Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves

Farr, Will M.; Stevenson, Simon; Mandel, Ilya; Farr, Ben; Vecchio, Alberto; Miller, M Coleman

DOI:
[10.1038/nature23453](https://doi.org/10.1038/nature23453)

License:
None: All rights reserved

Document Version
Peer reviewed version

Citation for published version (Harvard):
Farr, WM, Stevenson, S, Mandel, I, Farr, B, Vecchio, A & Miller, MC 2017, 'Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves', *Nature*, vol. 548, pp. 426-429.
<https://doi.org/10.1038/nature23453>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves

Will M. Farr¹, Simon Stevenson^{1,3}, M. Coleman Miller², Ilya Mandel^{1,3}, Ben Farr⁴ & Alberto Vecchio¹

¹*Birmingham Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Birmingham, B15 2TT, United Kingdom*

²*Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, MD 20742–2421, United States*

³*Kavli Institute for Theoretical Physics, Santa Barbara, CA 93106*

⁴*Enrico Fermi Institute and Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, United States*

The first direct detections of gravitational waves^{1–4} from merging binary black holes open a unique window into the binary black hole formation environment. One promising environmental signature is the angular distribution of the black hole spins; systems formed through dynamical interactions among already-compact objects are expected to have isotropic spin orientations^{5–9} whereas binaries formed from pairs of stars born together are more likely to have spins preferentially aligned with the binary orbital angular momentum^{10–14}. We consider existing gravitational wave measurements of the binary effective spin, the best-measured combination of spin parameters^{3,4}, in the four likely binary black hole detections GW150914, LVT151012, GW151226, and GW170104. If binary black hole spin magnitudes

extend to high values we show that the data exhibit a 2.4σ (0.015 odds ratio¹) preference for an isotropic angular distribution over an aligned one. By considering the effect of 10 additional detections¹⁵, we show that such an augmented data set would enable in most cases a preference stronger than 5σ (2.9×10^{-7} odds ratio). The existing preference for either an isotropic spin distribution or low spin magnitudes for the observed systems will be confirmed (or overturned) confidently in the near future.

Following the detection of a merging binary black hole system, parameter estimation tools compare model gravitational waveforms against the observed data to obtain a posterior distribution on the parameters that describe the compact binary source. The spin parameter with the largest effect on waveforms, and a correspondingly tight constraint from the data³, is a mass-weighted combination of the components of the dimensionless spin vectors of the two black holes that are aligned with the orbital axis, the “effective spin,” $-1 < \chi_{\text{eff}} < 1$ (see Methods Section 2).

Figure 1 shows an approximation to the posterior inferred on χ_{eff} for the four likely gravitational wave (GW) detections GW150914, GW151226, GW170104, and LVT151012 from Advanced LIGO’s first and second observing runs (O1 and O2)^{3,4}. Because samples drawn from the posterior on χ_{eff} are not publicly released at this time, we have approximated the posterior as a Gaussian distribution with the same mean and 90% credible interval, truncated to $-1 < \chi_{\text{eff}} < 1$. None of the χ_{eff} posteriors are consistent with two black holes with large aligned spins, $\chi_{1,2} \gtrsim 0.5$; this contrasts with the large spins inferred for the majority of black holes in X-ray binaries with claimed spin measurements¹⁶ (see below). The analysis here is relatively insensitive to the precise details

¹An odds ratio of r with $r \ll 1$ is equivalent to $x\sigma$ with $x = \Phi^{-1}(1 - r/2)$, where Φ is the unit normal CDF.

of the posterior distributions; other conclusions are more sensitive. In particular, our Gaussian approximation does permit $\chi_{\text{eff}} = 0$ for GW151226 while the true posterior rules this out at high confidence^{2,3}.

Figure 1 Approximate posteriors on χ_{eff} from the Advanced LIGO O1 and GW170104 observations^{3,4}. We approximate the posteriors reported using Gaussians with the same median and 90% credible interval. It is notable that none of the χ_{eff} posteriors support high black hole (BH) spin magnitudes with aligned spins, suggested by observations of stellar-mass black holes in X-ray binaries¹⁶.

Small values of χ_{eff} as exhibited in these systems can result from either intrinsically small

spins or larger spins whose direction is mis-aligned with the orbital angular momentum of the binary (i.e. spin vectors with small z -components). Mis-alignment is capable of producing *negative* values of χ_{eff} , however, whereas aligned spins will always have $\chi_{\text{eff}} \geq 0$. This difference provides strong discriminating power between the two angular distributions, even without good information about the magnitude distribution; to the extent that data favour negative χ_{eff} they weigh heavily against aligned models. To quantify the degree of support for these two alternate explanations of small χ_{eff} values in the merging binary black hole population, we compared the Bayesian evidence for various simple models of the spin population using the GW data set.

Each of our models for the merging binary black hole spin population assumes that the merging black holes are of equal mass (this is marginally consistent with the observations^{3,4}, and the χ_{eff} distribution is not sensitive to the mass ratio—see Methods Section 6). We assume that the population spin distribution factorises into a distribution for the spin magnitude a and a distribution for the spin angles. Finally, we assume that the distribution of spins is common to each component in a merging binary (the distributions of spin for each component in the binary could differ systematically due to different formation histories). Choosing one of three magnitude distributions (see Methods Section 2), “low” (mean $a = 0.33$, standard deviation 0.24), “flat” (mean $a = 0.5$, standard deviation 0.29), “high” (mean $a = 0.67$, standard deviation 0.24) and pairing with an isotropic angular distribution or a distribution that generates perfect alignment yields six different models for the χ_{eff} distribution. These models are shown in Figure 2.

Figure 2 The models for the population distribution of χ_{eff} considered in this paper. In all models we assume that the binary mass ratio $q \equiv m_1/m_2 = 1$ and that the distribution of spin vectors is the same for each component. The “flat” (blue lines), “high,” (green lines), and “low” (red lines) magnitude distributions are defined in Eq. (3). Solid lines give the χ_{eff} distribution under the assumption that the orientations of the spins are isotropic; dashed lines give the distribution under the assumption that both objects’ spins are aligned with the orbital angular momentum. The isotropic distributions are readily distinguished from the aligned distributions by the production of negative χ_{eff} values, while the distinction between the three models for the spin magnitude distribution is less sharp.

These magnitude distributions are not meant to represent any particular physical model, but rather to capture our uncertainty about the spin magnitude distribution; neither observations nor population synthesis codes can at this point authoritatively suggest *any* particular spin distribution¹⁶. Our models, however, allow us to see how sensitive the χ_{eff} distribution is to spin alignment given uncertainties about the spin magnitudes.

We fit hierarchical models of the three existing LIGO O1 and GW170104 observations using these six different, zero-parameter population distributions (see Methods Section 4). We also fit three mixture models for the population, where the angular distribution is a weighted sum of the isotropic and aligned distributions. The evidence, or marginal likelihood, for each of the models is shown in Figure 3. For all three magnitude distributions, the mixture models’ posterior on the mixing fraction peaks at 100% isotropic. Not surprisingly, given the small χ_{eff} values in the three detected systems, the most-favoured model among those with an isotropic angular distribution has the “low” magnitude distribution; the most favoured model among those with an aligned distribution also has the “low” magnitude distribution. The odds ratio between the “low” aligned and “low” isotropic models is 0.015, or 2.4σ ; thus the data favour isotropic spins among our suite of models. While the data favour spin amplitude distributions with small spin magnitudes, note that a model with all binary black hole systems having zero spin is ruled out by the GW151226 measurements, which bound at least one black hole to have spin magnitude ≥ 0.2 at 99% credibility².

Figure 3 Odds ratios among our models using the approximations to the posteriors on χ_{eff} from the O1 and GW170104 observations shown in Figure 1. The flat (“F”), high (“H”), and low (“L”) spin magnitude distributions (see Eq. (3)) are paired with isotropic (“I”) and aligned (“A”) angular distributions, as well as a mixture model of the two (“M”). The most-favoured models have the “low” distribution of spin magnitudes. The odds ratio between the best aligned and best isotropic models is 0.015, or 2.4σ . For all magnitude distributions the pure-isotropic models are preferred over the mixture models; correspondingly, the posterior on the mixture fraction peaks at 100% isotropic.

Estimates of the rate of binary black hole coalescences give a reasonable chance of 10 addi-

tional binary black hole detections in the next three years^{3,15}. Assuming 10 additional detections drawn from each of our six zero-parameter models in addition to the four existing detections from O1 and GW170104, with observational uncertainties drawn randomly from the three Gaussian widths used to approximate the χ_{eff} posteriors in Figure 1², we find the odds ratios shown in Figure 4. We find that most scenarios with an additional 10 detections allow the simulated angular distribution to be inferred with greater than 5σ (2.9×10^{-7} odds) credibility. In the most pessimistic case the distinction is typically 2.4σ (0.016 odds ratio). While such future detections should permit a confident distinction between *angular* distributions, we would remain much less certain about the *magnitude* distribution among the three options considered here until we have a larger number of observations.

²The measurement uncertainty in χ_{eff} depends on the other parameters of the merging binary black hole system, particularly on the mass ratio. Our assumption about future observational uncertainties is appropriate if the parameters of the three detected events are representative of the parameters of future detections. See Methods Section 8 for further discussion.

Figure 4 Distribution of odds ratios predicted with 10 additional observations above the four discussed above. Each panel corresponds to additional observations drawn from one of the χ_{eff} distribution models. The model from which the additional observations are drawn is outlined in red. The height of the blue bar gives the median odds ratio relative to the model from which the additional observations are drawn; the green line gives the 68% (1σ) symmetric interval of odds ratios over 1000 separate draws from the model distribution. The closest median ratio between the most-favoured isotropic model and the most-favoured aligned model is 0.016, corresponding to 2.4σ preference for the correct angular distribution; most models result in more than 5σ preference for the correct angular distribution. Because the four existing observations are included in each data

set the “correct” model is not necessarily preferred over the others, particularly when that model uses the “high” magnitude distribution, which is strongly dis-favoured from the O1 and GW170104 observations alone.

Most of our resolving power for the spin angular distribution is a result of the fact that our “aligned” models cannot produce $\chi_{\text{eff}} < 0$ (see Figure 2). If spins are intrinsically very small, with $a \lesssim 0.2$, then it is no longer possible to resolve the negative effective spin with a small number of observations (see Methods Section 5). As noted below, however, spins observed in X-ray binaries are typically large. Additionally, models which do not permit *some* spins with $\chi_{\text{eff}} \gtrsim 0.1$ are ruled out by the GW151226 observations². An “aligned” model with spin magnitudes from our “flat” distribution but permitting spin vectors oriented anti-parallel to the orbital angular momentum (leading to the possibility of positive *or* negative χ_{eff}) can only be distinguished from an isotropic true population at $\sim 3\sigma$ after 10–20 observations¹⁷; our flat aligned model can be distinguished from such a population at better than 5σ (odds $< 10^{-8}$) after 10 observations, emphasizing the information content of the bound $\chi_{\text{eff}} > 0$ for our aligned models.

Observational data on spin magnitudes in black hole systems is sparse¹⁶. Most of the systems studied are low-mass X-ray binaries rather than the high-mass X-ray binaries that are likely to be the progenitors of double black hole binaries. In addition, there are substantial systematic errors that can complicate these analyses¹⁶ and selection effects could yield a biased distribution. Nonetheless, if we take the reported spin magnitudes as representative then we find that there is a preference for high spins; for example, 14 of the 19 systems with reported spins have dimensionless

spin parameters in excess of 0.5. It is usually argued that the masses and spin parameters of stellar-mass black holes are unlikely to be altered significantly by accretion¹⁸, but this may not be true for all systems¹⁹. Thus the current spin parameters are probably close to their values upon core collapse, at least in high-mass X-ray binaries. However, the specific processes involved in the production of black hole binaries from isolated binaries could alter the spin magnitude distribution of those holes relative to the X-ray binary systems; for example, close tidal interactions could spin up the core, or stripping of the envelope could reduce the available angular momentum^{20–22}.

The spin directions in isolated binary black holes^{10–14} are usually expected to be preferentially aligned. Despite observed spin-orbit misalignments in massive stellar binaries²³, mass transfer and tidal interactions will tend to realign the binary. On the other hand, there is some evidence of spin-orbit misalignment in black hole X-ray binaries^{24,25}. This is consistent with the expectation that a supernova natal kick (if any) can change the orbital plane and misalign the binary²⁶; the supernova can also tilt the spin angle²⁷. Evolutionary processes, such as wind-driven mass loss and post-collapse fallback, can couple the spin magnitude and direction distributions, contrary to our simplified assumptions. A small misalignment at wide separation can also evolve to a more significant misalignment in component spins as the binary spirals in through GW emission²⁸, but χ_{eff} is approximately conserved through this evolution.

The spin directions of binary black holes formed dynamically through interactions in dense stellar environments^{5–8} are expected to be isotropic given the absence of a preferred direction⁹ and the persistence of an isotropic distribution through post-Newtonian evolution^{29,30}.

Acknowledgements We thank Richard O’Shaughnessy, Christopher Berry, Davide Gerosa, and Salvatore Vitale for discussions and comments on this work. WF, SS, IM and AV were supported in part by the STFC. MCM acknowledges support of the University of Birmingham Institute for Advanced Study Distinguished Visiting Fellows program. SS and IM acknowledge support from the National Science Foundation under Grant No. NSF PHY11-25915.

Supplementary Information Supplementary Information is linked to the online version of the paper at <http://www.nature.com/nature>.

Author Contributions All authors contributed to the work presented in this paper.

Author Information Reprints and permissions information is available at <http://www.nature.com/reprints>. The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to w.farr@bham.ac.uk.

Methods

1 Code Availability

This analysis used the Julia language³¹, Python libraries NumPy and SciPy^{32,33}, the plotting library Matplotlib³⁴, and performed computations in IPython notebooks³⁵. A repository containing the code and notebooks used for this analysis, together with the $\text{\LaTeX}$ source for this document, can be found under an open-source “MIT” license at <https://github.com/farr/AlignedVersusIsoSpin>.

2 Effective Spin and Spin Magnitude Distributions

The effective spin is defined by³⁶

$$\chi_{\text{eff}} = \frac{c}{GM} \left(\frac{\vec{S}_1}{m_1} + \frac{\vec{S}_2}{m_2} \right) \cdot \frac{\vec{L}}{|\vec{L}|} \equiv \frac{1}{M} (m_1 \chi_1 + m_2 \chi_2), \quad (1)$$

where $m_{1,2}$ are the gravitational masses of the more-massive (1) and less-massive (2) components, $M = m_1 + m_2$ is the total mass, $\vec{S}_{1,2}$ are the spin angular momentum vectors of the black holes in the binary, $\vec{L}$ is the orbital angular momentum vector, assumed to point in the $\hat{z}$ direction, and $\chi_{1,2}$ are the corresponding dimensionless projections of the individual BH spins. Because the dimensionless spin parameter,

$$a_{1,2} = \frac{c}{Gm_{1,2}^2} |\vec{S}_{1,2}|, \quad (2)$$

of each black hole is bounded by $0 \leq a_{1,2} < 1$, the projections along the orbital axis are bounded by $-1 < \chi_{1,2} < 1$, and $-1 < \chi_{\text{eff}} < 1$.

We form the population distributions of χ_{eff} shown in Figure 2 by assuming that each black hole in a binary has a dimensionless spin magnitude drawn from one of three distributions,

$$p(a) = \begin{cases} 2(1-a) & \text{“low”} \\ 1 & \text{“flat”} \\ 2a & \text{“high”} \end{cases}, \quad (3)$$

referred to as “low,” “flat,” and “high” in the text above. These distributions are shown in Extended Data Figure 1.

Extended Data Figure 1 Distributions of spin magnitudes. See Eq. (3) for the definition of the low (blue line), flat (green line), and high (red line) magnitude distributions used here. The distributions have mean spin 0.33, 0.5, and 0.67 and standard deviations 0.24, 0.29, and 0.24.

3 Mixture model

While we carried out Bayesian comparisons between isotropic and aligned spin distributions under various assumptions, a preference for one of the considered models over the others does not necessarily indicate that it is the correct model. All of the considered models could be inaccurate for the actual distribution, especially since all of the considered models are based on a number of ad-

ditional assumptions, such as decoupled spin magnitude and spin misalignment angle distributions and identical distributions for primary and secondary spins.

We now partly relax the simplified assumptions made earlier by considering the possibility that the true distribution of BBH spin-orbit misalignments observed by LIGO is a mixture of binaries with aligned spins and binaries with isotropic spins.

Extended Data Figure 2 Fraction of the BBH population coming from an isotropic distribution under a mixture model. The dotted line shows the flat prior on the fraction of BBHs coming from an isotropic distribution, f_i , under the mixture model. The 3 red lines show the posterior on f_i after O1 and GW170104 with our various assumptions regarding BH spin magnitudes. The solid line shows the posterior assuming that all BHs

have their spin magnitude drawn from the “flat” distribution. The dashed line assumes the “high” BH spin magnitude distribution $p(a) = 2a$. The dot-dash line assumes the “low” distribution $p(a) = 2(1 - a)$. We see that for a wide range of assumptions regarding BH spin magnitudes, the fraction coming from an isotropic distribution f_i peaks at 1.

We fit a mixture model³⁷ (labelled model ‘M’ in Figure 3) where a fraction f_i of BBHs have spins drawn from an isotropic distribution, whilst a fraction $1 - f_i$ have their spins aligned with the orbital angular momentum. We assume a flat prior on the fraction f_i . To test the robustness of our result, we vary the distribution we assume for BH spin magnitude distributions as with the aligned and isotropic models. We use the “flat”, “high” and “low” distributions (Equation 3), assuming all BHs have their spin magnitude drawn from the same distribution for both the aligned and isotropic populations. We calculate and plot the posterior on f_i given by Equation 6 ($f_i = \lambda$ in the derivation) in Extended Data Figure 2. We find the mean fraction of BBHs coming from an isotropic distribution is 0.70, 0.77 and 0.81 assuming the “low”, “flat” and “high” distributions for spin magnitudes respectively, compared to the prior mean of 0.5. The lower 90% limits are 0.38, 0.51 and 0.60 respectively, compared to the prior of 0.1. In all cases, the posterior peaks at $f_i = 1$. Thus, for these spin magnitude distributions we find that the current O1 and GW170104 LIGO observations constrain the majority of BBHs to have their spins drawn from an isotropic distribution. The evidence ratios of these mixture models to the isotropic distribution with “low” spin magnitudes are 0.43, 0.20 and 0.10 for the “low”, “flat” and “high” spin magnitude models. Thus we cannot rule out a mixture with the current data. If several different components contribute significantly to the true spin distribution it may take tens to hundreds of detections to accurately

determined the mixing fraction, depending on the distribution of spin magnitudes^{17,37}.

4 Hierarchical Modelling

LIGO measures χ_{eff} better than any other spin parameter, but still with significant uncertainty³⁸, so we need to properly incorporate measurement uncertainty in our analysis; thus our analysis must be *hierarchical*^{39,40}. In a hierarchical analysis, we assume that each event has a true, but unknown, value of the effective spin, drawn from the population distribution, which may have some parameters λ ; then the system is observed, represented by the likelihood function, which results in a distribution for the true effective spin (and all other parameters describing the system) consistent with the data. Combining, the joint posterior on each system's χ_{eff}^i parameters and the population parameters λ implied by a set of observations each with data d^i , is

$$p(\{\chi_{\text{eff}}^i\}, \lambda \mid \{d^i\}) \propto \left[\prod_{i=1}^{N_{\text{obs}}} p(d^i \mid \chi_{\text{eff}}^i) p(\chi_{\text{eff}}^i \mid \lambda) \right] p(\lambda). \quad (4)$$

The components of this formula are

- The GW (marginal) likelihood, $p(d \mid \chi_{\text{eff}})$. Here we use “marginal” because we are (implicitly) integrating over all parameters of the signal but χ_{eff} . Note that it is the likelihood rather than the posterior that matters for the hierarchical analysis; if we are given posterior distributions or posterior samples, we need to re-weight to “remove” the prior and obtain the likelihood.
- The population distribution for χ_{eff} , $p(\chi_{\text{eff}} \mid \lambda)$. This function can be parameterised by population-level parameters, λ . (In the cases discussed above, there are no parameters for

the population.)

- The prior on the population-level parameters, $p(\lambda)$.

If we do not care about the individual event χ_{eff} parameters, we can integrate them out, obtaining

$$p(\lambda \mid \{d^i\}) \propto \left[\prod_{i=1}^{N_{\text{obs}}} \int d\chi_{\text{eff}}^i p(d^i \mid \chi_{\text{eff}}^i) p(\chi_{\text{eff}}^i \mid \lambda) \right] p(\lambda). \quad (5)$$

If we are given posterior samples of χ_{eff}^{ij} (i labels the event, j labels the particular posterior sample) drawn from an analysis using a prior $p(\chi_{\text{eff}})$, then we can approximate the integral by a re-weighted average of the population distribution over the samples (here $p(\chi_{\text{eff}}^{ij})$ is the prior used to produce the posterior samples):

$$p(\lambda \mid \{d^i\}) \propto \left[\prod_{i=1}^{N_{\text{obs}}} \frac{1}{N_i} \sum_{j=1}^{N_i} \frac{p(\chi_{\text{eff}}^{ij} \mid \lambda)}{p(\chi_{\text{eff}}^{ij})} \right] p(\lambda). \quad (6)$$

Order of Magnitude Calculation It is possible to estimate at an order-of-magnitude level the rate at which evidence accumulates in favour of or against the isotropic models as more systems are detected. Based on Figure 2, approximate the isotropic population χ_{eff} distribution as uniform on $\chi_{\text{eff}} \in [-0.25, 0.25]$ and the aligned population χ_{eff} distribution as uniform on $\chi_{\text{eff}} \in [0, 0.5]$. Then the odds ratio between the isotropic and aligned models for each event is approximately

$$\frac{p(d \mid I)}{p(d \mid A)} \simeq \frac{P(-0.25 \leq \chi_{\text{eff}} \leq 0.25)}{P(0 \leq \chi_{\text{eff}} \leq 0.5)}, \quad (7)$$

where $P(A \leq \chi_{\text{eff}} \leq B)$ is the posterior probability (here used to approximate the likelihood) that χ_{eff} is between A and B . Using our approximations to the χ_{eff} posteriors described above, this gives an odds ratio of 5 in favour of the isotropic models, which is about a factor of two smaller

than the ratio in the more careful calculation described above. This is a satisfactory answer at an order-of-magnitude level.

If the true distribution is isotropic and follows this simple model, and our measurement uncertainties on χ_{eff} are $\simeq 0.1$, then the geometric mean of each subsequent measurement’s contribution to the overall odds is ~ 3 . After ten additional events, then, the odds ratio becomes $5 \times 3^{10} \simeq 3 \times 10^5$, or 4.6σ , consistent with the results of the more detailed calculation described above. If the true distribution of spins becomes half as wide ($\chi_{\text{eff}} \in [-0.125, 0.125]$ for isotropic and $\chi_{\text{eff}} \in [0, 0.25]$ for aligned spins), with the same uncertainties, then the existing odds ratio becomes 1.08, and each subsequent event drawn from the isotropic distribution contributes on average a factor of 1.6. In this case, after 10 additional events, the odds ratio becomes 150, or 2.7σ . With small spin magnitudes, our angular resolving power vanishes, as discussed in more detail in Methods Section 5.

Accumulation of evidence In Table 1 we show how the evidence for an isotropic distribution increases when including: only the 2 confirmed events—GW150914 and GW151226—from O1; all O1 events (including LVT151012); and all 4 likely binary black hole mergers, including GW170104.

5 Effect of small spin magnitudes

In the main text we considered three models for BH spin magnitudes: “low”, “flat” and “high”. These were intended to capture some of the uncertainty regarding the BH spin magnitude distribution. We may remain observationally uncertain about the spin magnitude distribution until we have $\mathcal{O}(100)$ observations^{41,42}.

Here we extend the “low” model as:

$$p(a) \propto (1 - a)^\alpha \quad (8)$$

When $\alpha = 0$, this recovers the “flat” distribution, whilst $\alpha = 1$ recovers the “low” distribution. For higher values of α , this distribution becomes more peaked towards $a = 0$.

Extended Data Figure 3 Effect of small spins on evidence ratio of isotropic against aligned models. The blue line shows the evidence ratio (plotted as the equivalent sigma) between a model where all systems are from an isotropic distribution, versus one where all systems are aligned, as a function of the power law α corresponding to

Equation 8. The top axis shows the mean spin magnitude $\bar{a}$ which this α corresponds to. We see that for mean spin magnitudes $\lesssim 0.2$ we find no evidence for either distribution over the other.

In Extended Data Figure 3 we plot the evidence ratio of isotropic to aligned distributions (plotted as the equivalent sigma) with spin magnitudes given by this model with α in the range 0–6. The top axis shows the mean spin magnitude that value of α corresponds to (e.g., for the “flat” distribution $\alpha = 0$, the mean spin magnitude is 0.5). We see that if typical BH spins are $\lesssim 0.2$ we have no evidence for one model over the other.

6 Mass Ratio

Extended Data Figure 4 shows the distributions of χ_{eff} that would obtain with a mass ratio $q = m_2/m_1 = 0.5$ compared to the distributions with $q = 1$ used above. The details of the distribution are sensitive to the mass ratio, but in our analysis we are primarily sensitive to the changing *sign* of χ_{eff} under the isotropic models. This latter property is insensitive to mass ratio. As an example, the distinction between the three different spin amplitude distributions after ten additional detections is quite weak compared to the aligned/isotropic distinction in Figure 4. The differences in the χ_{eff} distribution between $q = 1$ and $q = 0.5$ are even smaller than the differences between the different magnitude distributions.

Extended Data Figure 4 Distributions of χ_{eff} assuming all merging black holes have equal masses ($q = 1$) or a 2:1 mass ratio ($q = 0.5$). The details of the distribution are sensitive to the mass ratio, but in our analysis we are primarily sensitive to the changing *sign* of χ_{eff} under the isotropic models. This latter property is unchanged under changing mass ratio.

7 Approximations in the Gravitational Waveform and Selection Effects

While the Advanced LIGO searches use spin-aligned templates they are efficient in detecting misaligned binary black hole systems⁴³; we assume here that the χ_{eff} distribution of observed sources follows the true population.

The model waveforms used to infer the χ_{eff} of the three LIGO events incorporate approximations to the true behaviour of the merging systems that are expected to break down for sufficiently high mis-aligned spins. The effect of these approximations on inference on the parameters describing GW150914 has been investigated in detail⁴⁴. For this source, statistical uncertainties dominate over any waveform systematics. Detailed comparisons with numerical relativity computations using no approximations to the dynamics⁴⁵ also suggest that statistical uncertainties dominate the systematics for this system. Systematics may dominate for signals with this large SNR ($\simeq 23$) when the source is edge-on or has high spins⁴⁴. The other two events discussed in this paper are at much lower SNR, with correspondingly larger statistical uncertainties, and are probably similarly oriented and with similarly small spins, so we do not expect systematic uncertainties to dominate.

We assume here that measurements made in the future are not dominated by systematic errors, but this assumption would need to be revisited for high-SNR, edge-on, or high-spin sources detected in the future.

8 Precision of χ_{eff} measurements

Throughout this work we have made the simplifying assumption that the precision to which χ_{eff} can be constrained for individual binaries is independent of the binary’s properties. In practice, our ability to constrain χ_{eff} is dependent on the system’s properties, in particular its true χ_{eff} and mass ratio, which we illustrate in Extended Data Figure 5.

For this figure a detected population³ of 500 binaries was selected from a population with

³We qualify a system as “detected” if it produces a SNR above 8 in the second-loudest detector to select only

component masses distributed uniformly between 1 and $30 M_{\odot}$ with $m_1 + m_2 < 30 M_{\odot}$, locations distributed uniformly in volume, and orientations distributed isotropically. Data were simulated for each binary, and posteriors were estimated using the LIGO-Virgo parameter estimation library `LALInference`³⁸ using inspiral-only waveform models (merger and ringdown effects can provide additional information for some binaries, but we ignore those effects here). χ_{eff} is better constrained for binaries with high effective spins and high ($\sim$ equal) mass ratios.

We do not expect these effects to qualitatively affect our conclusions, though they could affect predictions for the total number of detections necessary to constrain the population. For example, if the universe preferentially forms asymmetric binaries with low mass ratios, individual χ_{eff} constraints will be systematically worse, requiring more binaries to infer the properties of the population.

coincident events.

Extended Data Figure 5 Widths of the 90% credible intervals for χ_{eff} for 500 binaries in a simulated detected population. χ_{eff} is better constrained for systems with high χ_{eff} and high mass ratio.

1. Abbott, B. P. *et al.* Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters* **116**, 061102 (2016). 1602.03837.
2. Abbott, B. P. *et al.* GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence. *Physical Review Letters* **116**, 241103 (2016). 1606.04855.
3. Abbott, B. P. *et al.* Binary Black Hole Mergers in the First Advanced LIGO Observing Run.

- Physical Review X* **6**, 041015 (2016). 1606.04856.
4. Abbott, B. P. *et al.* Gw170104: Observation of a 50-solar-mass binary black hole coalescence at redshift 0.2. *Phys. Rev. Lett.* **118**, 221101 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevLett.118.221101>.
 5. Sigurdsson, S. & Hernquist, L. Primordial black holes in globular clusters. *Nat.* **364**, 423–425 (1993).
 6. Portegies Zwart, S. F. & McMillan, S. L. W. Black hole mergers in the universe. *Astrophys. J.* **528**, L17 (2000).
 7. Rodriguez, C. L. *et al.* Binary Black Hole Mergers from Globular Clusters: Implications for Advanced LIGO. *Physical Review Letters* **115**, 051101 (2015). 1505.00792.
 8. Stone, N. C., Metzger, B. D. & Haiman, Z. Assisted inspirals of stellar mass black holes embedded in AGN discs: solving the ‘final au problem’. *Mon. Not. R. Astron. Soc.* **464**, 946–954 (2017). 1602.04226.
 9. Rodriguez, C. L., Zevin, M., Pankow, C., Kalogera, V. & Rasio, F. A. Illuminating Black Hole Binary Formation Channels with Spins in Advanced LIGO. *Astrophys. J. Lett.* **832**, L2 (2016). 1609.05916.
 10. Tutukov, A. V. & Yungelson, L. R. The merger rate of neutron star and black hole binaries. *Monthly Notices of the Royal Astronomical Society* **260**, 675–678 (1993).

11. Belczynski, K., Holz, D. E., Bulik, T. & O’Shaughnessy, R. The first gravitational-wave source from the isolated evolution of two stars in the 40-100 solar mass range. *Nat.* **534**, 512–515 (2016). 1602.04531.
12. Stevenson, S. *et al.* Formation of the first three gravitational-wave observations through isolated binary evolution. *Nat. Commun.* **8**, 14906 (2017). 1704.01352.
13. Mandel, I. & de Mink, S. E. Merging binary black holes formed through chemically homogeneous evolution in short-period stellar binaries. *Mon. Not. R. Astron. Soc.* **458**, 2634–2647 (2016). 1601.00007.
14. Marchant, P., Langer, N., Podsiadlowski, P., Tauris, T. M. & Moriya, T. J. A new route towards merging massive black holes. *Astron. and Astrophys.* **588**, A50 (2016). 1601.03718.
15. Abbott, B. P. *et al.* The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914. *Astrophys. J. Lett.* **833**, L1 (2016). 1602.03842.
16. Miller, M. C. & Miller, J. M. The masses and spins of neutron stars and stellar-mass black holes. *Phys. Rep.* **548**, 1–34 (2015). 1408.4145.
17. Vitale, S., Lynch, R., Sturani, R. & Graff, P. Use of gravitational waves to probe the formation channels of compact binaries. *Classical and Quantum Gravity* **34**, 03LT01 (2017). 1503.04307.
18. King, A. R. & Kolb, U. The evolution of black hole mass and angular momentum. *Mon. Not. R. Astron. Soc.* **305**, 654–660 (1999). astro-ph/9901296.

19. Fragos, T. & McClintock, J. E. The Origin of Black Hole Spin in Galactic Low-mass X-Ray Binaries. *Astrophys. J.* **800**, 17 (2015). 1408.2661.
20. Kushnir, D., Zaldarriaga, M., Kollmeier, J. A. & Waldman, R. GW150914: spin-based constraints on the merger time of the progenitor system. *Mon. Not. R. Astron. Soc.* **462**, 844–849 (2016). 1605.03839.
21. Zaldarriaga, M., Kushnir, D. & Kollmeier, J. A. Research Note: The Expected Spins of Gravitational Wave Sources With Isolated Field Binary Progenitors. *ArXiv e-prints* (2017). 1702.00885.
22. Hotokezaka, K. & Piran, T. Implications of the low binary black hole spins observed by LIGO. *ArXiv e-prints* (2017). 1702.03952.
23. Albrecht, S., Reffert, S., Snellen, I. A. G. & Winn, J. N. Misaligned spin and orbital axes cause the anomalous precession of DI Herculis. *Nat.* **461**, 373–376 (2009). 0909.2861.
24. Martin, R. G., Reis, R. C. & Pringle, J. E. Misalignment of the microquasar V4641 Sgr (SAX J1819.3-2525). *Mon. Not. R. Astron. Soc.* **391**, L15–L18 (2008). 0808.2139.
25. Morningstar, W. R. & Miller, J. M. The Spin of the Black Hole 4U 1543-47. *Astrophys. J. Lett.* **793**, L33 (2014). 1408.7028.
26. Kalogera, V. Spin-Orbit Misalignment in Close Binaries with Two Compact Objects. *Astrophys. J.* **541**, 319–328 (2000). astro-ph/9911417.

27. Farr, W. M., Kremer, K., Lyutikov, M. & Kalogera, V. Spin Tilts in the Double Pulsar Reveal Supernova Spin Angular-momentum Production. *Astrophys. J.* **742**, 81 (2011). 1104.5001.
28. Gerosa, D., Kesden, M., Sperhake, U., Berti, E. & O’Shaughnessy, R. Multi-timescale analysis of phase transitions in precessing black-hole binaries. *Phys. Rev. D* **92**, 064016 (2015). 1506.03492.
29. Schnittman, J. D. Spin-orbit resonance and the evolution of compact binary systems. *Phys. Rev. D* **70**, 124020 (2004). astro-ph/0409174.
30. Bogdanović, T., Reynolds, C. S. & Miller, M. C. Alignment of the Spins of Supermassive Black Holes Prior to Coalescence. *Astrophys. J. Lett.* **661**, L147–L150 (2007). astro-ph/0703054.
31. Bezanson, J., Karpinski, S., Shah, V. B. & Edelman, A. Julia: A fast dynamic language for technical computing (2012). 1209.5145.
32. van der Walt, S., Colbert, S. C. & Varoquaux, G. The numpy array: A structure for efficient numerical computation. *Computing in Science & Engineering* **13**, 22–30 (2011). URL <http://aip.scitation.org/doi/abs/10.1109/MCSE.2011.37>. <http://aip.scitation.org/doi/pdf/10.1109/MCSE.2011.37>.
33. Jones, E., Oliphant, T., Peterson, P. *et al.* SciPy: Open source scientific tools for Python (2001–). URL <http://www.scipy.org/>. [Online; accessed 24 May 2017].

34. Hunter, J. D. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering* **9**, 90–95 (2007). URL <http://aip.scitation.org/doi/abs/10.1109/MCSE.2007.55>.
35. Prez, F. & Granger, B. E. Ipython: A system for interactive scientific computing. *Computing in Science & Engineering* **9**, 21–29 (2007). URL <http://aip.scitation.org/doi/abs/10.1109/MCSE.2007.53>.
36. Abbott, B. P. *et al.* Properties of the Binary Black Hole Merger GW150914. *Physical Review Letters* **116**, 241102 (2016). 1602.03840.
37. Stevenson, S., Berry, C. P. L. & Mandel, I. Hierarchical analysis of gravitational-wave measurements of binary black hole spin-orbit misalignments. *ArXiv e-prints* (2017). 1703.06873.
38. Veitch, J. *et al.* Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. *Phys. Rev. D* **91**, 042003 (2015). 1409.7215.
39. Hogg, D. W., Myers, A. D. & Bovy, J. Inferring the Eccentricity Distribution. *Astrophys. J.* **725**, 2166–2175 (2010). 1008.4146.
40. Mandel, I. Parameter estimation on gravitational waves from multiple coalescing binaries. *Phys. Rev. D* **81**, 084029 (2010). 0912.5531.
41. Fishbach, M., Holz, D. & Farr, B. Are LIGO’s Black Holes Made From Smaller Black Holes? *ArXiv e-prints* (2017). 1703.06869.

- 42. Gerosa, D. & Berti, E. First or second generation? Black hole census with gravitational wave observations. *ArXiv e-prints* (2017). 1703.06223.
- 43. Abbott, B. P. *et al.* GW150914: First results from the search for binary black hole coalescence with Advanced LIGO. *Phys. Rev. D* **93**, 122003 (2016). 1602.03839.
- 44. Abbott, B. P. *et al.* Effects of waveform model systematics on the interpretation of GW150914. *ArXiv e-prints* (2016). 1611.07531.
- 45. Abbott, B. P. *et al.* Directly comparing GW150914 with numerical solutions of Einstein's equations for binary black hole coalescence. *Phys. Rev. D* **94**, 064035 (2016). 1606.01262.

Events	$\sigma_{I/A}$	$\sigma_{I/A}$	$\sigma_{I/A}$
	“Low”	“Flat”	“High”
GW150914 and GW151226	1.3	2.2	3.7
All O1 events	1.7	2.7	4.4
All O1 events and GW170104	2.4	3.6	5.4

Table 1: Significance $\sigma_{I/A}$ of the odds ratio between the isotropic and aligned models using just GW150914 and GW151226, all 3 O1 events, and all 4 currently observed events (including GW170104). The number in bold is the result we quote in the main text.