



Constraining Black Hole Spins with Gravitational-wave Observations

Vaibhav Tiwari¹, Stephen Fairhurst¹, and Mark Hannam¹Cardiff School of Physics and Astronomy Cardiff University, Queens Buildings, The Parade Cardiff CF24 3AA, UK; tiwariv@cardiff.ac.uk

Received 2018 July 26; revised 2018 October 13; accepted 2018 October 14; published 2018 December 3

Abstract

The observation of gravitational-wave signals from merging black hole binaries enables direct measurement of the properties of the black holes. An individual observation allows measurement of the black hole masses, but only limited information about either the magnitude or orientation of the black hole spins is available, primarily due to the degeneracy between measurements of spin and binary mass-ratio. Using the first six black hole merger observations, we are able to constrain the distribution of black hole spins. We perform model selection between a set of models with different spin population models, combined with a power-law mass distribution, to make inferences about the spin distribution. We assume a fixed power-law mass distribution for the black holes, which is supported by the data and provides a realistic distribution of the binary mass-ratio. This allows us to accurately account for selection effects due to variations in the signal amplitude with spin magnitude, and provides an improved inference on the spin distribution. We conclude that the first six LIGO and Virgo observations disfavor highly spinning black holes against low spins by an odds ratio of 15:1, thus providing strong constraints on spin magnitudes from gravitational-wave observations. Furthermore, we are able to rule out a population of binaries with completely aligned spins, even when the spins of the individual black holes are low, at an odds ratio of 22000:1, significantly strengthening earlier evidence against aligned spins. These results provide important information that will aid in our understanding of the formation processes of black holes.

Key words: black hole physics – gravitational waves

1. Introduction

Gravitational waves (GWs) emitted by merging black holes are identified in the LIGO and Virgo data through the use of search analysis pipelines, which use the known waveform morphology to identify weak signals in the data (Abbott et al. 2016a, 2017a, 2017b, 2017c). The observations are followed by parameter estimation analyses that extract posterior probability distributions for the parameters of the binary—the masses and spins of the component black holes as well as the distance, sky location, and orientation of the binary (Cutler & Flanagan 1994; Veitch et al. 2015; Abbott et al. 2016b). While some parameters are extracted with good precision, others cannot be accurately measured, and several sets of parameters are strongly correlated, such as, for example, the distance with binary orientation, and mass-ratio with black hole spins. Nonetheless, the observed parameters from several observations can be combined to obtain the underlying astrophysical distributions of black hole masses and spins. In this paper we use publicly available measurements from the first six GW signals observed from merging black holes (GW150914; LVT151012; GW151226 and GW170104; and GW170608 and GW170814; respectively in Abbott et al. 2016a, 2017a, 2017b, 2017c) to draw inferences about the underlying spin distribution of black holes.¹

1.1. Model Selection

For much of the parameter space it is not possible to accurately measure the individual spins with LIGO-Virgo observations at typical SNRs (Pürrer et al. 2016), but only two mass-weighted combinations of the spins: an effective spin, χ_{eff} , which describes the effect the aligned components have on the inspiral rate at which the binary’s orbit decays, and the

maximum orbital frequency prior to merger (Ajith et al. 2011), and an effective precessing spin, χ_p , which determines the effects of precession and leads to modulations in the observed waveform (Apostolatos et al. 1994; Schmidt et al. 2015). For the reported observations, there has been no discerning evidence of precession (the posterior distributions for χ_p are consistent with the priors). Consequently, we restrict attention to the observed values of χ_{eff} .

The effective spin used in LIGO-Virgo analyses is related to the individual spins of the two black holes in the binary by (Damour 2001; Ajith et al. 2011)

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

where m_1 and m_2 are the component masses of the binary and χ_1 and χ_2 are the components of the (dimensionless) spins aligned with the angular momentum defined as $\chi = c\mathbf{S} \cdot \hat{\mathbf{L}} / (Gm^2) \leq 1$.

We are interested in inferring the spin distribution of the merging black hole binaries in the Universe by comparing the observed distribution of effective spins with those predicted by astrophysical models. At present, we have only a limited number of observations, therefore we focus on a discrete set of models. We follow (Farr et al. 2017) and introduce six possible spin distributions for comparison. We use three distributions for spin magnitude:

$$\begin{aligned} p_{\text{low}}(a) &= 2(1 - a) \\ p_{\text{flat}}(a) &= 1 \\ p_{\text{high}}(a) &= 2a, \end{aligned} \quad (2)$$

where $a = |c\mathbf{S} / (Gm^2)|$ is the spin magnitude, ranging from zero to one. In addition, we consider both isotropic and aligned-spin models (Vitale et al. 2014, 2017a; Farr et al. 2017;

¹ We generate the parameter distribution using the confidence intervals reported in observation papers—please see Section 2 for more details.

Stevenson et al. 2017; Talbot & Thrane 2017). In the isotropic model each of the component spins is isotropically distributed on the sphere. In the aligned-spin model, the spin is assumed to be aligned with the orbital angular momentum, i.e., $\chi \equiv a$. We assume a specific distribution of all the parameters, varying only the distribution of the effective spin, χ_{eff} (as defined in Equation (1)) between the models and perform model selection between these to determine which of the models are preferred by the observations to date. The figures in this paper use acronyms for model names based on their initials, i.e., isotropic models have initial (I) and aligned models have initial (A).

1.2. Two Primary Effects

Before we proceed to model selection we must consider two primary effects black hole spins have on the waveform.

First, we must note the degeneracy between measurements of the effective spin parameter χ_{eff} and the mass-ratio $q = m_2/m_1$ of the binary (where we require $m_2 \leq m_1$ so that $q \leq 1$). A significant source of measurement error in χ_{eff} is due to the partial degeneracy between mass-ratio and effective spin. While the two black holes are in a slowly inspiralling orbit, the emitted waveform is well approximated by the post-Newtonian approximation. There, the dominant term that affects the phase evolution of a binary is its chirp mass, $\mathcal{M} = m_1^{3/5} m_2^{3/5} (m_1 + m_2)^{-1/5}$. The following terms are dependent upon the binary's mass-ratio and the black hole spins (Cutler & Flanagan 1994; Poisson & Will 1995; Baird et al. 2013). At this order, a high χ_{eff} -low q binary is indistinguishable from a low χ_{eff} -high q binary. Consequently, for lower-mass black hole binaries, where the majority of the gravitational-wave signal is observed in the inspiral, the effective spin and mass-ratio will be degenerate. For higher-mass binaries, only the last few orbits of the binary are observed and the merger and ringdown of the binary provide a significant contribution to the observed signal. The merger and ringdown parts of the waveform depend primarily on the total mass of the system. Thus, for higher-mass binaries, the degeneracy between mass-ratio and spin is less pronounced (Haster et al. 2016). The spin distribution can be obtained by marginalizing over the joint estimate of the mass and spin distribution $p(m_1, m_2, \chi_1, \chi_2)$. Thus, given the degeneracies between measurements of spins and masses, the assumed mass distribution of binaries in the universe will have a significant impact upon the inferred spins.

Second, there is a selection bias affecting inference of the spin population. The overall amplitude of the emitted gravitational wave depends upon both the masses and effective spin of the binary. The amplitude of the gravitational wave scales with the mass of the system. In addition, a spin-orbit coupling causes binaries with a positive χ_{eff} to undergo a greater number of orbits prior to merger and emit larger amplitude gravitational waves, compared to those with a zero or negative χ_{eff} (Campanelli et al. 2006). For example, a $40 M_{\odot}$ - $30 M_{\odot}$ binary, with maximal aligned spins, can be observed to a distance 1.6 times as large as one with maximal anti-aligned spins, leading to a factor of four increase in the observed rate for a population uniformly distributed in volume. We are biased toward observing high- χ_{eff} binaries, so the fact that binaries with high, aligned spins can be observed at greater distances means that these will be preferentially observed, and their non-observation after six detections suggests that

high- χ_{eff} binaries are rare. Moreover, the increase of the distance at the which binary is observable, caused due to an increase in the value of χ_{eff} , is also mass-dependent. Heavier binaries get a greater push in distances to which they can be observed than the lighter ones. Thus, selection effects will also depend on the assumed mass distribution.

Ideally, a combined analysis of masses and spins will naturally account for these effects. A flexible non-parametric prior that maximizes the overall probability of observing all the gravitational-wave signals can be used to obtain the parameter distributions. Such an analysis will require hundreds of events (see also Wysocki et al. 2018, which constructs a phenomenological distribution with a limited number of gravitational-wave observations).

At present, we have only a limited number of observations and therefore we account for these effects by imposing an astrophysically motivated mass distribution on the universe: $p(m_1) \propto m_1^{-2.3}$ with m_2 uniformly distributed between $5 M_{\odot}$ and m_1 . The choice is based on astrophysical observations that support the stellar initial mass function to follow a power-law distribution. Moreover, the power-law model provides a binary mass-ratio distribution supported by the population synthesis models Dominik et al. (2012) and Rodriguez et al. (2016). Furthermore, independent of the assumed spin distribution, the GW measurement also supports the power-law model of Abbott et al. (2017a). In summary, the power-law model is among other simple models that are supported by the data.

2. Method

Using the observed measurements of the effective spin from the six BBH mergers considered here, we use Bayesian model selection to calculate the odds ratio between the different models. While model selection is quite standard, care must be taken to ensure that the selection effects and mass priors are correctly incorporated; see also Loredo (2004) and Mandel et al. (2016). We are interested in calculating

$$p(\lambda|\{\mathbf{d}\}) = \frac{p(\{\mathbf{d}\}|\lambda)p(\lambda)}{p(\{\mathbf{d}\})}, \quad (3)$$

where $\{\mathbf{d}\}$ denotes the set of observations, $p(\lambda)$ is the prior on the model λ and $p(\{\mathbf{d}\})$ is formally given as the integral over λ of the numerator:

$$p(\{\mathbf{d}\}) = \int d\lambda p(\{\mathbf{d}\}|\lambda)p(\lambda). \quad (4)$$

Since the model, λ , gives a distribution for the parameters of the signal, θ , we can express the probability of obtaining a given data set \mathbf{d} corresponding to a single observation as

$$p(\mathbf{d}|\lambda) = \int d\theta p(\mathbf{d}|\theta)p(\theta|\lambda), \quad (5)$$

where the distribution of \mathbf{d} , given θ , is calculated from the Gaussian likelihood as

$$p(\mathbf{d}|\theta) \propto \prod_{X \in \text{dets}} \exp\left[-\frac{1}{2}\langle d_X - h_X(\theta) | d_X - h_X(\theta) \rangle\right], \quad (6)$$

where d_X denotes the data in detector ‘‘X’’ and $h_X(\theta)$ is the gravitational waveform expected in detector ‘‘X’’ from a binary with parameters θ . The product is over detectors in the network; $\langle a|b \rangle$ is the noise weighted inner product, defined in

the frequency domain as

$$\langle a|b \rangle = 4 \operatorname{Re} \int_0^{f_{\max}} \frac{\tilde{a}(f)\tilde{b}(f)^*}{S(f)} df, \quad (7)$$

and $S(f)$ is the power spectrum of the detector noise (Cutler & Flanagan 1994).

We must also take into account the fact that there is a separate threshold on the search. This arises in the normalization of the probability density for \mathbf{d} above. When there is no threshold, the probability distribution in Equation (5) is correctly normalized. However, when we impose a threshold, we must take into account that not all sets of parameters θ are equally likely to lead to the identification of a gravitational-wave signal. Thus, to normalize the probability, we must integrate the probability over all realizations of the data, \mathbf{d} , which produces an event above the threshold ρ_* :

$$\begin{aligned} p_{\text{det}}(\lambda) &:= \int_{\rho(\mathbf{d}) > \rho_*} d\mathbf{d} \int d\theta p(\mathbf{d}|\theta) p(\theta|\lambda) \\ &= \int d\theta \left[\int_{\rho(\mathbf{d}) > \rho_*} d\mathbf{d} p(\mathbf{d}|\theta) \right] p(\theta|\lambda) \\ &= \int d\theta p_{\text{det}}(\theta) p(\theta|\lambda), \end{aligned} \quad (8)$$

where we have introduced the quantity $p_{\text{det}}(\theta)$ that encodes the probability of an event with parameters θ being observed above threshold. Thus, the normalization factor is simply the probability of an event drawn from the population described by λ leading to a signal in the data that is above the detection threshold.

In principle, θ includes all parameters required to fully describe the system: sky location, distance, binary orientation, masses, and spins. However, for the discussion here, it is more useful to restrict to $\theta = (m_1, m_2, \chi_{\text{eff}})$. Then, the detection probability $p_{\text{det}}(\theta)$ can also be expressed as the fraction of sources with a given set of masses and spins that are observed, where we marginalize over the other parameters. We assess this by distributing sources uniformly in spacetime volume and orientation, and evaluating the fraction that will be observed above the detection threshold, i.e.

$$p_{\text{det}}(\theta) = \frac{1}{V_0} \int_0^{z_{\max}} dz \frac{dV_c}{dz} \frac{1}{1+z} f(z, \theta), \quad (9)$$

where V_c is the comoving volume, $0 \leq f(z, \theta) \leq 1$ is the selection function giving the probability of detecting a source with parameters θ at redshift z and the factor of $1+z$ is due to time dilation caused by the expanding universe. V_0 is the total volume within redshift z_{\max} and serves as a normalization factor. Then $p_{\text{det}}(\lambda)$ is the population-averaged detection probability, which is proportional to the sensitive volume $V_{\text{pop}}(\lambda)$, defined as

$$\frac{V_{\text{pop}}(\lambda)}{V_0} = p_{\text{det}}(\lambda) = \int d\theta p(\theta|\lambda) p_{\text{det}}(\theta), \quad (10)$$

where

$$V_0 = \int_0^{z_{\max}} \frac{dV_c}{dz} \frac{1}{1+z} dz \quad (11)$$

is the total volume. Sensitive volume is a primary ingredient in accounting for the selection effects and can be estimated using

semi-analytical (Abbott et al. 2016c) or numerical methods (Tiwari 2018).

We can now express the distribution for the data \mathbf{d} , corresponding to a single observation, given the model λ , as

$$p(\mathbf{d}|\lambda) = \frac{\int d\theta p(\mathbf{d}|\theta) p(\theta|\lambda)}{p_{\text{det}}(\lambda)}. \quad (12)$$

The expression generalizes in a straightforward manner to a population of observed events, as we assume that the parameters of the signals are independent, such that

$$p(\{\mathbf{d}\}|\lambda) = \prod_{i=1}^N p(\mathbf{d}_i|\lambda). \quad (13)$$

Finally, we can use Equation (3) to obtain the probability of a model λ , given the set of observations $\{\mathbf{d}\}$, as

$$p(\lambda|\{\mathbf{d}\}) = \frac{p(\lambda)}{[p_{\text{det}}(\lambda)]^N p(\{\mathbf{d}\})} \prod_{i=1}^N \int d\theta p(\mathbf{d}_i|\theta) p(\theta|\lambda). \quad (14)$$

This can then be used in a straightforward manner to perform model selection between two models λ_1 and λ_2 as

$$\begin{aligned} \frac{p(\lambda_1|\{\mathbf{d}\})}{p(\lambda_2|\{\mathbf{d}\})} &= \left[\frac{V_{\text{pop}}(\lambda_1)}{V_{\text{pop}}(\lambda_2)} \right]^{-N} \prod_{i=1}^N \left[\frac{\int d\theta p(\mathbf{d}_i|\theta) p_{\text{pop}}(\theta|\lambda_1)}{\int d\theta p(\mathbf{d}_i|\theta) p_{\text{pop}}(\theta|\lambda_2)} \right] \\ &\quad \times \left[\frac{p(\lambda_1)}{p(\lambda_2)} \right]. \end{aligned} \quad (15)$$

The three terms in the odds ratio are easily understood. The final term is simply the ratio of the priors for the two models. In this paper, when comparing models, we take an equal prior probability for the models, so this term is equal to unity. The middle term is the probability of observing the data \mathbf{d}_i given the model λ , and the first term arises as a normalization due to the threshold in the identification of signals in the data. We note that the overall prior on the data $p(\{\mathbf{d}\})$ cancels as it appears, the same way for both the models.

Rather than repeating the parameter estimation with a number of different prior distributions for θ , it is more straightforward to perform it once, with a simple prior, and then to re-weight the samples. In particular, let us assume that we obtain an estimate of the parameters $p_{\text{PE}}(\theta|\mathbf{d})$ given a prior $\pi(\theta)$. We can then use this to obtain an estimate of the conditional probability for \mathbf{d} :

$$p(\mathbf{d}|\theta) = \frac{p_{\text{PE}}(\theta|\mathbf{d}) p(\mathbf{d})}{\pi(\theta)}. \quad (16)$$

When performing parameter estimation, we obtain a set of posterior samples θ^j that describe the posterior distribution for θ . Explicitly, we can approximate an integral over the parameter space as

$$\int d\theta p_{\text{PE}}(\theta|\mathbf{d}) f(\theta) \approx \sum_j f(\theta^j). \quad (17)$$

Thus, the integral in Equation (15) can be well approximated by a sum over the (appropriately weighted) posterior samples

$$\frac{p(\lambda_1|\{\mathbf{d}\})}{p(\lambda_2|\{\mathbf{d}\})} \approx \left[\frac{V_{\text{pop}}(\lambda_1)}{V_{\text{pop}}(\lambda_2)} \right]^{-N} \prod_{i=1}^N \left[\frac{\sum_j p(\theta_i^j|\lambda_1)/\pi(\theta_i^j)}{\sum_j p(\theta_i^j|\lambda_2)/\pi(\theta_i^j)} \right] \times \left[\frac{p(\lambda_1)}{p(\lambda_2)} \right]. \quad (18)$$

The first term gives the ratio of the sensitive volumes for the two models, and favors the model with the lower sensitive volume. The second term sums over the re-weighted posterior samples, where the re-weighting factor is simply the ratio of the desired prior to the one used in the parameter estimation. The final term is the ratio of the priors of the two models. Let us now look in detail at the impact of the three factors when calculating the odds ratios. As discussed previously, we will always assume an equal prior between models, so the final term is unity.

In order to estimate the sensitive volumes of the different population models, defined in Equation (10), we perform Monte Carlo integration. To do so, we sample from the astrophysically expected distributions of parameters and determine the fraction of sources that would be observed. The six populations used in the analysis follow the same mass distribution but different spin distributions. Random samples of the masses and spins are drawn from the population and are assigned randomly chosen orientations and sky locations. Samples are distributed in redshift as determined by standard cosmology. The expected signal-to-noise ratio (S/N) of the signals produced by these binaries at the detectors are calculated and signals that cross a certain S/N threshold are labeled as recovered. Since all of the events apart from GW170814 were observed by only the LIGO detectors, for simplicity we estimate the sensitive volume for the LIGO Hanford—LIGO Livingston detector network. We choose a fixed power spectrum for the detector noise operating close to the sensitivity of the LIGO detectors during the first observing run. While the sensitivity of the detectors varies over the runs, and between the first and second observing runs, the ratio of sensitive volumes for the different population models is relatively insensitive to changes in the detector sensitivity. To optimize the calculation, we estimate the S/N for face-on signals on a fiducial grid of binary masses and spins. The expected S/N of a binary with arbitrary masses, spins, location, and orientation is calculated by linear interpolation in the mass and spin space and incorporating the loss in S/N due to the arbitrary orientation Schutz (2011). To test the efficacy of the procedure, we compare our result with the results reported in reference (Tiwari 2018). Our results are within 10% of the reported values.

The Monte Carlo equivalent of Equation (10) is given by

$$V_{\text{pop}}(\lambda) = V_0 \frac{N_{\text{rec}}}{N}, \quad (19)$$

where N is the number of samples used in simulating the population and N_{rec} is the number of recovered samples.

Next, let us consider the re-weighting of the posterior samples. In particular, the prior, $\pi(\theta)$ is usually taken to be flat in m_1 and m_2 , subject to the condition that $m_1 > m_2$ and flat in the z -components of the spin (Abbott et al. 2016b, 2016a, 2017a). In

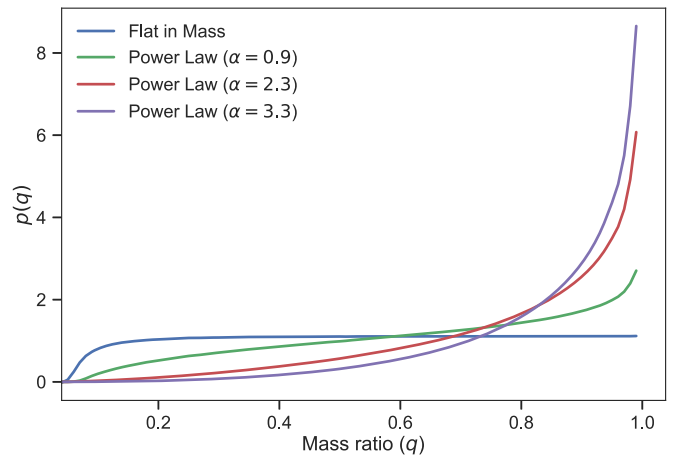


Figure 1. Distribution of mass-ratio for various mass distributions. We select $p(m_1) \propto m_1^{-\alpha}$ and $p(m_2)$ uniform in m_2 between $5 M_{\odot}$ and m_1 . The mass-ratio distribution for different values of the power-law index $\alpha = 0.9, 2.3,$ and 3.3 is shown. At larger values of α , binaries of comparable mass are favored. We also show the mass-ratio distribution from flat m_1 and m_2 distributions.

computing the probabilities for the various models under consideration, we must vary the spin prior to match one of the six distributions under consideration. In addition, we would like to use a different mass prior, which is astrophysically motivated (Fishbach & Holz 2017; Talbot & Thrane 2018). In particular, we select $p(m_1) \propto m_1^{-\alpha}$ and $p(m_2)$ uniform in m_2 between $5 M_{\odot}$ and m_1 (Abbott et al. 2016a, 2016c). In Figure 1 we show the prior distribution for the mass-ratio, given the flat prior $\pi(\theta)$. In addition, we show the distribution that is obtained with the power-law prior, with $\alpha = 2.3$, the value used to obtain the results, as well as $\alpha = 0.9$ and $\alpha = 3.3$. These values, with a mean at $\alpha = 2.3$, cover the 1σ confidence interval of the possible values of α that are consistent with the observations (Abbott et al. 2017d).

The mass-ratio distribution based on an astrophysical model is significantly different from one obtained with a flat prior on the component masses. In particular, close to equal mass-ratio systems are significantly more likely, while binaries with mass ratios greater than 5:1 are down-weighted by a factor between 1.5 and 30, depending upon the value of α . The higher the value of the power-law slope α , the more the distribution is skewed toward equal masses. Since there is a degeneracy between mass-ratio and aligned-spin, a preference for close to equal mass binaries will provide tighter posteriors on the spin. So, applying an astrophysically motivated mass prior leads to a preference for lower spin values. This has a significant impact on down-weighting the high-spin distributions when summing over the re-weighted posterior distributions in Equation (18)).

Finally, we discuss the generation of the posterior samples. The full LIGO-Virgo analysis produces thousands of posterior samples, which are used to obtain the parameter distributions presented in the results papers. In principle, those could be used directly in the calculation of the odds ratio in Equation (18). However, at present only samples corresponding to LIGO's first scientific run are publicly available (Abbott et al. 2018), so we instead generate distributions that mimic the full parameter estimation results, based upon the parameter values and uncertainties presented in the results papers. As we are interested in only the masses and spins, we do not perform a full parameter estimation analysis, but rather use the fact that

Table 1

Comparison between Reported Credible Intervals of Gravitational-wave Observations and Those Obtained from the Approximate Posterior Samples Generated Using the Methods Discussed in the Paper

	Primary Mass $m_1^{\text{source}}/M_\odot$	Secondary Mass $m_2^{\text{source}}/M_\odot$	Effective Spin χ_{eff}
GW150914	$36.0^{+6.0}_{-5.0}/36.2^{+5.2}_{-3.8}$	$27.0^{+4.0}_{-4.0}/29.1^{+3.7}_{-4.4}$	$-0.05^{+0.11}_{-0.12}/-0.06^{+0.14}_{-0.14}$
LVT151012	$24^{+11}_{-6}/23^{+18}_{-6}$	$13^{+4}_{-4}/13^{+4}_{-5}$	$0.1^{+0.2}_{-0.2}/0.0^{+0.3}_{-0.2}$
GW151226	$15.0^{+12.0}_{-4.0}/14.2^{+8.3}_{-3.7}$	$7.0^{+3.0}_{-2.0}/7.5^{+2.3}_{-2.3}$	$0.22^{+0.29}_{-0.13}/0.2^{+0.21}_{-0.1}$
GW170104	$29.0^{+9.0}_{-4.0}/31.2^{+8.4}_{-6}$	$20.0^{+4.0}_{-4.0}/19.4^{+5.3}_{-5.9}$	$-0.09^{+0.13}_{-0.18}/-0.12^{+0.21}_{-0.2}$
GW170608	$13^{+12}_{-3}/12^{+7}_{-2}$	$7^{+2}_{-3}/7^{+2}_{-2}$	$0.13^{+0.34}_{-0.12}/0.07^{+0.23}_{-0.09}$
GW170814	$30.0^{+7.0}_{-4.0}/30.5^{+5.7}_{-3.0}$	$23.0^{+3.0}_{-4.0}/25.3^{+2.8}_{-4.2}$	$0.07^{+0.13}_{-0.13}/0.06^{+0.12}_{-0.12}$

Note. In each entry in the table, the first number and uncertainty give the median and 90% range obtained from the posterior samples generated for our analysis. The second number and interval give the same values as derived directly from the gravitational-wave data (Abbott et al. 2016a, 2017a, 2017b, 2017c). The median values are all in good agreement (less than 1σ deviation in all cases), and the intervals are also comparable.

the measurement of masses and spins is largely independent of the sky location and binary orientation (Singer & Price 2016).

We approximate the detector data \mathbf{d} using a gravitational waveform (Husa et al. 2016; Khan et al. 2016) corresponding to the median value of the reported values of the masses: the aligned-spin magnitude for both black holes set equal to the reported median value of χ_{eff} , as reported in the observations (Abbott et al. 2016a, 2017a, 2017b, 2017c); and the zero-precessing component of the spins. In addition, we generate the waveform associated with a face-on binary (which is consistent with all of the observed signals), and at a distance D that produces an S/N equal to the reported value. To obtain parameter estimates, we then draw a large number of samples in mass, spin and distance, with the distributions following the prior distribution $\pi(\theta)$. Specifically, we use flat distributions of component masses (with limits between $5M_\odot$ and $75M_\odot$), dimensionless aligned-spin magnitudes (between 0 and 0.95), and distances distributed uniformly in volume (between $D/2$ and $2D$). For each of these points, we generate a face-on signal with the appropriate masses, spins, and distance and calculate the Gaussian likelihood defined in Equation (6). We are only interested in mapping the degeneracy between mass-ratio and spin and as we assign the same noise to the detectors in the network, the properties of the detector network are not important and we perform the calculation using a single detector, with a sensitivity matching that obtained by the LIGO detectors during O1. Finally, we estimate the posteriors on the parameters by performing rejection sampling based on the value of the calculated likelihood. By assuming face-on signals, we will obtain a more tightly constrained distance distribution than is allowed with arbitrary orientations. While our calculations do not directly use the distance, the distribution will impact the inferred mass distributions, as the signal observed at the detector depends upon the redshifted masses $m_i(1+z)$. However, in this work, we are primarily concerned with the mass-ratio and spin of the binary, both of which are unaffected by redshift effects. Thus, this simplified method of obtaining posterior samples will provide accurate distributions for the most important parameters in this analysis.

Table 1 compares the credible intervals for the announced GW observations (Abbott et al. 2016a, 2017a, 2017b, 2017c) with the credible intervals of the samples obtained using the method described above. The intervals obtained for the masses

and effective spins are comparable. We have made several approximations, with the expectation that they will generate differences at the observed levels. In addition, the reported gravitational-wave results make use of the average results from both spin-aligned and full-spin precessing models, while in the posteriors that we generate we consider only aligned spins. The mass ratios of the approximated posteriors are nearly equal to the reported values. We have verified that our results are insensitive to shifts in the posterior distributions at the levels reported in Table 1.

3. Results and Discussion

We now investigate the impact of both mass priors and selection effects on the inferred spin distribution for black holes in binaries. Some of the previous works have used χ_{eff} in inferring spin distributions (Farr et al. 2017, 2018); however, the inclusion of these two effects allows tighter constraints on the fraction of binaries with large values of χ_{eff} . The effect of using an astrophysically motivated mass prior can be seen in Figure 2, which shows the mass-ratio and effective spin contours for GW150914 and the much lighter GW151226. As expected, the low-mass systems display significant correlations that are less evident at higher masses. The figure shows the inferred mass-ratio and effective spin distributions for two choices of prior: a flat-in-mass prior and the astrophysically motivated power law. The flat-in-mass distribution allows for more extreme mass ratios, and consequently larger values of the effective spin. For example, the flat prior on masses, which was used in producing parameter estimates in the LIGO-Virgo results papers, follows spins up to 0.6, while imposing the astrophysically motivated mass distribution removes support for $\chi_{\text{eff}} > 0.4$ for GW151226. However, for high-mass systems, such as GW150914, where there is less degeneracy between masses and spins, the effect is not significant.² Figure 2 also plots contours produced from publicly available parameter samples. Also, considering Table 1, mock samples show good agreement with the true samples.

The impact of selection effects is shown in Figure 3 which plots the ratio of the sensitive volume of the spin models to the sensitive volume of the low isotropic model. Binaries with a

² A similar analysis of the effect of mass priors has been performed in Vitale et al. (2017b).




Figure 2. Contour plots for the posterior samples on the $q-\chi_{\text{eff}}$ plane. The contours show the 90% credible regions of the posterior probability. The solid line corresponds to the posterior probability estimated using a publicly available posterior (Abbott et al. 2018). Samples have been produced using a flat prior distribution for the component masses. The dashed line corresponds to the posterior probability obtained by re-weighting the samples with a power-law distribution for m_1 : $p(m_1) \propto m_1^{-2.3}$, and m_2 uniformly distributed between $5 M_{\odot}$ and m_1 . The posterior probability on the effective spin is affected by the choice of prior on the mass-ratio. The power-law prior disfavors a small mass-ratio, which restricts the inferred spin distribution. For comparison, the plot also includes dotted contours prepared using mock samples that are used in this analysis. The true and mock samples are in good agreement.




Figure 3. Sensitive volumes of the spin models compared to the low isotropic model. The figure plots the ratio of sensitive volumes of various models with respect to the low isotropic spin model. This ratio raised to the power by the number of observation accounts for the selection effects.




Figure 4. Odds ratios for different spin magnitude and orientation distributions in reference to the low isotropic model. All of the aligned-spin distributions are excluded at $>4\sigma$; the high isotropic distribution is disfavored by a factor of 15:1, which equates to about 1.8σ .

higher spin magnitude can be observed at a greater distance, so we expect that the low isotropic model, which leads to the population with the smallest spin magnitudes, to have the lowest sensitive volume. Thus, all other things being equal, for each event observed, the model with the lower sensitive volume is preferred. This has a significant impact on the aligned-spin models, but the volume ratio between models with isotropic spin distributions is close to unity. Thus, this does not have a significant impact for a small number of events, but does give a factor of 5 contribution to the odds ratio with 50 events.

The results of the analysis are shown in Figure 4. This shows the odds ratio between each of the six models discussed above and the low isotropic model. The low isotropic distribution is preferred. All models with aligned spins are disfavored at greater than 22000:1, or, equivalently, 4σ . (The flat- and high-aligned distributions are disfavored at $>5\sigma$.) This provides

strong evidence against aligned spins. The improvement arise from three different factors. The inclusion of two additional events (GW170608 and GW170814) increases the odds ratio by a factor of six, while an accurate treatment of the selection effects and the use of an approximately good mass-ratio distribution to handle the mass-ratio-spin degeneracy increase the odds by a factor of sixty. Incorporation of selection effects increases the odds ratio by a factor of four while accounting for correlations between mass-ratio and spin increases the odds ratio by a factor of sixteen. Thus, these three effects combined explain the factor of around four hundred improvement in our ability to exclude aligned-spin models in favor of an isotropic distribution of spins.

Although the LA model does not represent a viable astrophysical model for the spin distribution, the systematic and selection effects are most visible for this model. Standard

Table 2
Estimated Odds Ratio for Different Values of α and Maximum Mass

Spin Model	$\alpha = 0.9$	$\alpha = 2.3$	$\alpha = 3.3$
High Isotropic (Maximum Primary Mass = $75 M_{\odot}$)	14.3:1	15:1	15.5:1
Low Aligned (Maximum Primary Mass = $75 M_{\odot}$)	14000:1	22000:1	20000:1
High Isotropic (Maximum Primary Mass = $95 M_{\odot}$)	14.8:1	16.1:1	16.2:1
Low Aligned (Maximum Primary Mass = $95 M_{\odot}$)	15900:1	25000:1	30000:1

Note. Only the odds ratios between LI versus HI, and LI versus LA models are included.

posterior samples of both GW151226 and GW170608 show evidence of alignment and have an odds ratio of greater than one for the LA spin model, compared to the LI model. Observations are best described by a mixture of LA and LI spin models and a spin analysis aimed at producing mixture models is the best approach. Another important factor to consider is misalignment between the orbital angular momentum vector and the total angular momentum vector. We note that allowing spin misalignments at mergers of up to 30° has only a minimal effect on the probability distribution of χ_{eff} . The overall magnitude of χ_{eff} will reduce by a few percent due to this misalignment. Rigorous analyses that deal with misalignment and mixture models can be found in Stevenson et al. (2017) and Talbot & Thrane (2017).

However, as noted in this paper, systematic and selection effects will affect the inference made by a mixture model, as a mixture model will likely introduce a bias toward LA models, e.g., the odds ratio for GW151226 reduces for the LA model upon using re-weighted posterior samples instead of standard posterior samples, and the odds ratio for GW170608 reduces to close to unity.

This result is only mildly dependent upon the value of the slope, α . As we increase the value of α , we favor lower-mass black hole binaries in the population while lower values of α have a larger fraction of high-mass binaries. The impact of spin on visible volume is more significant for higher masses. Consequently, small values of α lead to larger difference in sensitive volumes between different spin distributions. However, for larger α , the mass-ratio and aligned-spin are more tightly constrained. This increases the importance of re-weighting the posterior samples. Thus, overall, changing the value of α has a limited impact on the results. We also note that choice of the minimum and the maximum masses in the power-law model has only a mild effect on the results. Table 2 lists the odds ratios of the HI and LA spin models in reference to the LI model for different values of α and the maximum mass of the primary component of the binary.

More significantly, we are able to show that low spins are preferred to high spins, even when isotropically distributed. Specifically, the low isotropic distribution is preferred over the high isotropic distribution at 15:1 (around 1.8σ). While highly spinning components can produce a binary with minimal χ_{eff} , this requires a detailed balance between the two spins of the black holes—either the positive spin from one black hole must cancel with the negative spin of the second or the majority of the spin must be in the plane of the orbit. Thus, the fact that χ_{eff} was small for all of the first six black hole binary mergers provides evidence that highly spinning stellar-mass black holes are rare. As discussed in detail in the methods section, this result *does not* depend strongly on the choice of astrophysical mass distribution.

We have shown that the first six observations of black hole binary mergers can be used to place a limit on the magnitude of black hole spins based on gravitational-wave observations. The data show strong evidence for isotropic, rather than aligned, spins. Furthermore, there is emerging evidence that small low-spin magnitudes are preferred to high-spin magnitudes (see Figure 8 of Farr et al. 2018 as well as of Wysocki et al. 2018). This contrasts the nominal spin magnitudes inferred from X-ray binary observations, which are more consistent with high spins.

We emphasize that these conclusions depend on our choices of possible spin distributions. If, for example, our low-spin distribution was restricted to much lower values of χ_{eff} , then aligned-spin configurations would not be so strongly disfavored. However, this would only *strengthen* our main conclusion, which is the preference for low-spin magnitudes.

Distributions of black hole spins will be further refined through future gravitational-wave observations. Tens of black hole binary mergers are expected to be observed in the third advanced LIGO-Virgo observing run. To get a sense of what we might expect, we simulated 30 observations from the low isotropic spin distribution and combined them with the six observations discussed above. With this set of observations, we would be able to exclude a population of black hole binaries where 20% have aligned spins and 80% have isotropic spins with a confidence of at least 4σ . Furthermore, we would obtain an odds ratio in favor of low isotropic spins over high isotropic spins of around 130,000:1 (4.5σ) and flat isotropic spins of around 300:1 (3.0σ). Finally, we note that we have not made use of the precessing component of the spin. A clear observation of precession will give irrefutable evidence of spin misalignment, while observations of χ_p consistent with zero will provide further evidence against high-spin magnitudes.

We thank the following people for interesting and useful discussions: Will Farr, Frank Ohme, Richard O’Shaughnessy, Simon Stevenson, Eric Thrane, and Vivien Raymond. The authors were funded by the Science and Technology Facilities Council (STFC) grants ST/L000962/1 and ST/N005430/1. M.D.H. and S.F. were supported by the European Research Council Consolidator Grant 647839.

ORCID iDs

Vaibhav Tiwari  <https://orcid.org/0000-0002-1602-4176>
 Stephen Fairhurst  <https://orcid.org/0000-0001-8480-1961>
 Mark Hannam  <https://orcid.org/0000-0001-5571-325X>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *PhRvX*, **6**, 041015
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *PhRvL*, **116**, 241102
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016c, *ApJL*, **833**, L1

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, *PhRvL*, **118**, 129901
- Abbott, B. P., Abbott, R., Abbott, T. D., et al., et al. 2017b, *ApJL*, **851**, L35
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017c, *PhRvL*, **119**, 141101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017d, *PhRvD*, **118**, 121101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, LIGO Document T1800235-v14, <https://dcc.ligo.org/LIGO-T1800235/public>
- Ajith, P., Hannam, M., Husa, S., et al. 2011, *PhRvL*, **106**, 241101
- Apostolatos, T., Cutler, C., Sussman, G. J., & Thorne, K. S. 1994, *PhRvD*, **49**, 6274
- Baird, E., Fairhurst, S., & Hannam, M. 2013, *PhRvD*, **87**, 024035
- Campanelli, M., Lousto, C. O., & Zlochower, Y. 2006, *PhRvD*, **74**, 041501
- Cutler, C., & Flanagan, E. E. 1994, *PhRvD*, **49**, 2658
- Damour, T. 2001, *PhRvD*, **64**, 124013
- Dominik, M., Belczynski, K., Fryer, C., et al. 2012, *ApJ*, **759**, 52
- Farr, B., Holz, D. E., & Farr, W. M. 2018, *ApJL*, **854**, L9
- Farr, W. M., Stevenson, S., Miller, M. C., et al. 2017, *Natur*, **548**, 426
- Fishbach, M., & Holz, D. 2017, *ApJL*, **851**, L25
- Haster, C.-J., Wang, Z., Berry, C. P. L., et al. 2016, *MNRAS*, **457**, 4499
- Husa, S., Khan, S., Hannam, M., et al. 2016, *PhRvD*, **93**, 044006
- Khan, S., Khan, S., Hannam, M., et al. 2016, *PhRvD*, **93**, 044007
- Loredo, T. J. 2004, in AIP Conf. Proc. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering (Melville, NY: AIP), 195
- Mandel, I., Farr, W. M., & Gair, J. 2016, LIGO Document P1600187-v1, <https://dcc.ligo.org/LIGO-P1600187/public>
- Poisson, E., & Will, C. M. 1995, *PhRvD*, **52**, 848
- Pürrer, M., Hannam, M., & Ohme, F. 2016, *PhRvD*, **93**, 084042
- Rodriguez, C. L., Chatterjee, S., & Rasio, F. A. 2016, *PhRvD*, **93**, 084029
- Schmidt, P., Ohme, F., & Hannam, M. 2015, *PhRvD*, **91**, 024043
- Schutz, B. 2011, *CQGra*, **28**, 125023
- Singer, L. P., & Price, L. R. 2016, *PhRvD*, **93**, 024013
- Stevenson, S., Berry, C. P. L., & Mandel, I. 2017, *MNRAS*, **471**, 2801
- Talbot, C., & Thrane, E. 2017, *PhRvD*, **96**, 023012
- Talbot, C., & Thrane, E. 2018, *ApJ*, **856**, 173
- Tiwari, V. 2018, *CQGra*, **35**, 145009
- Veitch, J., Raymond, V., Farr, B., et al. 2015, *PhRvD*, **91**, 042003
- Vitale, S., Gerosa, D., Haster, C. J., et al. 2017b, *PhRvL*, **119**, 251103
- Vitale, S., Lynch, R., Sturani, R., & Graff, P. 2017a, *CQGra*, **34**, 03LT01
- Vitale, S., Lynch, R., Veitch, J., et al. 2014, *PhRvD*, **112**, 251101
- Wysocki, D., Lange, J., & O. 'saughnessy, R. 2018, arXiv:1805.06442