

LETTER TO THE EDITOR

Spin it as you like: The (lack of a) measurement of the spin tilt distribution with LIGO-Virgo-KAGRA binary black holes

Salvatore Vitale^{1,2} , Sylvia Biscoveanu^{1,2} , and Colm Talbot^{1,2} 

¹ LIGO Laboratory, Massachusetts Institute of Technology, 185 Albany St, Cambridge, MA 02139, USA

² Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139, USA
e-mail: salvo@mit.edu

Received 29 September 2022 / Accepted 17 November 2022

ABSTRACT

Context. The growing set of gravitational-wave sources is being used to measure the properties of the underlying astrophysical populations of compact objects, black holes, and neutron stars. Most of the detected systems are black hole binaries. While much has been learned about black holes by analyzing the latest LIGO-Virgo-KAGRA (LVK) catalog, GWTC-3, a measurement of the astrophysical distribution of the black hole spin orientations remains elusive. This is usually probed by measuring the cosine of the tilt angle ($\cos \tau$) between each black hole spin and the orbital angular momentum, with $\cos \tau = +1$ being perfect alignment.

Aims. The LVK Collaboration has modeled the $\cos \tau$ distribution as a mixture of an isotropic component and a Gaussian component with mean fixed at +1 and width measured from the data. We want to verify if the data require the existence of such a peak at $\cos \tau = +1$.

Methods. We used various alternative models for the astrophysical tilt distribution and measured their parameters using the LVK GWTC-3 catalog.

Results. We find that (a) augmenting the LVK model, such that the mean μ of the Gaussian is not fixed at +1, returns results that strongly depend on priors. If we allow $\mu > +1$, then the resulting astrophysical $\cos \tau$ distribution peaks at +1 and looks linear, rather than Gaussian. If we constrain $-1 \leq \mu \leq +1$, the Gaussian component peaks at $\mu = 0.48^{+0.46}_{-0.99}$ (median and 90% symmetric credible interval). Two other two-component mixture models yield $\cos \tau$ distributions that either have a broad peak centered at $0.19^{+0.22}_{-0.18}$ or a plateau that spans the range $[-0.5, +1]$, without a clear peak at +1. (b) All of the models we considered agree as to there being no excess of black hole tilts at around -1. (c) While yielding quite different posteriors, the models considered in this work have Bayesian evidences that are the same within error bars.

Conclusions. We conclude that the current dataset is not sufficiently informative to draw any model-independent conclusions on the astrophysical distribution of spin tilts, except that there is no excess of spins with negatively aligned tilts.

Key words. gravitational waves – black hole physics – methods: data analysis

1. Introduction

More than 90 binary black holes (BBHs) have been detected in the data of the ground-based gravitational-wave detectors LIGO (Aasi et al. 2015) and Virgo (Acernese et al. 2015) by the LIGO-Virgo-Kagra (LVK) collaboration and other groups (Abbott et al. 2021b; Nitz et al. 2021; Olsen et al. 2022). This dataset has been used to infer the properties of the underlying population – or populations – of BBHs. Among the parameters of interest, the masses and spins of the black holes play a prominent role since they can shed light on the binary formation channels¹ (e.g., Vitale et al. 2017; Farr et al. 2017; Zevin et al. 2017, 2021; Farr et al. 2018; Wong et al. 2021; Bouffanais et al. 2021).

There currently exist a few approaches toward measuring the population properties. The first is to use a functional form for a reasonable population distribution, parameterized by some phenomenological parameters. For example, the LVK parameterized the primary mass distribution of the black holes as a mixture of a power law distribution and a Gaussian component (Abbott et al. 2021e; Talbot & Thrane 2018; Fishbach & Holz 2017). This model includes several hyperparameters (since they pertain to the population as a whole, not to the individual events), which are measured from the data: the slope of the power law, the minimum and maximum black hole mass (including a smoothing parameter), the mean and standard deviation of the Gaussian component, and the branching ratio between the power law and Gaussian components. It is worth noting that not all parametric models make equally strong assumptions: an example of a more flexible parametrization is the beta distribution that the LVK has used to describe the population distribution of black hole spin magnitudes (Wysocki et al. 2019). Those more elastic models might be more appropriate if one does not have strong observational or theoretical expectations about what the astrophysical

¹ Eccentricity is also a powerful indicator of a binary formation channel (Peters 1964; Hinder et al. 2008; Morscher et al. 2015; Samsing 2018; Rodriguez et al. 2018b,a; Gondán & Kocsis 2019; Zevin et al. 2017), but it is currently harder to measure due to the limited sensitivity of ground-based detectors at frequencies below 20 Hz.

distribution of a parameter should look like, or simply if they want to be more conservative. Just as Bayesian priors can significantly affect the posterior when the likelihood does not have a strong peak when analyzing individual compact binary coalescences, a hyper-model that is too strong could leave imprints on the inferred hyperparameters.

A second approach is to use nonparametric models, based on Gaussian processes, for example (Tiwari 2021; Edelman et al. 2022; Rinaldi & Del Pozzo 2021; Mandel et al. 2017; Vitale et al. 2019). Those usually have many more free parameters, which allows them to fit features in the data that parameterized models might miss. However, their larger number of parameters implies they might need more sources to reach a level of precision comparable to parametric models. Ideally, when strong parametric models are used, one would like to check that the results are not impacted by the model itself, and instead reveal features that are genuinely present in the data. A possible approach is to run multiple models (parametric and nonparametric) and verify that they agree to within statistical uncertainties.

Finally, recent work has focused on using the predictions of numerical simulations as a model. This typically involves applying machine learning (Wong et al. 2021) or density estimation techniques (Zevin et al. 2017; Bouffanais et al. 2021) to the binary parameter distributions' output by rapid population synthesis codes or N -body simulations. While this approach is more astrophysically motivated, the population synthesis simulations have their own uncertainties and assumptions and often include so many free parameters, such that a complete exploration of the model space is not possible with current computational techniques (e.g., Broekgaarden et al. 2021). The impact of modeling on astrophysical inference of gravitational-wave sources has already become apparent in recent months. Several groups have investigated whether there is evidence for a fraction of black hole spin magnitude to be vanishingly small, finding results that depend on the model to a large extent (Callister et al. 2022; Galaudage et al. 2021; Tong et al. 2022; Roulet et al. 2021; Mould et al. 2022). We note that Callister et al. (2022) also provide a comprehensive summary of the status of that measurement.

In this Letter we focus on the inference of the population distribution of tilt angles for the black hole binaries in the latest LVK catalog. This is the angle that each of the black hole spin vectors forms with the orbital angular momentum at some reference frequency (following the LVK data release, we use 20 Hz for the reference frequency²). In their latest catalog of BBHs, the LVK collaboration characterized this distribution by using a mixture model composed of an isotropic distribution plus a Gaussian distribution that peaks at $\cos \tau = 1$ (i.e., when the spin vector and the angular momentum are aligned) with an unknown width (Talbot & Thrane 2017; Abbott et al. 2021e). This model reflects expectations from astrophysical binary modeling. Indeed, numerical simulations suggest that binaries that formed in galactic fields should have spins preferentially aligned with the angular momentum (e.g., Tutukov & Yungelson 1993; Kalogera 2000; Belczynski et al. 2020; Zaldarriaga et al. 2018; Stevenson et al. 2017; Gerosa et al. 2018), whereas binaries that formed

dynamically (i.e., in globular or star clusters) should have randomly oriented tilts (e.g., Portegies Zwart & McMillan 2002; Rodriguez et al. 2015; Antonini & Rasio 2016; Rodriguez et al. 2019; Gerosa & Fishbach 2021).

While this is a reasonable model, we are interested in verifying whether it is actually supported by the available data, or whether we are instead getting posteriors that are driven by that model. Callister et al. (2022) and Tong et al. (2022) recently considered alternative models for the tilt angles, but they focused on whether there is a cutoff at negative cosine tilts (i.e., for antialigned spins), and if that answer depends on the model for the spin magnitude. On the other hand, we do not limit our investigation to the existence of negative tilts, but instead are interested in what-if anything—can be said about the tilt inference that is not strongly dependent on the model being used. We have considered different alternative models³ and verified that all of them yield Bayesian evidences (and maximum log likelihood values), which are comparable with the default model used by the LVK. We have also included models that allow for a correlation between $\cos \tau$ and other parameters (binary mass, mass ratio, or spin magnitude, in turn) and find that those too yield similar evidences. Critically, all of these alternative – and equally supported by the data – models yield noticeably different posterior distributions for the tilt population relative to the default LVK model. In particular, different models give different support to the existence and position of a feature at positive $\cos \tau$. On the other hand, they all agree as to there being no excess of systems with $\cos \tau \simeq -1$. We conclude that the current constraints on the distribution of tilt angles are significantly affected by the model used, and that more sources (or weaker models) are needed before any conclusions can be drawn about the astrophysical distribution of BBH tilts.

2. Methods

We used hierarchical Bayesian analysis – extensively described in Appendix A – to infer the astrophysical properties of the BBHs reported in GWTC-3. To represent the astrophysical distribution of primary mass, mass ratio, redshift, and spin magnitudes, we used the flagship models used by the LVK in Abbott et al. (2021e). Namely, the primary mass distribution is their “Power law + Peak” (Talbot & Thrane 2018); the mass ratio is a power law (Fishbach & Holz 2020); the two spin magnitudes are independently and identically distributed according to a beta distribution (Wysocki et al. 2019); and the redshift is evolving with a power law (Fishbach et al. 2018). For the (cosine)⁴ tilt distributions, we considered several models of increasing complexity.

3. Results

In Appendix B we report our results when we used, for the tilt distribution, the “DEFAULT” spin model of Appendix B-2-a of Abbott et al. (2021e) (referred to as LVK default in the rest of the paper): this serves as a useful comparison for the more complex models described in the remainder of this work. In this section we reanalyze the GWTC-3 BBH with different parameterized two-component mixture models for $\cos \tau$. We report results for three-component models in Appendix C. It is

² For the BBHs detected in the second half of the third observing run (O3b), the LVK has also released tilt posteriors evaluated at minus infinity, i.e., at very large orbital separations (Mould & Gerosa 2022). We ran the Isotropic + Beta model of Sect. 3.2 on O3b sources only, and found that the analyses with tilts calculated at 20 Hz and minus infinity yield the same astrophysical $\cos \tau$ distribution. We also find that using O3b only sources the $\cos \tau$ distribution moves toward the left, compared to what is shown in Fig. 6, and peaks closer to 0.

³ A public repository with the hyperposteriors used in this work is available at <https://doi.org/10.5281/zenodo.7305735>.

⁴ Even though we only report results for the cosine of the tilt angle – $\cos \tau$ – we might occasionally refer to tilts only, to lighten the text.

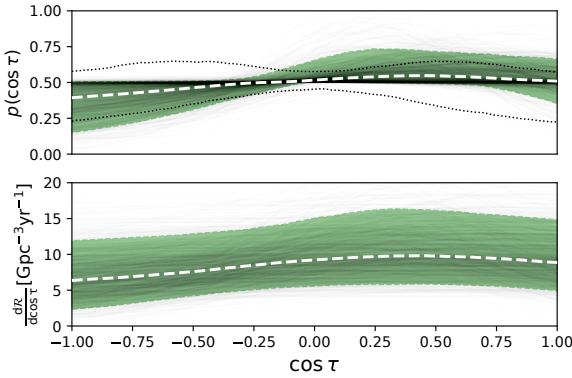


Fig. 1. Posterior for $\cos \tau$ (top panel) and differential merger rate per unit $\cos \tau$ (bottom panel) obtained using the Isotropic + Gaussian model when the mean of the Gaussian component is allowed to vary in the range $\mu \in [-1, 1]$. The two thin black dotted lines in the top panel show the 90% CI obtained by drawing the model's hyperparameters from their priors. In both panels, the thin black lines represent individual posterior draws, whereas the colored band shows the 90% CI. The thick dashed line within the band is the median.

assumed that the $\cos \tau$ distribution is not correlated to any other of the astrophysical parameters. This assumption is revisited in Appendix D.

3.1. Isotropic + Gaussian model

To check if the data require that the Gaussian component of the LVK default mixture model peaks at $\cos \tau = 1$ (Appendix B), we relaxed the assumption that the normal distribution must be centered at +1, that is to say we treated the mean of the Gaussian component μ as another model parameter:

$$p(\cos \tau_1, \cos \tau_2 | \mu, \sigma, g) = \frac{1-g}{4} + g \prod_j^2 \mathcal{N}(\cos \tau_j, \mu, \sigma). \quad (1)$$

The prior for μ is uniform in the range $[-1, 1]$ (Table G.2 reports the priors for the hyperparameters of all models used in the paper).

The top panel of Fig. 1 shows the resulting posterior for the tilt angle. For the mean of the Gaussian component, we measured $\mu = 0.48^{+0.46}_{-0.99}$ (unless otherwise stated, we quote the median and 90% symmetric credible interval). Some of the uncertainty in this measurement is due to our choice to allow for large σ , since Gaussians with large σ are rather flat and can be centered anywhere without significantly affecting the likelihood. If we restrict the prior space to only allow for narrower Gaussians, then μ is much more constrained. For example, if we only keep samples with $\sigma < 0.5$ ($\sigma < 1$), then $\mu = 0.28^{+0.35}_{-0.31}$ ($\mu = 0.41^{+0.47}_{-0.36}$). This can also be seen in a corner plot of the mean and standard deviation of the Gaussian component (Fig. 2, dark green). The data prefer positive means with standard deviations in the approximate range $\sigma \in [0.25, 1.5]$. Smaller values of σ are excluded, as are Gaussians centered at negative values of μ , in other words such that χ_{eff} – the mass-weighted projection of the total spin along the angular momentum (Damour 2001) – would be negative. The bottom panel of Fig. 1 shows the posterior on the differential merger rate per unit $\cos \tau$ for the same model.

It is worth noting that the astrophysical $\cos \tau$ posterior changes entirely if the prior for μ is extended to allow for values outside of the range $[-1, 1]$; whereas, the resulting population model is, of course, still truncated and properly normalized in

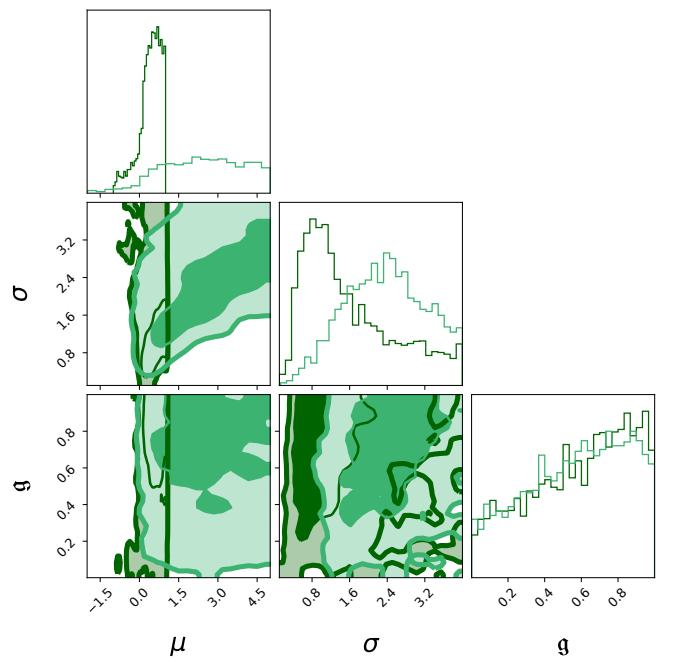


Fig. 2. Joint and marginal posteriors for the mean and standard deviation of the Gaussian component for the Isotropic + Gaussian model, as well as for the branching ratio g , when the mean of the Gaussian component is allowed to vary in the range $\mu \in [-1, 1]$ (dark green) or $\mu \in [-5, 5]$ (light green).

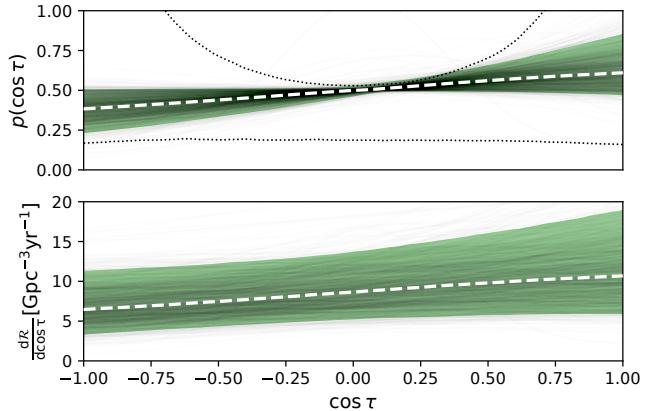


Fig. 3. Same as Fig. 1, but for the Isotropic + Gaussian model, when the mean of the Gaussian component is allowed to vary in the range $\mu \in [-5, 5]$.

that domain. For example, Fig. 3 shows the posterior for the tilt distribution obtained with a wider μ prior, uniform in the range $[-5, 5]$. Here again it is the case that the distribution is consistent with having a peak for aligned spins. A look at the joint distribution of μ and σ , Fig. 2, reveals that the peak at $\mu = 1$ is not obtained because the Gaussian component peaks there, but rather by truncating Gaussians that peak at $\mu > 1$ and have large standard deviations. This results in a much steeper shape for the posterior near the $\cos \tau = +1$ edge than what could possibly be obtained by forcing $-1 \leq \mu \leq 1$. The only feature that seems solid against model variations is the fact that there is no excess of systems at negative $\cos \tau$.

We can better visualize what happens at the edges of the $\cos \tau$ domain – that is for values of tilts close to aligned ($\cos \tau \simeq 1$) or antialigned ($\cos \tau \simeq -1$) – by plotting the ratio of the posterior support for aligned spin versus antialigned spins. This is

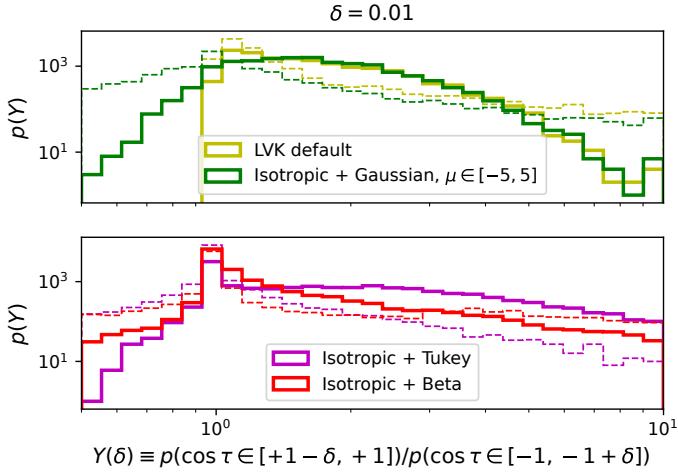


Fig. 4. Unnormalized histogram of $Y(\delta)$, the ratio between the probability of $\cos \tau$ for $\cos \tau \in [1 - \delta, 1]$ and $\cos \tau \in [-1, -1 + \delta]$ for the models of Sects. 3.1–3.3 for $\delta = 0.01$. We also show the result for the LVK default model (Appendix B) for comparison. The solid curves were obtained by sampling the hyperposterior. The dashed line reports the same quantity, but drawing the model’s hyperparameters from their priors. Values of $Y > 1$ imply more support for aligned than antialigned spins.

equivalent to making a histogram of the ratio of the value that the thin black curves in Fig. 3 take on the far right and far left side. Specifically, for each of the posterior draws, we calculated an asymmetry coefficient Y defined as

$$Y(\delta) \equiv \frac{p(\cos \tau \in [1 - \delta, 1])}{p(\cos \tau \in [-1, -1 + \delta])} \quad (2)$$

and plotted the resulting histogram. When $Y = 1$, the $\cos \tau$ distribution takes the same value at both of the edges; when $Y > 1$ ($Y < 1$), the $\cos \tau$ distribution has more support for aligned (antialigned) systems than for antialigned (aligned) ones. We notice that the LVK default model excludes an a priori excess of antialigned tilts, that is $Y < 1$. This is instead not true for the other models we consider in this section. This is shown in the top panel of Fig. 4, for $\delta = 0.01$ (i.e., considering ~ 8 degree widths around $\pm L$, where L is the angular momentum vector). We see that the prior (dashed green curve) of Y for the Isotropic + Gaussian model can extend to values smaller than 1, that is it can produce more antialigned spins than aligned ones (readers can compare this with the blue dashed curve and Appendix B). In fact, we see that the prior for this asymmetry probe is much less strong than in the default LVK default model as it does not exclude $Y < 1$. As for the LVK default, the posterior of Y is not inconsistent with 1. Values of Y smaller than 1 – that is an excess of antialigned spins – are severely suppressed relative to the prior, and so is a large excess of positive tilts. The posterior for Y has a rather broad peak in the range $\sim [1, 2]$, corresponding to distributions that are consistent with being either isotropic or having a mild excess of positive alignment. In Appendix F we show similar plots for different values of δ .

In the top panel of Fig. 5, we show the marginalized posteriors of the branching ratio for the isotropic component, i , of the models described in this section, together with the reference LVK default model. While small variations exist, they all have support across the whole prior range, with a preference for small values of i . By comparing Figs. 1, 3 and 5, one may be surprised that the branching ratio posteriors for the two Isotropic + Gaussian runs are basically the same, and yet Fig. 1 seems to

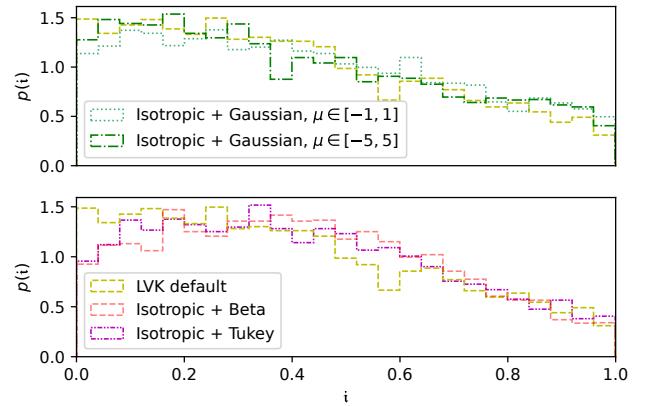


Fig. 5. Marginalized posterior for the branching ratio of the isotropic component – i – for all of the uncorrelated two-component models (Sects. 3.1–3.3). The figure is split into two panels to enhance clarity. In both panels, we also report the posterior obtained with the LVK default model (yellow dashed line) for comparison.

have a higher density of horizontal (i.e., isotropic) curves. This happens because the Gaussian component becomes isotropic for small means μ and large standard deviations σ . Comparing the two distributions for μ in the top-left panel of Fig. 2, we see that indeed the run where μ is constrained to $[-1, 1]$ has more support at small means μ , which, together with large standard deviations σ , yield nearly horizontal posteriors in Fig. 1. To a different extent, the same is true for the models described below: there exist corners of the parameter space where the nonisotropic component can in fact generate a flat distribution, and unless otherwise said we do not a priori exclude that possibility.

3.2. Isotropic + Beta model

To allow for a more elastic model for the nonisotropic component, we now replace the Gaussian component of the previous section with a beta distribution. We have:

$$p(\cos \tau_1, \cos \tau_2 | \alpha, \beta, b) = \frac{1-b}{4} + b \prod_j^2 \mathcal{B}(\cos \tau_j, \alpha, \beta). \quad (3)$$

We offset the input of the beta distribution and scaled its maximum value such that it spans the domain $[-1, 1]$. We stress that we did not limit the range of α and β to nonsingular values, that is we did allow them to be smaller than 1 (Table G.2). In turn, this implies that we can get posteriors for $\cos \tau$ that peak at the edges of the range. The resulting posterior for $\cos \tau$ is shown in Fig. 6, which also shows, for comparison, the 90% credible interval (CI) when drawing hyperparameters from their priors (thin dashed lines). With this model, we recovered a broad peak at small positive values of $\cos \tau$. One can convert the α and β parameters of our rescaled beta distribution to the corresponding mean as

$$\mu_\beta = -1 + 2 \frac{\alpha}{\alpha + \beta}. \quad (4)$$

We find $\mu_\beta = 0.19^{+0.22}_{-0.18}$. While some of the posterior draws do peak at +1, overall the upper edge of the 90% CI band does not show a peak in that region. Compared to what has been seen in the previous section, this model finds more support for antialigned tilts, with a median value that is at the lower edge of the 90% CI for the Isotropic + Gaussian models. In the

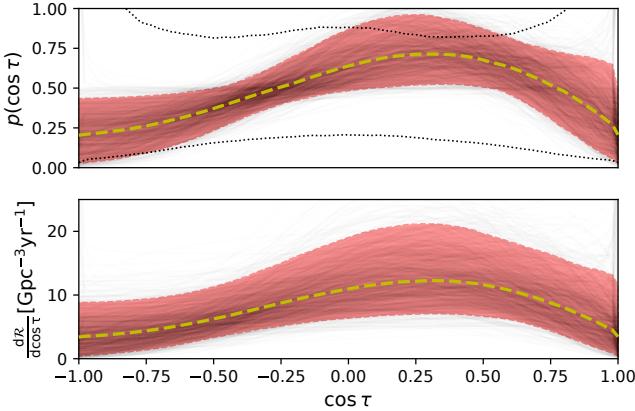


Fig. 6. Same as Fig. 1, but for the Isotropic + Beta model. It is important to note the different scale for the y axis of the bottom panel compared with similar figures for other models.

bottom panel of Fig. 4, we show in red the posterior (solid line) and prior (dashed line) of the asymmetry coefficient Y defined in Eq. (2). Also for this model, we observed that the posterior disfavors configurations with an excess of antialigned black hole tilts, relative to the prior. We noticed that for the Isotropic + Beta model, values of $Y > 1$ do not necessarily represent a peak at positive tilts (see Fig. 6), but only that negative tilts are even more suppressed.

The marginalized posterior for the branching ratio of the isotropic component is shown in the bottom panel of Fig. 5 (histogram with dotted hatches). We find that, unlike the Gaussian-based models of Sect. 3.1 or the LVK default model, it does feature a very broad peak in the middle of the range.

3.3. Isotropic + Tukey model

We concluded our exploration of two-component models with a mixture of an isotropic distribution and a distribution based on the Tukey window function. Mathematically,

$$p(\cos \tau_1, \cos \tau_2 | t, T_{x0}, T_k, T_r) = \frac{1-t}{4} + t \prod_j^2 \mathcal{T}(\cos \tau_j, T_{x0}, T_k, T_r), \quad (5)$$

where the exact expression for the functional form of the window function and a few examples are given in Appendix E. Figure 7 reports the resulting posterior distribution for $\cos \tau$, together with the prior (dotted black lines). The 90% CI band features a plateau that extends from $\cos \tau \simeq -0.5$ to $+1$.

In Fig. 8, we show the posterior distribution for the parameters controlling the Tukey channel, together with the corresponding branching ratio. The Tukey component is centered at T_{x0} , whose fifth and 95th percentile are -0.37 and 0.94 , respectively. The marginal posterior for T_k prefers values close to 1.8 , implying wider Tukey distributions. Smaller values of T_k are possible only for small t , as is expected given that for small t the data cannot constrain the Tukey component, and the posterior must then resemble the uniform prior, which includes small T_k . The posteriors of these two parameters are correlated such that when t is large, T_k is also large, meaning the resulting $\cos \tau$ distribution more closely resembles an isotropic distribution. However, when T_k gets larger than ~ 2 , then it is not constrained at all. This happens because when T_k is that large, the Tukey component is extremely close to an isotropic distribution, at which point the whole model is isotropic, and the branching ratio stops being a meaningful parameter. Finally, the posterior for T_r is wide,

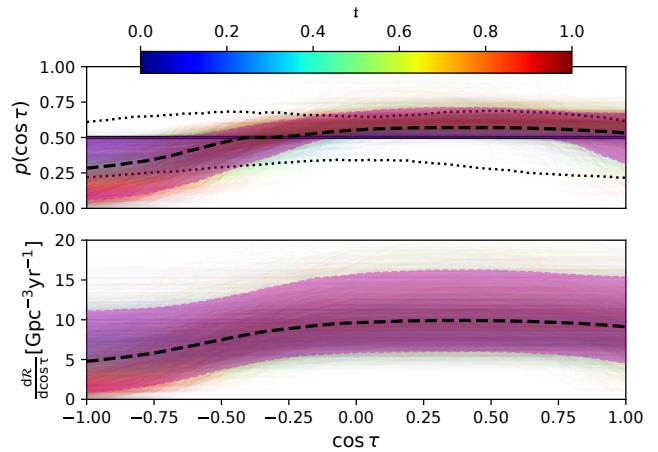


Fig. 7. Same as Fig. 1, but for the Isotropic + Tukey model. Individual posterior draws are colored according to the corresponding value of the branching ratio for the Tukey component, t .

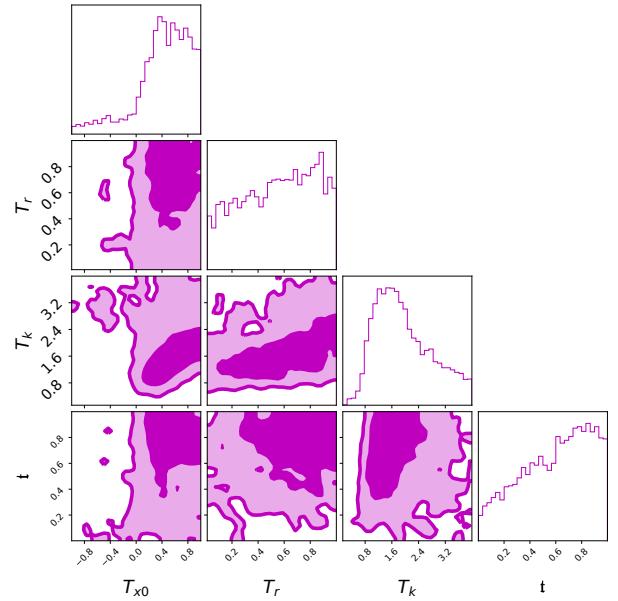


Fig. 8. Joint and marginal posteriors for the hyperparameters and branching ratio of the Tukey component in the Isotropic + Tukey model.

with a preference for larger values, implying a Tukey distribution that ramps up and down smoothly rather than producing sharp features.

Figure 8 also reveals that large values of T_k are responsible for the near entirety of the support at $T_{x0} < 0$, since a Tukey that is basically a uniform distribution can be centered anywhere without affecting the likelihood. If we restrict the analysis to $T_k \leq 2$ ($T_k \leq 1$), we get to see that T_{x0} is much better constrained to $T_{x0} = 0.46^{+0.44} (T_{x0} = 0.30^{+0.41})$, which excludes negative values at $\sim 90\%$ credibility. The fact that our generous hyperparameter priors allow for Tukey distributions that resemble isotropic ones also explain the peak at $Y = 1$ in the bottom panel of Fig. 4 (purple lines, solid for the posterior and dashed for the prior). We find, once again, that the data do not exclude that the $\cos \tau$ distribution is in fact isotropic, and the only solid conclusion one can make seems to be that there is no excess of systems at $\cos \tau \simeq -1$, since $p(Y)$ is heavily suppressed relative to its prior for $Y \lesssim 1$.

4. Conclusions

In this paper we have reanalyzed the LVK’s 69 BBHs of GWTC-3, using different models for the astrophysical distribution of the black hole spin tilt angle, in other words the angle the spin vector forms with the orbital angular momentum at a reference frequency (20 Hz). Black hole spin tilts can yield precious information about their astrophysical formation channels. It is usually expected that dynamical formation of binaries results in an isotropic distribution of the spin vectors (e.g., Portegies Zwart & McMillan 2002; Rodriguez et al. 2015; Antonini & Rasio 2016; Rodriguez et al. 2019; Gerosa & Fishbach 2021). On the other hand, for black hole binaries that formed in the field via isolated binary evolution, it is expected that the spins are nearly aligned with the angular momentum (e.g., Tutukov & Yungelson 1993; Kalogera 2000; Belczynski et al. 2020; Zaldarriaga et al. 2018; Stevenson et al. 2017; Gerosa et al. 2018), that is to say that tilts are small, as the angular momenta of the progenitor stars are aligned by star-star and star-disk interactions (Hut 1981; Packet 1981). Indeed, if the binary forms in the field, the only mechanism that could yield significant black hole tilts are asymmetries in the supernovae explosions that create the black holes. These asymmetries can impart a natal kick large enough to tilt the orbital plane (Katz 1975; Kalogera 2000; Hurley et al. 2002). However, the black hole natal kick distribution is poorly understood both theoretically (e.g., Dominik et al. 2012; Zevin et al. 2017; Mapelli & Giacobbo 2018; Repetto et al. 2012; Giacobbo & Mapelli 2020; Fragos et al. 2010) and observationally (Brandt et al. 1995; Nelemans et al. 1999; Mirabel et al. 2001, 2002; Wong et al. 2014).

These expectations explain why in their most recent catalog, the LVK has modeled the astrophysical tilt distribution as a mixture of two components: an isotropic part and a Gaussian distribution centered at $\cos \tau = 1$ and with a width that is measured from the data. This is a rather strong model, as it forces a Gaussian onto the data that must be centered at +1. This might not be advisable because individual tilt measurements are usually broad, which implies that the functional form of the assumed astrophysical model can leave a discernible imprint on the posterior. Given the rather large uncertainties as to how much spin misalignment can be produced in each channel, and the fact that the BBH population currently available might contain black holes that formed in different channels (Zevin et al. 2021; Wong et al. 2021; Bouffanais et al. 2021; Franciolini & Pani 2022), it is legitimate to question whether the data require that the $\cos \tau$ distribution peaks at +1, as opposed to just being consistent with it. We find that it does not.

We consider three two-component mixture models: **Isotropic + Gaussian**, made of an isotropic component and a Gaussian component whose mean is not fixed at +1, but rather measured from the data; **Isotropic + Beta**, made of an isotropic component and a (singular) beta component; and **Isotropic + Tukey**, made of an isotropic component and a distribution based on the Tukey window. We find that the only model with a posterior that peaks at $\cos \tau = 1$ is the **Isotropic + Gaussian**, and only if we allow the mean of the Gaussian to take values outside of the $\cos \tau$ domain $[-1, 1]$: Fig. 2 shows that in this case the Gaussian component peaks at $\mu > 1$ and is very broad, yielding a peak at +1 (Fig. 3). The other two models yield either a posterior that peaks at $\cos \tau \approx 0.2$ and no peak at +1, or a plateau from $\cos \tau \approx -0.5$ to +1. For all of the above models, the data do not decisively rule out a fully isotropic tilt distribution, but they are inconsistent with an excess of

systems with large and negative tilts. The models we considered found features—in addition to an isotropic distribution—whose exact shape depends on the model’s flexibility to a large extent. Therefore, the fact that the LVK default model finds support for a peak at $\cos \tau = 1$ can be explained because it can only add support at $\cos \tau = 1$ when trying to match any population features in addition to the isotropic distribution. Our results agree with previous results for population models fitting the distribution of χ_{eff} , which find only a small fraction of sources with negative χ_{eff} , implying negative tilts (Roulet & Zaldarriaga 2019; Miller et al. 2020; Roulet et al. 2021; Callister et al. 2022). Indeed, if we recast our inferred distributions for $\cos \tau$ and spin magnitude to the resulting χ_{eff} distribution, we would obtain results consistent with Abbott et al. (2021e). In the appendices below, we report on other models. In Appendix C we consider three three-component mixture models: **Isotropic + 2 Gaussians**, made of an isotropic component and two Gaussian components; **Isotropic + Gaussian + Beta**, made of an isotropic component a (nonsingular) beta component and a Gaussian component; and **Isotropic + Gaussian + Tukey**, made of an isotropic component, a distribution based on the Tukey window and a Gaussian component. These more elastic models are consistent with a broad plateau in the $\cos \tau$ posterior that extends from ~ -0.5 to +1. Whether there are also peaks or features at small positive $\cos \tau$ and/or at +1 depends on the exact model.

In Appendix D we augment some of the two-component mixture models to allow for correlations between the tilt angles and another of the binary parameters: component masses, component spins, mass ratio, and total mass in turn. One might expect some correlation as the mechanisms that can misalign the binary orbital plane or the black hole spins are affected by the binary parameters (e.g., Janka & Mueller 1994; Burrows & Hayes 1996; Fryer & Kusenko 2006; Gerosa & Fishbach 2021). To keep the number of model parameters limited – consistent with the relatively limited number of sources – we only consider linear correlations, and only correlated the tilt distribution with one other parameter at a time. However, even with these limitations, we find that the current dataset cannot significantly constrain eventual correlations. For all of the models considered in this work, we report Bayesian evidences (Table G.1), which might be used to calculate odd ratios. We find that – within sampling and numerical uncertainties – all of the models are equally supported by the data.

We conclude that the current dataset is not yet large and informative enough to prove that the astrophysical tilt distribution has features, nor that if features exist they manifest as an excess of systems with nearly aligned spin vectors. On the contrary, most of the models we considered yield a broad peak in the astrophysical $\cos \tau$ distribution at small and positive values. The only conclusion that is consistently found across all models is that there is no excess of systems with negative tilts, relative to what is expected in an isotropic distribution. Our results agree with the literature (e.g., Callister et al. 2022; Tong et al. 2022; Abbott et al. 2021e; Mould et al. 2022) as to the lack of an excess of $\cos \tau \approx -1$, but disagree as to other details (e.g., whether there is a hard cutoff in the $\cos \tau$ distribution at $\cos \tau < 0$ (cfr Callister et al. 2022)). The point of this work is to show that those disagreements are to be expected, given the information in the current dataset. The next observing run of LIGO, Virgo, and KAGRA is scheduled to start in early 2023 (Abbott et al. 2018) and should yield hundreds of BBH sources. Those may yield a first firm measurement of the astrophysical distribution of the tilt

angle, and possibly allow us to begin probing correlations with other astrophysical parameters.

Acknowledgements. The authors would like to thank C. Adamcewicz, V. Baibhav, T. Dent, S. Galauadage, C. Rodriguez and M. Zevin for useful comments and discussion. We would in particular like to thank T. Callister and D. Gerosa for many insightful comments and suggestions. We would like to thank the anonymous A&A Referee, whose comments helped improve the manuscript. S.V. is supported by NSF through the award PHY-2045740. S.B. is supported by the NSF Graduate Research Fellowship under Grant No. DGE-1122374. CT is supported by the MKI Kavli Fellowship. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation. We used publicly-available programs BILBY (Ashton et al. 2019; Romero-Shaw et al. 2020), DYNESTY (Speagle 2020) and GWPOLULATION (Talbot et al. 2019). This paper carries LIGO document number LIGO-P2200275.

References

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, *Class. Quant. Grav.*, 32, 074001

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, *Liv. Rev. Rel.*, 21, 3

Abbott, R., Abbott, T. D., Acernese, F., et al. 2020a, *GWTC-2 Data Release: Parameter Estimation Samples and Skymaps*, <https://dcc.ligo.org/LIGO-P2000223/public>

Abbott, R., Abbott, T. D., Acernese, F., et al. 2020b, *Parameter estimation sample release for GWTC-1*, <https://dcc.ligo.org/LIGO-P1800370/public>

Abbott, R., Abbott, T. D., Acernese, F., et al. 2021a, *GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run - Parameter Estimation Data Release*

Abbott, R., Abbott, T. D., Acernese, F., et al. 2021b, ArXiv e-prints [arXiv:2111.03606]

Abbott, R., Abbott, T. D., Acernese, F., et al. 2021c, *GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run - O3 search sensitivity estimates*

Abbott, R., Abbott, T. D., Acernese, F., et al. 2021d, *GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run - Parameter Estimation Data Release*

Abbott, R., Abbott, T. D., Acernese, F., et al. 2021e, ArXiv e-prints [arXiv:2111.03634]

Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, *Class. Quant. Grav.*, 32, 024001

Adamcewicz, C., & Thrane, E. 2022, *MNRAS*, 517, 3928

Antonini, F., & Rasio, F. A. 2016, *ApJ*, 831, 187

Ashton, G., Huebner, M., Lasky, P. D., et al. 2019, *ApJS*, 241, 27

Belczynski, K., Klencki, J., Fields, C. E., et al. 2020, *A&A*, 636, A104

Biscoveanu, S., Isi, M., Vitale, S., & Varma, V. 2021, *Phys. Rev. Lett.*, 126, 171103

Biscoveanu, S., Callister, T. A., Haster, C.-J., et al. 2022, *ApJ*, 932, L19

Bouffanais, Y., Mapelli, M., Santoliquido, F., et al. 2021, *MNRAS*, 507, 5224

Brandt, W. N., Podsiadlowski, P., & Sigurdsson, S. 1995, *MNRAS*, 277, L35

Broekgaarden, F. S., Berger, E., Stevenson, S., et al. 2021, *MNRAS*, 516, 5737

Burrows, A., & Hayes, J. 1996, *Phys. Rev. Lett.*, 76, 352

Callister, T. A., Haster, C.-J., Ng, K. K. Y., Vitale, S., & Farr, W. M. 2021, *ApJ*, 922, L5

Callister, T. A., Miller, S. J., Chatzioannou, K., & Farr, W. M. 2022, *ApJ*, 937, L13

Damour, T. 2001, *Phys. Rev. D*, 64, 124013

Dominik, M., Belczynski, K., Fryer, C., et al. 2012, *ApJ*, 759, 52

Edelman, B., Doctor, Z., Godfrey, J., & Farr, B. 2022, *ApJ*, 924, 101

Farr, W. M. 2019, *Res. Notes Am. Astron. Soc.*, 3, 66

Farr, W. M., Stevenson, S., Coleman Miller, M., et al. 2017, *Nature*, 548, 426

Farr, B., Holz, D. E., & Farr, W. M. 2018, *ApJ*, 854, L9

Fishbach, M., & Holz, D. E. 2017, *ApJ*, 851, L25

Fishbach, M., & Holz, D. E. 2020, *ApJ*, 891, L27

Fishbach, M., Holz, D. E., & Farr, W. M. 2018, *ApJ*, 863, L41

Fragos, T., Tremmel, M., Rantsiou, E., & Belczynski, K. 2010, *ApJ*, 719, L79

Franciolini, G., & Pani, P. 2022, *Phys. Rev. D*, 105, 123024

Fryer, C. L., & Kusenko, A. 2006, *ApJ*, 163, 335

Galauadage, S., Talbot, C., Nagar, T., et al. 2021, *ApJ*, 921, L15

Gerosa, D., & Fishbach, M. 2021, *Nat. Astron.*, 5, 8

Gerosa, D., Berti, E., O'Shaughnessy, R., et al. 2018, *Phys. Rev. D*, 98, 084036

Giacobbo, N., & Mapelli, M. 2020, *ApJ*, 891, 141

Golomb, J., & Talbot, C. 2022, ArXiv e-prints [arXiv:2210.12287]

Gondán, L., & Kocsis, B. 2019, *ApJ*, 871, 178

Hinder, I., Vaishnav, B., Herrmann, F., Shoemaker, D. M., & Laguna, P. 2008, *Phys. Rev. D*, 77, 081502

Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, *MNRAS*, 329, 897

Hut, P. 1981, *A&A*, 99, 126

Janka, H. T., & Mueller, E. 1994, *A&A*, 290, 496

Kalogera, V. 2000, *ApJ*, 541, 319

Katz, J. I. 1975, *Nature*, 253, 698

Mandel, I., Farr, W. M., Colonna, A., et al. 2017, *MNRAS*, 465, 3254

Mandel, I., Farr, W. M., & Gair, J. R. 2019, *MNRAS*, 486, 1086

Mapelli, M., & Giacobbo, N. 2018, *MNRAS*, 479, 4391

Miller, S., Callister, T. A., & Farr, W. 2020, *ApJ*, 895, 128

Mirabel, I. F., Dhawan, V., Mignani, R. P., Rodrigues, I., & Guglielmetti, F. 2001, *Nature*, 413, 139

Mirabel, I. F., Mignani, R., Rodrigues, I., et al. 2002, *A&A*, 395, 595

Morscher, M., Patabiraman, B., Rodriguez, C., Rasio, F. A., & Umbreit, S. 2015, *ApJ*, 800, 9

Mould, M., & Gerosa, D. 2022, *Phys. Rev. D*, 105, 024076

Mould, M., Gerosa, D., Broekgaarden, F. S., & Steinle, N. 2022, *MNRAS*, 517, 2738

Nelemans, G., Tauris, T. M., & van den Heuvel, E. P. J. 1999, *A&A*, 352, L87

Nitz, A. H., Kumar, S., Wang, Y.-F., et al. 2021, *ApJ*, submitted [arXiv:2112.06878]

Olsen, S., Venumadhav, T., Mushkin, J., et al. 2022, *Phys. Rev. D*, 106, 043009

Packet, W. 1981, *A&A*, 102, 17

Peters, P. C. 1964, *Phys. Rev.*, 136, B1224

Portegies Zwart, S. F., & McMillan, S. L. W. 2002, *ApJ*, 576, 899

Repetto, S., Davies, M. B., & Sigurdsson, S. 2012, *MNRAS*, 425, 2799

Rinaldi, S., & Del Pozzo, W. 2021, *MNRAS*, 509, 5454

Rodriguez, C. L., Morscher, M., Patabiraman, B., et al. 2015, *Phys. Rev. Lett.*, 115, 051101; Erratum: 2016, 116, 029901

Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., et al. 2018a, *Phys. Rev. D*, 98, 123005

Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., & Rasio, F. A. 2018b, *Phys. Rev. Lett.*, 120, 151101

Rodriguez, C. L., Zevin, M., Amaro-Seoane, P., et al. 2019, *Phys. Rev. D*, 100, 043027

Romero-Shaw, I. M., Talbot, C., Biscoveanu, S., et al. 2020, *MNRAS*, 499, 3295

Roulet, J., & Zaldarriaga, M. 2019, *MNRAS*, 484, 4216

Roulet, J., Chia, H. S., Olsen, S., et al. 2021, *Phys. Rev. D*, 104, 083010

Safarzadeh, M., Farr, W. M., & Ramirez-Ruiz, E. 2020, *ApJ*, 894, 129

Samsing, J. 2018, *Phys. Rev. D*, 97, 103014

Speagle, J. S. 2020, *MNRAS*, 493, 3132

Stevenson, S., Vigna-Gómez, A., Mandel, I., et al. 2017, *Nat. Commun.*, 8, 14906

Talbot, C., & Thrane, E. 2017, *Phys. Rev. D*, 96, 043030

Talbot, C., & Thrane, E. 2018, *ApJ*, 856, 173

Talbot, C., Smith, R., Thrane, E., & Poole, G. B. 2019, *Phys. Rev. D*, 100, 043030

Tiwari, V. 2021, *Class. Quant. Grav.*, 38, 155007

Tong, H., Galauadage, S., & Thrane, E. 2022, Arxiv eprints [arXiv:2209.02206]

Tutukov, A. V., & Yungelson, L. R. 1993, *MNRAS*, 260, 675

Vitale, S., Lynch, R., Sturani, R., & Graff, P. 2017, *Class. Quant. Grav.*, 34, 03LT01

Vitale, S., Farr, W. M., Ng, K., & Rodriguez, C. L. 2019, *ApJ*, 886, L1

Vitale, S., Gerosa, D., Farr, W. M., & Taylor, S. R. 2020, in *Handbook of Gravitational Wave Astronomy*, eds. C. Bambi, S. Katsanevas, & K. D. Kokkotas (Springer), *Liv. Ref. Work.*, 45

Wong, T.-W., Valsecchi, F., Ansari, A., et al. 2014, *ApJ*, 790, 119

Wong, K. W. K., Breivik, K., Kremer, K., & Callister, T. 2021, *Phys. Rev. D*, 103, 083021

Wysocki, D., Lange, J., & O'Shaughnessy, R. 2019, *Phys. Rev. D*, 100, 043012

Zaldarriaga, M., Kushnir, D., & Kollmeier, J. A. 2018, *MNRAS*, 473, 4174

Zevin, M., Pankow, C., Rodriguez, C. L., et al. 2017, *ApJ*, 846, 82

Zevin, M., Baveria, S. S., Berry, C. P. L., et al. 2021, *ApJ*, 910, 152

Appendix A: Hierarchical inference

We aim to measure the hyperparameters λ that control the distribution of single-event parameters θ (the black hole masses, spins, redshifts, etc.) given the dataset D consisting of the 69 GWTC-3 BBHs with a false alarm ratio smaller than 1 per year— $D \equiv \{d_i, i = 1 \dots 69\}$ —reported by the LVK collaboration (Abbott et al. 2021e). The posterior for λ can be written as follows (Mandel et al. 2019; Fishbach et al. 2018; Vitale et al. 2020):

$$p(\lambda|D) \propto \pi(\lambda) \prod_{i=1}^{69} \frac{p(d_i|\lambda)}{\alpha(\lambda)},$$

where we have analytically marginalized over the overall merger rate, which is not relevant for our inference. The function $\alpha(\lambda)$ represents the detection efficiency, that is to say the fraction of BBHs that are detectable, given the population parameters λ ; $\pi(\lambda)$ is the prior for the population hyperparameters, and $p(d_i|\lambda)$ is the likelihood of the stretch of data containing the i -th BBH. This allowed us to account for selection effects and infer the properties of the underlying, rather than the observed, population.

Using Bayes' theorem and marginalizing over the single-event parameters, the single-event likelihood can be written as

$$p(d_i|\lambda) = \int d\theta p(d_i|\theta) \pi(\theta|\lambda) \propto \int d\theta \frac{p(\theta|d_i, \mathcal{H}_{\text{PE}}) \pi(\theta|\lambda)}{\pi(\theta|\mathcal{H}_{\text{PE}})}, \quad (\text{A.1})$$

where $p(\theta|d_i, \mathcal{H}_{\text{PE}})$ is the posterior distribution for the binary parameters θ of the i -th source. The population hyperparameters λ are typically inferred using a hierarchical process that first involves obtaining posteriors for θ for each individual event under a noninformative prior, $\pi(\theta|\mathcal{H}_{\text{PE}})$. The hypothesis \mathcal{H}_{PE} represents the settings that were used during this individual-event parameter estimation step. The last term, $\pi(\theta|\lambda)$, is the population prior, that is to say our model for how the parameters θ are distributed in the population, given the hyperparameters.

The integral in Eq. A.1 can be approximated as a discrete sum

$$\int d\theta p(d_i|\theta) \pi(\theta|\lambda) \simeq N_{\text{samples}}^{-1} \sum_k^{N_{\text{samples}}} \frac{\pi(\theta_i^k|\lambda)}{\pi(\theta_i^k|\mathcal{H}_{\text{PE}})},$$

where the N_{samples} samples are drawn from the posterior distribution of the i -th event. We used the posterior samples of the 69 BBHs reported in GWTC-3, as released in Abbott et al. (2020b,a, 2021a,d). For the sources reported in GWTC-1, we used the samples labeled `IMRPhenomPv2_posterior` in the data release; for GWTC-2 we used `PublicationSamples`; for GWTC-2.1 we used `PrecessingSpinIMRHM`; and for GWTC-3 we use `C01:Mixed`. To sample the hyperposterior, we used the DYNESTY (Speagle 2020) sampler available with the `GWPopulation` package (Talbot et al. 2019).

The detection efficiency $\alpha(\lambda)$ can also be calculated through an approximated sum starting from a large collection of simulated BBHs for which the signal-to-noise ratio (or another detection statistic) is recorded, as described in Farr (2019) and Abbott et al. (2021e). We used the `endo3_bbhpop-LIGO-T2100113-v12-1238166018-15843600.hdf5` sensitivity file released by the LVK (Abbott et al. 2021c) to calculate $\alpha(\lambda)$, using a false alarm threshold of one per year to identify detectable sources, consistently with Abbott et al. (2021e).

Appendix B: Reference tilt model

Here, we compare our results against the LVK's model (LVK default) of Abbott et al. (2021e): a mixture between an isotropic component and Gaussian distribution with $\mu = 1$ and an unknown standard deviation,

$$p(\cos \tau_1, \cos \tau_2 | \sigma, g) = \frac{1-g}{4} + g \prod_j^2 \mathcal{N}(\cos \tau_j, \mu = 1, \sigma). \quad (\text{B.1})$$

The Gaussian component was truncated and normalized in the range $[-1, 1]$. The two hyperparameters of LVK default are thus the branching ratio g of the Gaussian component and its standard deviation σ ; they are the same for both black holes. We notice that in Talbot & Thrane (2017), the two normal distributions can assume different values of σ . However, since in general the spins of the least massive objects are measured with extremely large uncertainty, there is no reason to expect that imposing the same distribution on both tilts would introduce biases.

In Fig. B.1 we show the resulting inference on the $\cos \tau$, which—modulo differences in sampling settings—is directly comparable to what is presented by the LVK in Abbott et al. (2021e) (their Fig. 15). The colored area shows the 90% CI, the thick dashed line is the median, and the thin lines represent individual draws from the posterior. The two dashed lines represent the edges of the 90% CI obtained by sampling the hyperparameters from their priors. It is worth noticing that the LVK default model excludes, a priori, the possibility of an excess of tilts relative to isotropy (i.e., a posterior larger than 0.5) at negative values, as well as a dearth of tilts relative to isotropy for $\cos \tau \gtrsim 0.45$. Just as Abbott et al. (2021e), we find that the posterior is not entirely inconsistent with a fully isotropic tilt distribution, while preferring an excess of positive alignment.

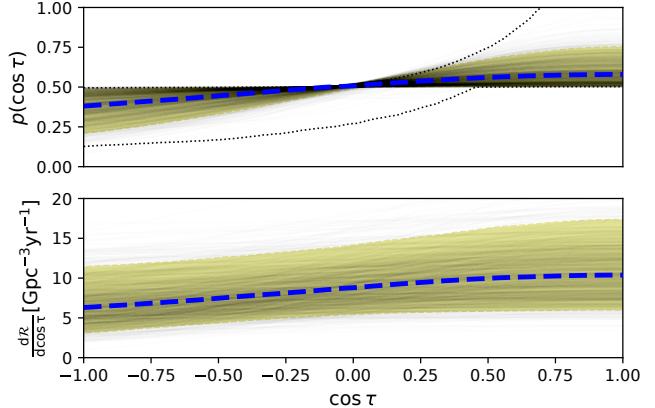


Fig. B.1. Posterior for $\cos \tau$ (top panel) and differential merger rate per unit $\cos \tau$ (bottom panel) obtained using the reference LVK default model. The two thin black dotted lines in the top panel show the 90% CI obtained by drawing the model's hyperparameters from their priors. In both panels, the thin black lines represent individual posterior draws, whereas the colored band shows the 90% CI. The thick dashed line within the band is the median.

This is shown in the top panel of Fig. 4, where the solid blue line was obtained using samples from the hyperparameters' posterior whereas the dashed blue line was obtained by sampling their priors. The fact that there is a hard cutoff at $Y = 1$ (the finite bin size causes the curves to extend to values slightly smaller than one) is just a symptom of the fact that the LVK default model excludes, a priori, an excess of negative tilts and a dearth at positive tilts, as mentioned above.

The curve is consistent with $Y = 1$, that is to say isotropic posteriors are perfectly consistent with the data, even though it should be appreciated that the model prefers that region a priori. The level of consistency can also be assessed with Figure 5, which reports, with dashed blue lines, the marginalized posterior on the branching ratio of the isotropic component (as opposed to the Gaussian component, to allow for direct comparisons with other models). While broad, it favors small values for the fraction of sources in the isotropic component, though fully isotropic distributions ($i = 1$) are not excluded. The other curves in the figure are discussed in the main body. For all of our models, Tab. G.1 reports the Bayesian evidence, maximum log likelihood, and the number of parameters for the $\cos \tau$ model, as a differential relative to the default LVK model. That table also includes a fully isotropic model (Isotropic, with $p(\cos \tau_1, \cos \tau_2) = 1/4$), which we include as a useful reference. The Isotropic model performs the worst, though not at the point that it can be ruled out with high confidence.

Appendix C: Three-component models

The results presented in Sec. 3 show that, depending on the exact model being used, the tilt distribution seems to show either a peak at $+1$, a peak at a smaller positive value of $\cos \tau$, or a broad plateau for positive $\cos \tau$. As this might suggest that two peaks, or features, are present in the data, in this Appendix we consider models that are comprised of an isotropic component, plus two other components. To explore the effect of the model on the resulting posterior, we consider different functional forms.

C.1. Isotropic + Gaussian + Beta model

We used a mixture model with an isotropic component, a beta distribution component, and a Gaussian component,

$$p(\cos \tau_1, \cos \tau_2 | g, b, \mu, \sigma, \alpha, \beta) = \frac{1 - g - b}{4} + b \prod_j^2 \mathcal{B}(\cos \tau_j, \alpha, \beta) + g \prod_j^2 \mathcal{N}(\cos \tau_j, \mu, \sigma). \quad (\text{C.1})$$

In order to reduce degeneracy between the Gaussian and the beta components, we set the uniform prior for the mean of the Gaussian component to $\mu \sim \mathcal{U}(0.9, 5)$; meanwhile, we restricted the prior of the beta parameters to nonsingular values, $\alpha, \beta \sim \mathcal{U}(1, 20)$. In practice, this reduces the possibility that the two components can both create peaks in the same region of the $\cos \tau$ domain, which would increase degeneracy and hence make sampling more inefficient.

The resulting $\cos \tau$ posterior is shown in Fig. C.1. The individual posterior draws are colored according to their value of g ; we stress that small (large) g does not necessarily imply large (small) b since the isotropic fraction $i \equiv 1 - g - b$ needs not be zero. The 90% CI shows traces of the two features we encountered previously, namely a peak at $+1$ and one at smaller positive values of $\cos \tau$. As with all of the other models explored in this paper (and, to our knowledge, in the literature), we find that the data exclude an excess of black holes with $\cos \tau \simeq -1$. The corner plot in Fig. C.2 shows the three branching ratios for this model. The prior for g and b was uniform in the plane, with the constraint that $g + b \leq 1$; we show the resulting marginal priors as dotted black lines in the diagonal panels. This model prefers small values of b coupled with large values of g , as shown in the top-left off-diagonal panel. There is little posterior support for even moderate values of b : the 95th percentile for the

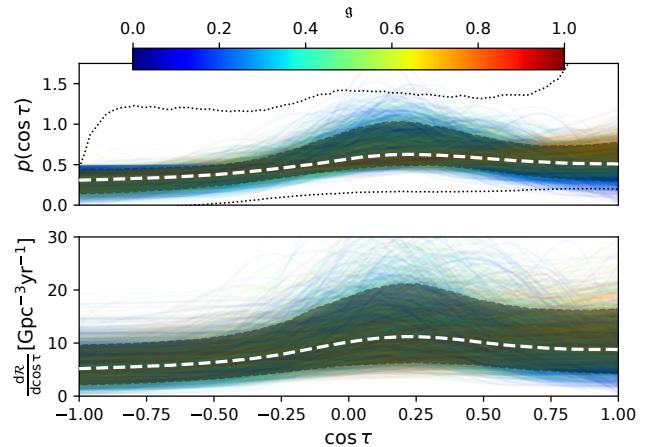


Fig. C.1. Same as Fig. 7, but for the Isotropic + Gaussian + Beta model. Individual posterior draws are colored according to the branching ratio of the Gaussian component, g . It is important to note the different scale for the y axis of the bottom panel compared with similar figures for other models.

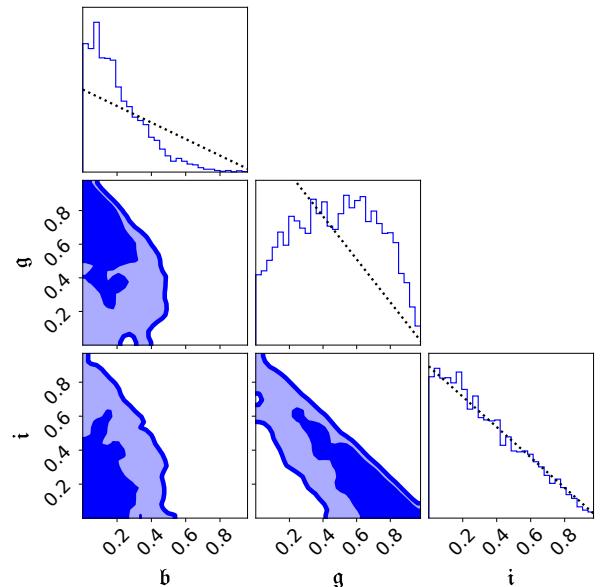


Fig. C.2. Joint and marginal posteriors for the branching ratios of all channels for the Isotropic + Gaussian + Beta model. The thin dashed lines in the diagonal plots are the corresponding priors.

marginal posterior $p(b)$ is 0.50. As already visible in Fig. C.1, the beta component peaks at small positive values of $\cos \tau$: we find $\mu_\beta = 0.15^{+0.39}_{-0.45}$.

C.2. Isotropic + 2 Gaussians model

Next, we used a mixture model, with an isotropic component, and two Gaussian distributions. Here too, to avoid perfect degeneracy, we somewhat restricted the allowed range of the Gaussian means. The Gaussian on the right (index “R”) has a mean that can only vary in the range $[0.9, 5]$. The prior for the Gaussian one on the left (index “L”) spans the range $[-1, 1]$; however, we set to 0 the likelihood for samples for which $\mu_L \notin [A, B]$, where A and B are hyperparameters of the model. In practice, this implies that the Gaussian on the left is truncated (and hence normalized) in the range $[A, B]$ and has a mean in the same range, for each

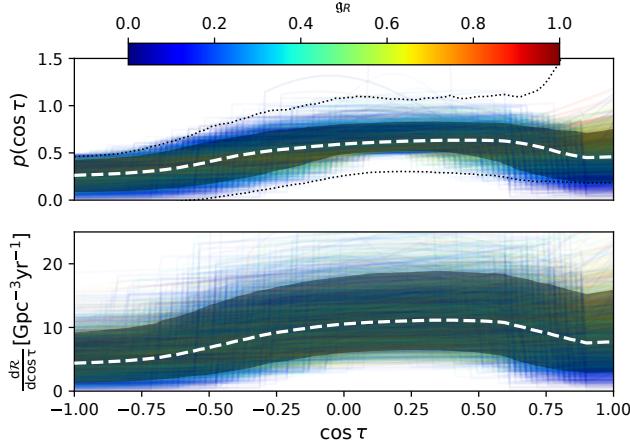


Fig. C.3. Same as Fig. 7, but for the Isotropic + 2 Gaussians model, when the left Gaussian component is truncated in the range $[A, B]$, with A and B model’s hyperparameters. Individual posterior draws are colored according to the branching ratio of the rightmost Gaussian component. It is important to note the different scale for the y axis of the bottom panel compared with similar figures for other models.

sample. Mathematically,

$$p(\cos \tau_1, \cos \tau_2 | g_L, g_R, \mu_L, \sigma_L, \mu_R, \sigma_R, A, B) = \frac{1 - g_L - g_R}{4} + g_L \prod_j^2 \mathcal{N}_{[A, B]}(\cos \tau_j, \mu_L, \sigma_L) + g_R \prod_j^2 \mathcal{N}(\cos \tau_j, \mu_R, \sigma_R). \quad (\text{C.2})$$

We explicitly added hyperparameters for the domain of the left Gaussian in order to verify if the data prefer solutions that do not add posterior support to the antialigned ($\cos \tau \gtrsim -1$) region (cfr. Callister et al. 2022). The resulting posterior for $\cos \tau$ is shown in Fig. C.3, where the individual posterior draws are colored according to the branching ratio of the right Gaussian – g_R . The 90% CI again shows two features: a rather broad peak for a small positive value of $\cos \tau$ and a second peak at +1.

We note that the data are informative for the parameters A and B (Fig. C.4). While their priors are uniform, the posteriors for both A and B show clear peaks. For A , we measured $A = -0.49^{+0.42}_{-0.40}$, which notably excludes -1 at 90% credibility. Meanwhile, the posterior for B rails against +1. The standard deviation for the left Gaussian, σ_L , is large (the 5th percentile of $p(\sigma_L | d)$ is 0.46), which implies that even though functionally speaking this component of our model is a Gaussian, the data seem to prefer very wide Gaussians, resembling pieces of segments. The mean of the left Gaussian prefers small positive values, $\mu_L = 0.13^{+0.52}_{-0.59}$, and shows no obvious correlation with σ_L .

Figure C.5 shows the branching ratios for the two Gaussian and the isotropic component $i \equiv 1 - g_R - g_L$, together with the corresponding priors (dashed lines). The measurement is not precise, and only small departures from the priors are apparent. In particular, for both g_R and g_L , the posteriors yield a wide peak at ~ 0.5 , whereas i peaks at 0 more than the prior.

We stress that the posterior on $\cos \tau$ we obtained for this model is heavily impacted by the fact that the left Gaussian is truncated in a range, whose position is measured from the data. If instead we set $-A = B = 1$, that is to say if we extend (and normalize) the left Gaussian to the full $\cos \tau$ range, we would obtain a radically different posterior (Fig. C.6). The branching ratio for the left Gaussian component in this case does not show

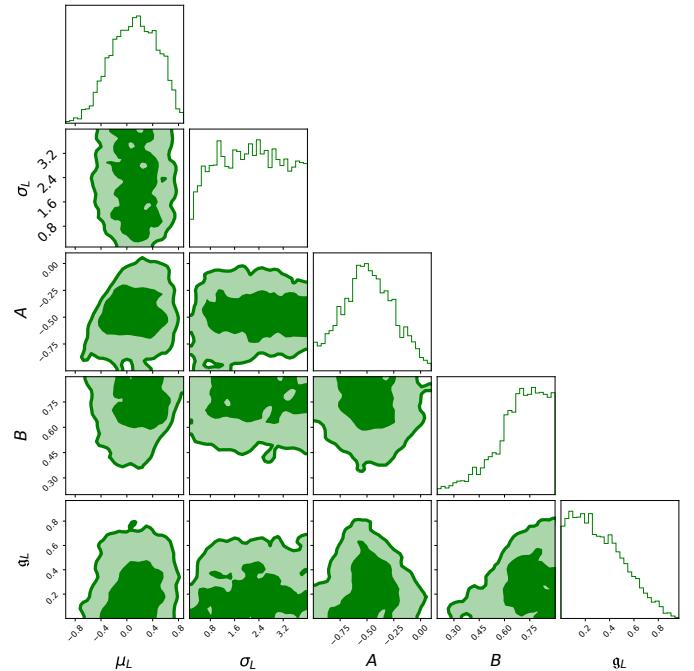


Fig. C.4. Joint and marginal posteriors for the hyperparameters and branching ratio associated with the left Gaussian of the Isotropic + 2 Gaussians model.

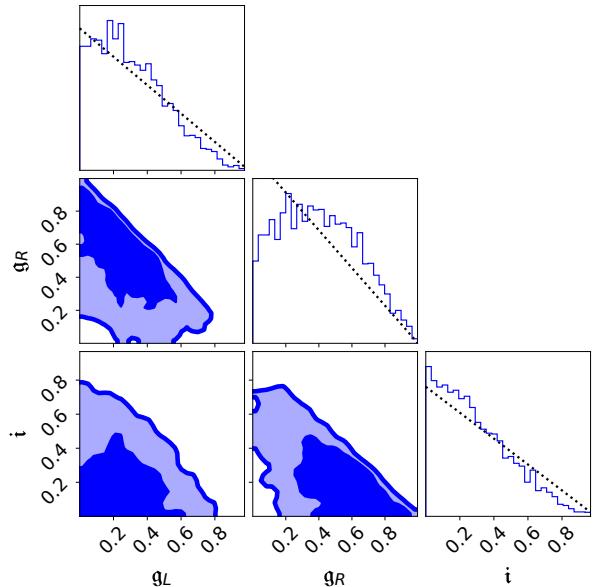


Fig. C.5. Same as Fig. C.2, but for the Isotropic + 2 Gaussians model.

significant differences relative to the prior. While some of the posterior draws show prominent peaks for small positive values of $\cos \tau$, those are not frequent enough to create a visible peak in the 90% CI band, as was instead the case in Fig. C.3.

Given that the only difference between the models behind Fig. C.6 and Fig. C.3 is the truncation of the left Gaussian’s domain, it is tempting to think that the tails of the left Gaussian—if free to extend all the way to $\cos \tau = -1$ —would give too much posterior weight in that region, which is not supported by the data. This explanation is also consistent with the fact that our model of Sec. C.1 does find the peak, since a beta distribution can produce tails that are less wide than a Gaussian.

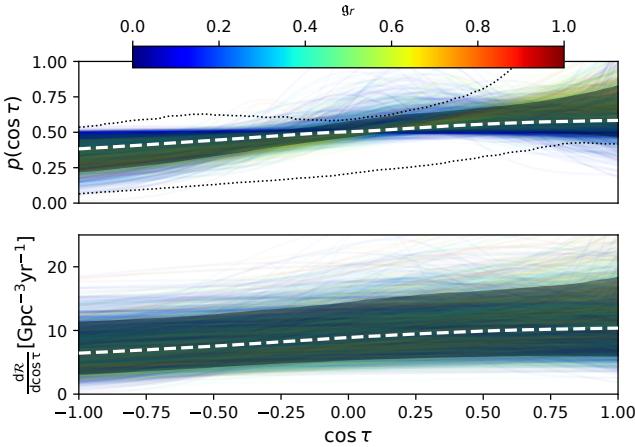


Fig. C.6. Same as Fig. C.3, but without truncating the left Gaussian (i.e., with $-A = B = 1$). Individual posterior draws are colored according to the branching ratio of the rightmost Gaussian component. It is important to note the different scale for the y axis of the bottom panel compared with similar figures for other models.

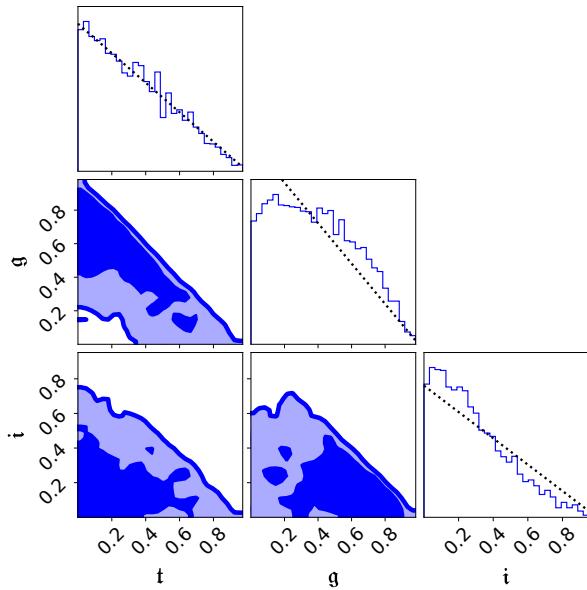


Fig. C.7. Same as Fig. C.2, but for the Isotropic + Gaussian + Tukey model.

C.3. Isotropic + Gaussian + Tukey model

To end our exploration of three-component models, we modified the model of the previous section and replaced the left Gaussian with a distribution based on the Tukey window function. Mathematically,

$$p(\cos \tau_1, \cos \tau_2 | t, g, T_{x0}, T_k, T_r, \mu, \sigma) = \frac{1-t-g}{4} + t \prod_j^2 \mathcal{T}(\cos \tau_j, T_{x0}, T_k, T_r) + g \prod_j^2 \mathcal{N}(\cos \tau_j, \mu, \sigma) \quad . \quad (\text{C.3})$$

The priors for all of the hyperparameters are uniform, with the exception of the branching ratios t and g , which are jointly uniform in the triangle $t + g \leq 1$.

Figure C.7 shows the posteriors for the branching ratios, including that of the isotropic component $i \equiv 1 - t - g$. As for

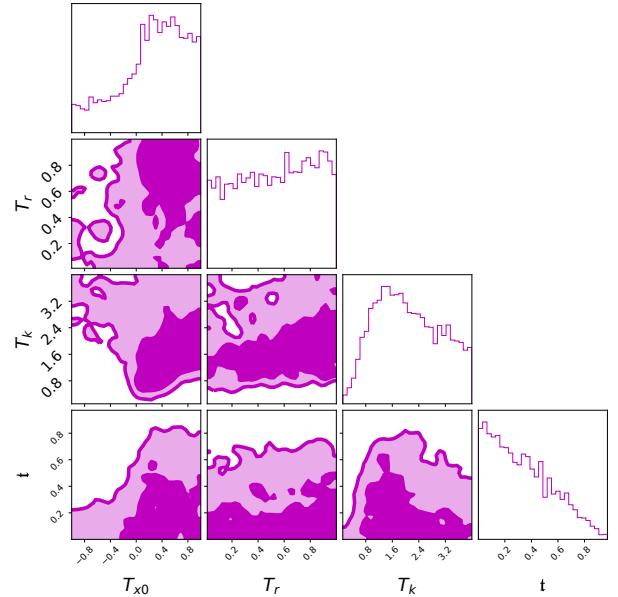


Fig. C.8. Joint and marginal posteriors for the hyperparameters and branching ratio of the Tukey component of the Isotropic + Gaussian + Tukey model.

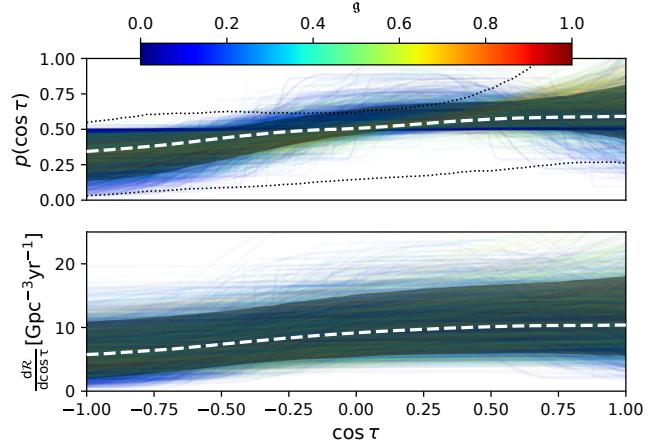


Fig. C.9. Same as Fig. 7, but for the Isotropic + Gaussian + Tukey model. Individual posterior draws are colored according to the branching ratio of the Gaussian component. It is important to note the different scale for the y axis of the bottom panel compared with similar figures for other models.

the Isotropic + 2 Gaussians model, the branching ratios are not measured with precision. The data prefer smaller values of t and g and moderate values of g . Comparing Fig. C.8 with the corresponding plot for the Isotropic + Tukey run (Fig. 8), we find qualitatively consistent results. In particular, T_{x0} mostly has support at positive values, except when T_k can take large values or when t is small. Using the full posterior, we find $T_{x0} = 0.34^{+0.58}_{-1.1}$, while restricting to samples with $T_k \leq 2$ ($T_k \leq 1$) yields $T_{x0} = 0.37^{+0.53}_{-0.81}$ ($T_{x0} = 0.26^{+0.54}_{-0.52}$), which is consistent with the simpler two-component model.

Similarly, we find that the posterior for $\cos \tau$ mainly differs from that of Fig. 7 because of some additional – but not large – support at $\cos \tau = 1$, due to the contribution of the Gaussian component.

Appendix D: Correlated mixture models

In this Appendix we revisit some of the models presented in the main body of the paper, and we extend them to allow for the possibility that the hyperparameters governing the population-level $\cos \tau$ distribution are correlated with some of the astrophysical binary parameters. Previous works have considered correlations between the effective aligned spin, χ_{eff} , and the BBH masses and redshifts (Safarzadeh et al. 2020; Callister et al. 2021; Franciolini & Pani 2022; Biscoveanu et al. 2022; Adamcewicz & Thrane 2022), but not a direct correlation between the tilts and these other intrinsic parameters. Given that the number of BBH sources is still relatively small, we only considered a subset of two-component models in order to keep the number of hyperparameters limited. For some of the correlated models, the distributions for the two tilt angles are not assumed to be identical (this happens when each tilt is allowed to be correlated with the corresponding component mass or spin magnitude): we only report the distribution for the tilt of a primary (i.e., most massive) black hole, as it is usually best measured.

D.1. Isotropic + correlated Gaussian model

We first allowed for the possibility that the mean and standard deviation of the Gaussian component might be correlated with other astrophysical parameters – κ , described below – since those should be related to the details of the supernovae explosions that would have tilted the orbit (e.g., Gerosa et al. 2018). We minimally modified the Isotropic + Gaussian model to allow the mean and standard deviation to linearly vary with the parameter that is correlated to the spin tilt. This introduces another set of hyperparameters, which control the linear part of the mean and standard deviation:

$$p(\cos \tau_1, \cos \tau_2 | \mu_a, \mu_b, \sigma_a, \sigma_b, g, \kappa_1, \kappa_2) = \frac{1-g}{4} + g \prod_j^2 \mathcal{N}(\cos \tau_j, \mu(\kappa_j), \sigma(\kappa_j)), \quad (\text{D.1})$$

where $\mu(\kappa) = \mu_a + \frac{\kappa}{N} \mu_b$ and $\sigma(\kappa) = \sigma_a + \frac{\kappa}{N} \sigma_b$.

It is important to notice that even though we could have also allowed for correlations in the branching ratio g , we decided not to, as that parameter is already very poorly measured (cfr. Fig 5). Similarly, and following Callister et al. (2021) and Biscoveanu et al. (2021), we only considered linear correlations. As more sources are detected, both of these assumptions might be trivially relaxed. The constant N was chosen to guarantee that the coefficient of μ_b, σ_b is always smaller than 1. We explored the following possible correlations:

- **Component masses**, $\kappa_1 = m_1, \kappa_2 = m_2, N = 100 M_\odot$
- **Component spins**, $\kappa_1 = \chi_1, \kappa_2 = \chi_2, N = 1$
- **Mass ratio**, $\kappa_1 = \kappa_2 = q, N = 1$
- **Total mass**, $\kappa_1 = \kappa_2 = m_{\text{tot}}, N = 200 M_\odot$

We found that we cannot constrain, in any significant way, the parameters that enact the correlations (i.e., μ_b and σ_b), for which we recovered posteriors which highly resemble the corresponding priors. This is shown in Fig. D.1, where we report the parameters of the Gaussian component for the analysis where they are correlated with the component masses. The dashed horizontal lines in the diagonal panels represent the corresponding

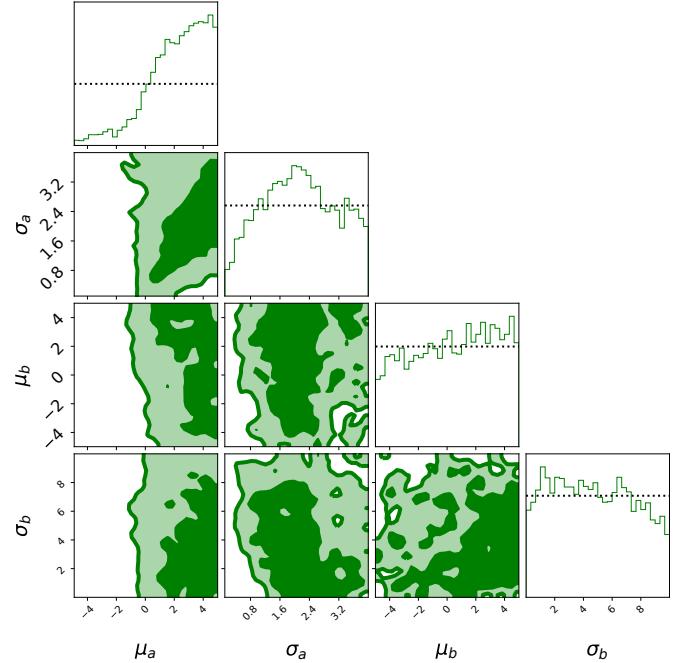


Fig. D.1. Joint and marginal posteriors for the Gaussian parameters obtained in the analysis where they are correlated with the component masses. Dashed lines in the diagonal panels represent the priors. While μ_a and σ_a resemble the corresponding posterior in the Isotropic + Gaussian model (Fig. 2), the terms that enact the correlation, μ_b and σ_b , are nearly unconstrained.

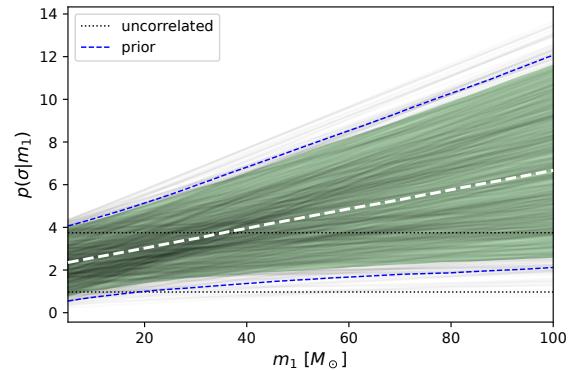


Fig. D.2. Posterior of the standard deviation of the Gaussian component for the Isotropic + correlated Gaussian model, when we allowed for correlation with the component masses, as a function of the primary mass m_1 . The thin black lines are individual posterior draws, the colored band is the 90% CI, and the thick dashed line is the median. The two horizontal think dotted lines enclose the 90% CI for the Isotropic + Gaussian model, which does not allow for correlations. Finally, the two thin blue dashed lines enclose the 90% CI obtained sampling the prior.

priors. Especially for the standard deviation term σ_b , no information is gained relative to the prior. Because we restricted the prior on σ_b to non-negative values (see Tab. G.2) to ensure that the width of the Gaussian does not become negative for any values of κ , this implies that the overall standard deviation for the Gaussian component can only increase with the mass (Fig. D.2). However, as made clear by comparing the 90% CI band with the extent of the 90% CI obtained with prior draws, the increase of σ is entirely prior-driven.

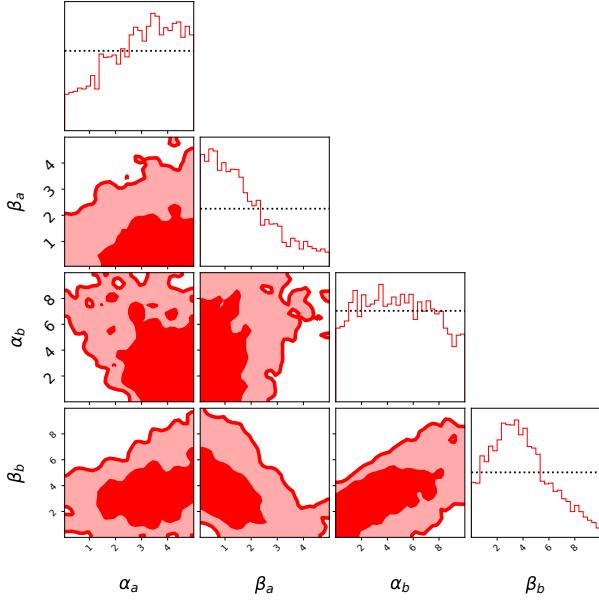


Fig. D.3. Joint and marginal posteriors for the beta parameters obtained in the analysis where they are correlated with the mass ratio. Dashed lines in the diagonal panels represent the priors.

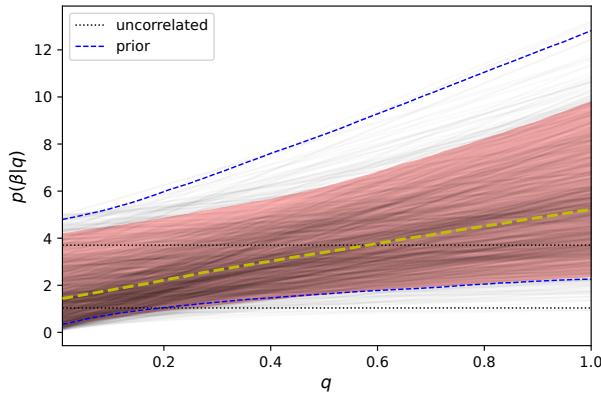


Fig. D.4. Same as D.2, but for the β parameter of the Isotropic + correlated Beta model, when correlated with the mass ratio q .

D.2. Isotropic + correlated Beta model

Finally, we augmented the Isotropic + Beta model of Sec. 3.2 to allow for correlations in the parameters that control the beta component:

$$p(\cos \tau_1, \cos \tau_2 | \alpha_a, \alpha_b, \beta_a, \beta_b, \kappa_1, \kappa_2, b) = \frac{1-b}{4} + b \prod_{j=1}^2 \mathcal{B}(\cos \tau_j, \alpha(\kappa_j), \beta(\kappa_j)), \quad (\text{D.2})$$

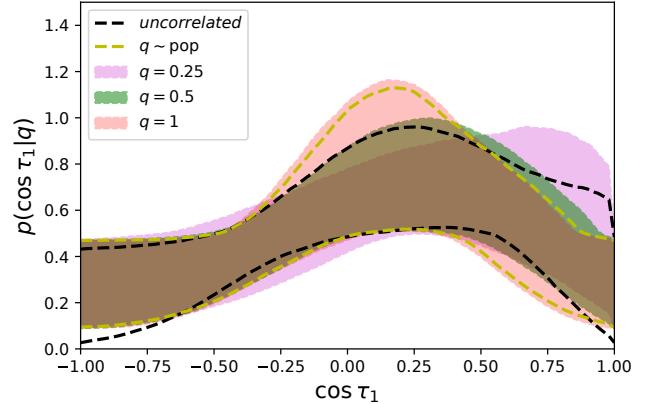


Fig. D.5. Conditional posteriors for $\cos \tau_1$ for the Isotropic + correlated Beta model, when $\cos \tau$ is correlated with the mass ratio. Colored bands show the 90% CI posterior for $\cos \tau_1$ conditional on a fixed value of the mass ratio; yellow dashed lines enclose the 90% CI obtained sampling the correlated parameter from its inferred astrophysical distribution; and black dashed lines enclose the 90% CI of the Isotropic + Beta model, which does not allow for correlations. The increased support at $+1$ as q decreases is mostly prior-driven.

with $\alpha(\kappa) = \alpha_a + \frac{\kappa}{N} \alpha_b$ and $\beta(\kappa) = \beta_a + \frac{\kappa}{N} \beta_b$.

We considered the same possible correlations (i.e., values of κ and N) described in the previous section. As for the previous correlated model, we restricted the priors for α_b and β_b to the non-negative domain to ensure that the overall α and β parameters of the beta distribution do not become negative, enforcing that only positive correlations can exist between the $\cos \tau$ distribution and κ . For this model, we found that the current dataset cannot significantly constrain the correlation parameters, even though we did not exactly recover the priors. For example, in Fig. D.3 we show the posterior and priors (thin dashed lines) for the parameters of the beta distribution when we allowed correlations with the mass ratio q . The parameters that enact the correlation, α_b and β_b , have wide posteriors, which, however, are not as similar as their prior as σ_b was for the Isotropic + correlated Gaussian models (Fig. D.1). Figure D.4 shows that the main impact of the measurement, relative to the prior, is to exclude large values of β . However, it is still the case that the overall trend in the 90% CI of the beta parameters are prior dominated. Just as for the Isotropic + correlated Gaussian models, this results in more support at $\cos \tau \simeq +1$ for small masses, mass ratios, or spins. Functionally, this happens because the parameters controlling the beta component take smaller values at small values of the correlated parameter, and that moves the peak toward the edge of the domain (for example Fig. D.5 for correlations with the mass ratio). However, just as for the Isotropic + correlated Gaussian model, these trends are mostly a result of the prior and of the model.

Appendix E: Tukey window

We implemented the Tukey window used in Eq. 5 as

$$\mathcal{T}(x, T_{x0}, T_k, T_r) \propto \begin{cases} 0, & x < \max(-1, T_{x0} - T_k) \\ \frac{1}{2} \left\{ 1 + \cos \left[\frac{\pi}{T_k T_r} (x - T_{x0} + T_k - T_k T_r) \right] \right\}, & \max(-1, T_{x0} - T_k) \leq x < T_{x0} - T_k(1 - T_r) \\ 1, & T_{x0} - T_k(1 - T_r) \leq x < T_{x0} + T_k(1 - T_r) \\ \frac{1}{2} \left\{ 1 + \cos \left[\frac{\pi}{T_k T_r} (x - T_{x0} - T_k - T_k T_r) \right] \right\}, & T_{x0} + T_k(1 - T_r) \leq x \leq \min(+1, T_{x0} + T_k) \\ 0, & x > \min(+1, T_{x0} + T_k) \end{cases} . \quad (\text{E.1})$$

This represents a Tukey window that is symmetric around T_{x0} and whose domain is $2T_k$ wide. The parameter T_r controls the shape of the window ($T_r = 0$ gives a rectangular window, while $T_r = 1$ gives a cosine). The distribution was then truncated and normalized in the range $[-1, 1]$. Figure E.1 shows four examples. Since the width, the shape, and the position can all be varied, this model is quite elastic and can latch onto both broad and narrow features. We highlight that in the default setting, we allowed the uniform prior of T_k to go up to 4 (Tab G.2). This

implies that just as for the Isotropic + Gaussian model, there are parts of the parameter space where the nonisotropic component can be made very similar to, or indistinguishable from, the isotropic component. In this case, that happens when T_k is large and T_r is small. This distribution can also produce curves that ramp up from zero to a plateau, with various degrees of smoothness: the thick green line in Fig. E.1 is an example and—if T_r were zero—would produce a $\cos \tau$ distribution similar to the second row in Fig. 5 of [Callister et al. \(2022\)](#).

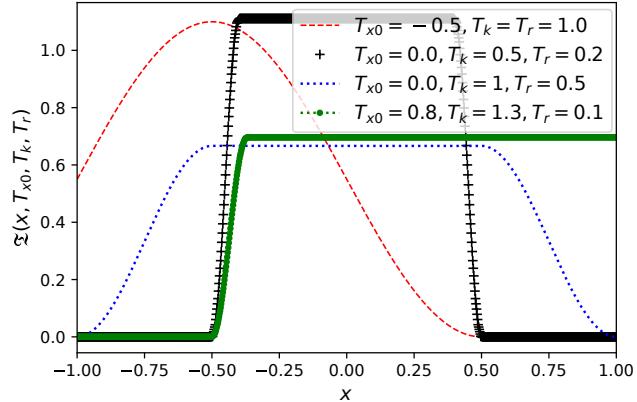


Fig. E.1. Four examples of the distribution in Eq. E.1.

Appendix F: Asymmetry $Y(\delta)$ for various values of δ

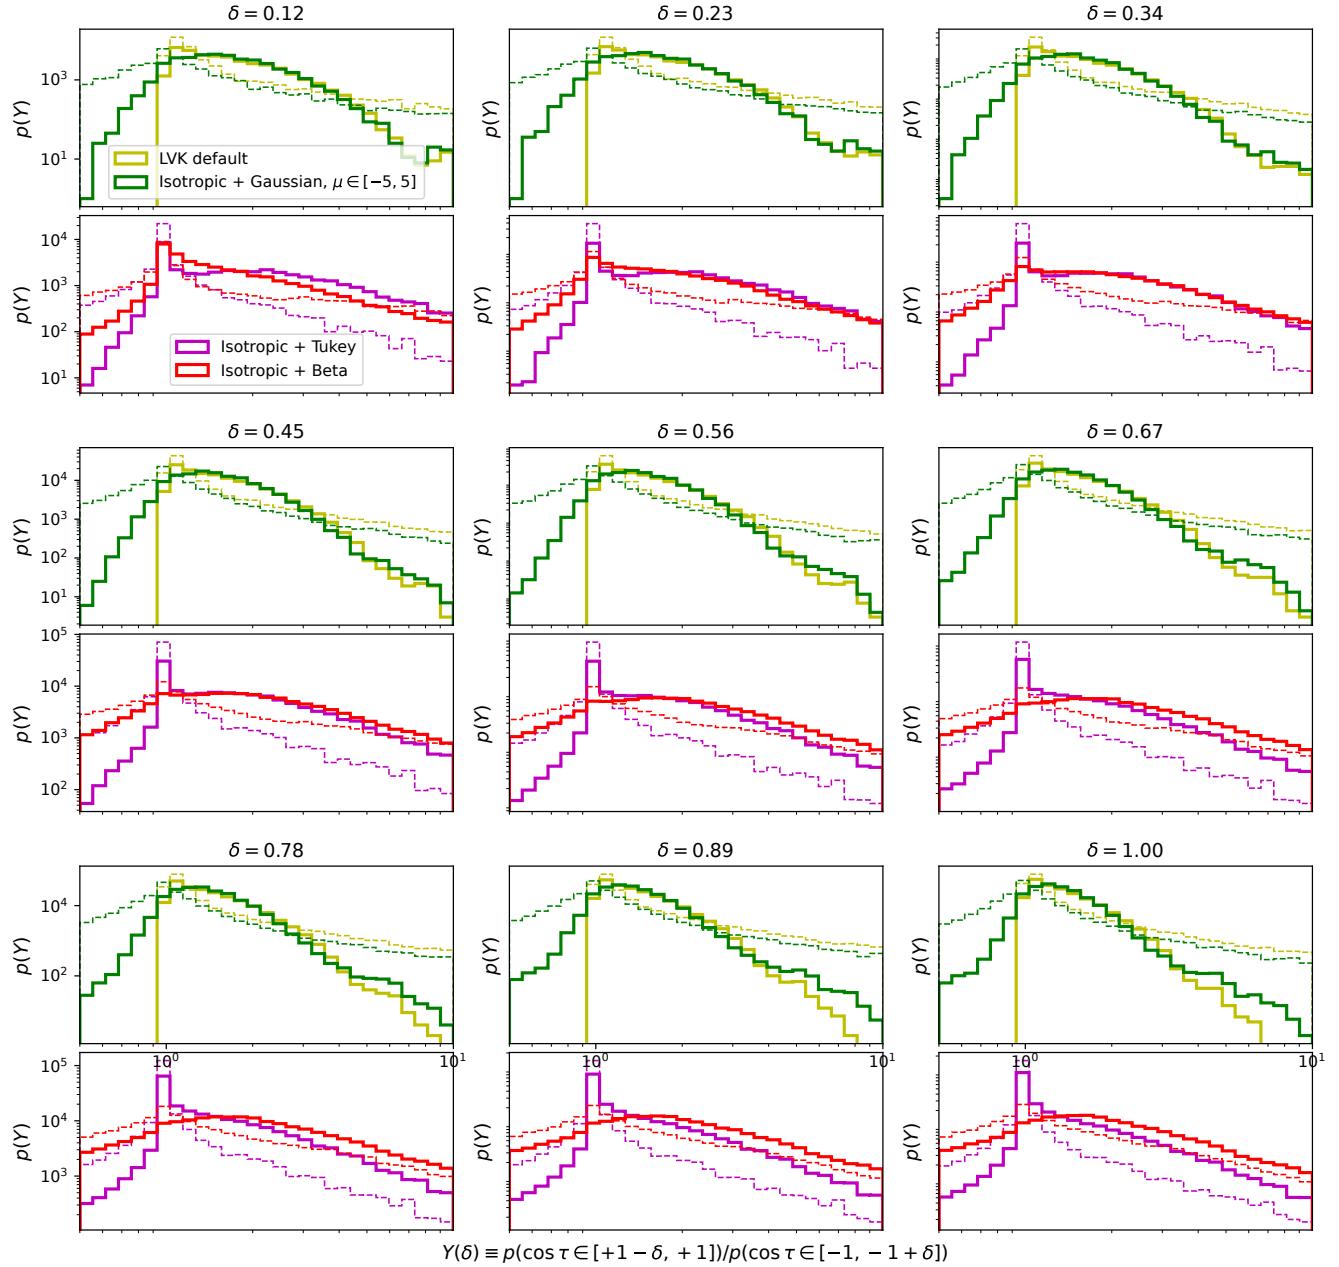


Fig. F.1. Same as Fig. 4, but for various values of δ .

Appendix G: Tables

Table G.1 reports, for all models (including those discussed later in other appendices), the number of spin parameters, the natural log of the Bayesian evidence, and the natural log of the maximum likelihood point. All quantities are expressed as deltas relative to the reference LVK `default` model. Table G.2 lists the priors used for the hyperparameters of all models.

Table G.1. Bayesian evidence, maximum log likelihood value, and number of parameters for the $\cos \tau$ models (relative to the reference LVK model of Eq. B.1).

| Run | InEvidence | InMaxL | # spin pars |
|--|------------|--------|-------------|
| Isotropic | -0.8 | -3.2 | -2 |
| LVK default | ref | ref | ref |
| Isotropic + Gaussian $\mu \in [-1, 1]$ | -0.3 | -0.4 | +1 |
| Isotropic + Gaussian $\mu \in [-5, 5]$ | -0.4 | -0.2 | +1 |
| Isotropic + Beta | +0.4 | +0.4 | +1 |
| Isotropic + Tukey | -0.2 | +0.7 | +2 |
| Isotropic + correlated Gaussian (m) | -0.4 | -0.6 | +3 |
| Isotropic + correlated Gaussian (g) | -0.5 | -0.1 | +3 |
| Isotropic + correlated Gaussian (χ) | -0.6 | -0.4 | +3 |
| Isotropic + correlated Gaussian (M_{tot}) | -0.7 | -0.1 | +3 |
| Isotropic + correlated Beta (m) | -0.4 | +0.0 | +3 |
| Isotropic + correlated Beta (g) | -0.1 | -0.3 | +3 |
| Isotropic + correlated Beta (χ) | +0.3 | -0.6 | +3 |
| Isotropic + correlated Beta (M_{tot}) | -0.3 | -0.7 | +3 |
| Isotropic + Gaussian + Beta | +0.3 | -0.6 | +4 |
| Isotropic + 2 Gaussians ($-A = B = 1$) | -0.1 | -0.3 | +4 |
| Isotropic + Gaussian + Tukey | +0.1 | +0.3 | +5 |
| Isotropic + 2 Gaussians | +0.6 | +0.2 | +6 |

Notes: With our settings, the evidences carry a statistical uncertainty of ± 0.15 . Additional uncertainties in the log likelihood – and hence evidence – arise from the numerical evaluation of the integral in Eq. A.1. Golomb & Talbot (2022) estimate this extra uncertainty to be roughly $\Delta \log \mathcal{L} \pm 1$ when using the publicly released LVK injection sets to estimate selection effects. We also note that nested sampling algorithms do not aim to find the highest likelihood point, so it is possible that for nested models (e.g., LVK `default` and Isotropic + Gaussian) the broader model finds a slightly lower maximum likelihood point. Given these uncertainties, the only reliable – yet unsurprising – conclusions one may draw is that a purely isotropic model yields the worst match to the data.

Table G.2. Priors used for the hyperparameters of the tilt models. All variables are dimensionless.

| LVK default - Eq. B.1 | |
|---|--|
| σ | $\mathcal{U}(0.1,4)$ |
| g | $\mathcal{U}(0,1)$ |
| Isotropic + Gaussian - Eq. 1 | |
| μ | $\mathcal{U}(-5,5)$ or $\mathcal{U}(-1,1)$ |
| σ | $\mathcal{U}(0.1,4)$ |
| g | $\mathcal{U}(0,1)$ |
| Isotropic + Beta - Eq. 3 | |
| α | $\mathcal{U}(0.05,5)$ |
| β | $\mathcal{U}(0.05,5)$ |
| b | $\mathcal{U}(0,1)$ |
| Isotropic + Tukey - Eq. 5 | |
| T_{x0} | $\mathcal{U}(-1,1)$ |
| T_r | $\mathcal{U}(0.01,1)$ |
| T_k | $\mathcal{U}(0.1,4)$ |
| t | $\mathcal{U}(0,1)$ |
| Isotropic + Gaussian + Beta - Eq. C.1 | |
| α | $\mathcal{U}(1,20)$ |
| β | $\mathcal{U}(1,20)$ |
| μ | $\mathcal{U}(0.9,5)$ |
| σ | $\mathcal{U}(0.1,5)$ |
| g, b | $\mathcal{U}(0,1), g + b \leq 1$ |
| Isotropic + 2 Gaussians - Eq. C.2 | |
| μ_L | $\mathcal{U}(-1,1), A < \mu_L < B$ |
| σ_L | $\mathcal{U}(0.1,4)$ |
| μ_R | $\mathcal{U}(0.9,5)$ |
| σ_R | $\mathcal{U}(0.1,5)$ |
| A | $\mathcal{U}(-1,0,1)$ |
| B | $\mathcal{U}(0.2,0,9)$ |
| g_L, g_R | $\mathcal{U}(0,1), g_L + g_R \leq 1$ |
| Isotropic + Gaussian + Tukey - Eq. C.3 | |
| T_{x0} | $\mathcal{U}(-1,1)$ |
| T_r | $\mathcal{U}(0.01,1)$ |
| T_k | $\mathcal{U}(0.1,4)$ |
| μ_R | $\mathcal{U}(0.9,5)$ |
| σ_R | $\mathcal{U}(0.1,5)$ |
| g, t | $\mathcal{U}(0,1), g + t \leq 1$ |
| Isotropic + correlated Gaussian - Eq. D.1 | |
| μ_a | $\mathcal{U}(-5,5)$ |
| μ_b | $\mathcal{U}(-5,5)$ |
| σ_a | $\mathcal{U}(0.1,4)$ |
| σ_b | $\mathcal{U}(0,10)$ |
| g | $\mathcal{U}(0,1)$ |
| Isotropic + correlated Beta - Eq. D.2 | |
| α_a | $\mathcal{U}(0.05,5)$ |
| α_b | $\mathcal{U}(0,10)$ |
| β_a | $\mathcal{U}(0.05,5)$ |
| β_b | $\mathcal{U}(0,10)$ |
| b | $\mathcal{U}(0,1)$ |