhRvD*, **80**, 124026
- Rodriguez, C. L., Morscher, M., Pattabiraman, B., et al. 2015, *PhRvL*, **115**, 051101
- Rodriguez, C. L., Zevin, M., Pankow, C., Kalogera, V., & Rasio, F. A. 2016, *ApJL*, **832**, L2
- Schnittman, J. D. 2004, *PhRvD*, **70**, 124020
- Sigurdsson, S., & Hernquist, L. 1993, *Natur*, **364**, 423
- Stevenson, S., Berry, C. P. L., & Mandel, I. 2017a, arXiv:1703.06873
- Stevenson, S., Vigna-Gómez, A., Mandel, I., et al. 2017b, *NatCo*, **8**, 14906
- Stone, N. C., Metzger, B. D., & Haiman, Z. 2017, *MNRAS*, **464**, 946
- Talbot, C., & Thrane, E. 2017, arXiv:1704.08370
- Tutukov, A. V., & Yungelson, L. R. 1993, *MNRAS*, **260**, 675
- Veitch, J., Raymond, V., Farr, B., et al. 2015, *PhRvD*, **91**, 042003
- Vitale, S., Lynch, R., Sturani, R., & Graff, P. 2017, *CQGrA*, **34**, 03LT01
- Zaldarriaga, M., Kushnir, D., & Kollmeier, J. A. 2018, *MNRAS*, **473**, 4174