

Explaining LIGO's observations via isolated binary evolution with natal kicks

Daniel Wysocki,^{1,*} Davide Gerosa,² Richard O'Shaughnessy,¹ Krzysztof Belczynski,³ Wojciech Gladysz,⁴ Emanuele Berti,^{5,6} Michael Kesden,⁷ and Daniel E. Holz^{8,9}

¹*Rochester Institute of Technology, Rochester, New York 14623, USA*

²*TAPIR 350-17, California Institute of Technology, 1200 E California Boulevard, Pasadena, California 91125, USA*

³*Nicolaus Copernicus Astronomical Centre, Polish Academy of Sciences, Ulica Bartycka 18, 00-716 Warsaw, Poland*

⁴*Astronomical Observatory, Warsaw University, Aleje Ujazdowskie 4, 00-478 Warsaw, Poland*

⁵*Department of Physics and Astronomy, The University of Mississippi, University, Mississippi 38677, USA*

⁶*CENTRA, Departamento de Física, Instituto Superior Técnico, Universidade de Lisboa, Avenida Rovisco Pais 1, 1049 Lisboa, Portugal*

⁷*Department of Physics, The University of Texas at Dallas, Richardson, Texas 75080, USA*

⁸*Enrico Fermi Institute, Department of Physics, Department of Astronomy and Astrophysics, and Kavli Institute for Cosmological Physics, University of Chicago, Chicago, Illinois 60637, USA*

⁹*Kavli Institute for Particle Astrophysics & Cosmology and Physics Department, Stanford University, Stanford, California 94305, USA*



(Received 8 September 2017; published 23 February 2018)

We compare binary evolution models with different assumptions about black-hole natal kicks to the first gravitational-wave observations performed by the LIGO detectors. Our comparisons attempt to reconcile merger rate, masses, spins, and spin-orbit misalignments of all current observations with state-of-the-art formation scenarios of binary black holes formed in isolation. We estimate that black holes (BHs) should receive natal kicks at birth of the order of $\sigma \approx 200$ (50) km/s if tidal processes do (not) realign stellar spins. Our estimate is driven by two simple factors. The natal kick dispersion σ is bounded from above because large kicks disrupt too many binaries (reducing the merger rate below the observed value). Conversely, the natal kick distribution is bounded from below because modest kicks are needed to produce a range of spin-orbit misalignments. A distribution of misalignments increases our models' compatibility with LIGO's observations, if all BHs are likely to have natal spins. Unlike related work which adopts a concrete BH natal spin prescription, we explore a range of possible BH natal spin distributions. Within the context of our models, for all of the choices of σ used here and within the context of one simple fiducial parameterized spin distribution, observations favor low BH natal spin.

DOI: [10.1103/PhysRevD.97.043014](https://doi.org/10.1103/PhysRevD.97.043014)

I. INTRODUCTION

The discovery and interpretation of gravitational waves (GW) from coalescing binaries [1] has initiated a revolution in astronomy [2]. Several hundred more detections are expected over the next five years [3–5]. Already, the properties of the sources responsible—the inferred event rates, masses, and spins—have confronted other observations of black hole (BH) masses and spins [4], challenged previous formation scenarios [2,4], and inspired new models [6–9] and insights [10,11] into the evolution of massive stars and the observationally accessible gravitational waves they emit [12,13]. Over the next several years, our understanding of the lives and deaths of massive stars

over cosmic time will be transformed by the identification and interpretation of the population(s) responsible for coalescing binaries, with and without counterparts, because measurements will enable robust tests to distinguish between formation scenarios with present [14,15] and future instruments [16,17], both coarsely and with high precision. In this work, we demonstrate the power of gravitational wave measurements to constrain how BHs form, within the context of one formation scenario for binary BHs: the isolated evolution of pairs of stars [18–28].

Within the context of that model, we focus our attention on the one feature whose unique impacts might be most observationally accessible: BH natal kicks. Observations strongly suggest that when compact objects like neutron stars are formed after the death of a massive star, their

*dw2081@rit.edu

birth can impart significant linear momentum or “kick.” For example, observations of pulsars in our galaxy suggest birth velocity changes as high as $v_k \sim 450$ km/s [29]. These impulsive momentum changes impact the binary’s intrinsic orbit and stability, changing the orbital parameters like semimajor axis and orbital plane [30,31], as well as causing the center of mass of the remnant BH binary (if still bound) to recoil at a smaller but still appreciable velocity. While no single compelling and unambiguous observation can be explained only with a BH natal kick, the assumption of small but nonzero BH natal kicks provides a natural explanation for several observations, including the posterior spin-orbit misalignment distribution of GW151226 and the galactic x-ray binary misalignment [32–35] and recoil velocity [36–41]. Modest BH natal kicks can be produced by, for example, suitable neutrino-driven supernova engines; see, e.g., [42] and references therein.

We compare binary formation models with different BH natal kick prescriptions to LIGO observations of binary black holes. Along with [42], our calculation is one of the first to perform this comparison while changing a single, physically well-defined and astrophysically interesting parameter: the BH natal kick strength. It is the first to self-consistently draw inferences about binary evolution physics by comparing observations simultaneously to the predicted detection rate; binary BH masses; and binary BH spins, accounting for both magnitude and misalignment.

This comparison is important because BH natal kicks introduce two complementary and unusually distinctive effects on the binary BHs that LIGO detects. On the one hand, strong BH natal kicks will frequently disrupt possible progenitor binary systems. As the strength of BH natal kicks increases, the expected number of coalescing binary BHs drops precipitously [19,20,43]. On the basis of observations to date, BH natal kicks drawn from a distribution with one-dimensional velocity dispersion σ greater than 265 km/s are disfavored [26]. On the other hand, BH natal kicks will tilt the orbital plane, misaligning the orbital angular momentum from the black hole’s natal spin direction—assumed parallel to the progenitor binary’s orbital angular momentum [15,31]. The imprint of these natal kicks on the binary’s dynamics is preserved over the aeons between the BH-BH binary’s formation and its final coalescence [30,44–46]. The outgoing radiation from each merger contains information about the coalescing binary’s spin (see, e.g., [47–49] and references therein), including conserved constants that directly reflect the progenitor binary’s state [50,51]. Several studies have demonstrated that the imprint of processes that misalign BH spins and the orbit can be disentangled [52–54].

In this work, we show that LIGO’s observations of binary black holes can be easily explained in the context of isolated binary evolution, if BH natal kicks act with the (modest) strength to misalign the orbital plane from the initial spin directions (presumed aligned). In this approach, the absence of large aligned spins either reflects fortuitous

but nonrepresentative observations or low natal BH spins. A companion study by Belczynski *et al.* [42] describes an alternative, equally plausible explanation: the BH natal spin depends on the progenitor, such that the most massive BHs are born with low natal spins. A longer companion study by Gerosa *et al.* [55] will describe the properties and precessing dynamics of this population in greater detail.

This paper is organized as follows. First, in Sec. II we describe the entire process used to generate and characterize a detection-weighted populations of precessing binary BHs, evaluated using different assumptions about BH natal kicks. As described in Sec. II A, we adopt previously studied binary evolution calculations to determine how frequently compact binaries merge throughout the universe. In Sec. II B, we describe how we evolve the binary’s precessing BH spins starting from just after it forms until it enters the LIGO band. In Sec. II C, we describe the parameters we use to characterize each binary: the component masses and spins, evaluated after evolving the BH binary according to the process described in Sec. II B. To enable direct comparison with observations, we convert from detection-weighted samples—the output of our binary evolution model—to a smoothed approximation, allowing us to draw inferences about the relative likelihood of generic binary parameters. In Sec. III we compare these smoothed models for compact binary formation against LIGO’s observations to date. We summarize our conclusions in Sec. VI. In Appendix A we describe the technique we use to approximate each of our binary evolution simulations. In Appendix B, we provide technical details of the underlying statistical techniques we use to compare these approximations to LIGO observations. To facilitate exploration of alternative assumptions about natal spins and kicks, we have made publicly available all of the marginalized likelihoods evaluated in this work, as Supplemental Material [56].

II. ESTIMATING THE OBSERVED POPULATION OF COALESCING BINARY BLACK HOLES

A. Forming compact binaries over cosmic time

Our binary evolution calculations are performed with the STARTRACK isolated binary evolution code [20,57], with updated calculation of common-envelope physics [23], compact remnant masses [58], and pair instability supernovae [59]. Using this code, we generate a synthetic universe of (weighted) binaries by Monte Carlo [24]. Our calculations account for the time- and metallicity- dependent star formation history of the universe, by using a grid of 32 different choices for stellar metallicity. As shown in Table I, we create synthetic universes using the same assumptions (M10) adopted by default in previous studies [26,42,59]. Again as in previous work, we explore a one-parameter family of simulations that adopt different assumptions about BH natal kicks (M13–M18). Each new model assumes all BHs receive natal kicks drawn from the same Maxwellian distribution, with one-dimensional velocity distribution parameterized

by σ (a quantity which changes from model to model). In the M10 model used for reference, BH kicks are also drawn from a Maxwellian distribution, but suppressed by the fraction of ejected material that is retained (i.e., does not escape to infinity, instead being accreted by the BH). Because the progenitors of the most massive BHs do not, in our calculations, eject significant mass to infinity, the heaviest BHs formed in this “fallback-suppressed kick” scenario receive nearly or exactly zero natal kicks.

These synthetic universes consist of weighted BH-BH mergers (indexed by i), each one acting a proxy for a part of the overall merger rate density in its local volume [25,39]. As our synthetic universe resamples from the same set of 32 choices for stellar metallicity, the same evolutionary trajectory appears many times, each at different redshifts and reflecting the relative probability of star formation at different times. The most frequent formation scenarios and the fraction of detected binaries from each channel are shown in Table II.

The underlying binary evolution calculations performed by STARTRACK effectively do not depend on BH spins at any stage.¹ We therefore have the freedom to reuse each calculation above with any BH natal spin prescription whatsoever. Unlike Belczynski *et al.* [42], we do not adopt a physically-motivated and mass-dependent BH natal spin, to allow us to explore all of the possibilities that nature might allow. Instead, we treat the birth spin for each BH as a parameter, assigning spins χ_1 and χ_2 to each black hole at birth. For simplicity and without loss of generality, for each event we assume a fixed BH spin for the first-born ($\chi_1 = |\mathbf{S}_1|/m_1^2$) and a potentially different spin for the second-born ($\chi_2 = |\mathbf{S}_2|/m_2^2$) BH. Both choices of fixed spin are parameters. By carrying out our calculations on a discrete grid in χ_1, χ_2 for each event—here, we use $\chi_{1,2} = 0.1\dots 1$ —we encompass a wide range of possible choices for progenitor spins, allowing us to explore arbitrary (discrete) natal spin distributions. For comparison, [53] adopted a fixed natal spin $\chi_i = 0.7$ for all BHs. Our choices for BH natal spin distributions are restricted only by our choice of discrete spins. Our model is also implicitly limited by requiring all BHs have natal spins drawn from the same mass-independent distributions. By design, our calculation did not include enough degrees of freedom to enable the natal spin distribution to change with mass, as was done for example in [42].

¹The response of the BH’s mass and spin to accretion depends on the BH’s spin. We adopt a standard procedure whereby the BH accretes from the innermost stable circular orbit. In our binary evolution code, this spin evolution is implemented directly via an ODE based on (prograde, aligned) ISCO accretion as in [60], though the general solution is provided in [61] and applied since, e.g., in [62,63]. For the purposes of calculating the final BH mass from the natal mass and its accretion history, we adopted a BH natal spin of $\chi = 0.5$; however, relatively little mass is accreted and the choice of spin has a highly subdominant effect on the BH’s evolution.

TABLE I. Properties of the formation scenarios adopted in this work. The first column indicates the model calculation name, using the convention of other work [26,59]. The second column provides the kick distribution width. Model M10 adopts mass-dependent, fallback suppressed BH natal kicks. For the BH population examined here, these natal kicks are effectively zero for massive BHs; see, e.g., [14]. The remaining scenarios adopt a mass-independent Maxwellian natal kick distribution characterized by the 1-d velocity dispersion σ , as described in the text. The third column quantifies how much the mass distribution changes as we change σ . To be concrete, we compare the (source frame) total mass distributions for the BH-BH binaries LIGO is expected to detect, using a Kullback–Leibler (KL) divergence [Eq. (4)]. If $p(M|\alpha)$ denotes the mass distribution for $\alpha = \text{M10, M18, M17, } \dots$, and α_* denotes M17, then the third column is the KL divergence $D_{KL}(M, \alpha) = \int dM p(M|\alpha) \ln[p(M|\alpha)/p(M|\alpha_*)]$. The fourth column is the KL divergence using the joint distribution of both binary masses: $D_{KL}(m_1, m_2|\alpha) = \int dm_1 dm_2 p(m_1, m_2|\alpha) \ln[p(m_1, m_2|\alpha)/p(m_1, m_2|\alpha_*)]$. Because M10 adopts fallback-suppressed natal kicks, while the remaining models assume fallback-independent natal kicks, we use the special symbol \emptyset to refer to M10 in subsequent plots and figures.

Name	σ (km/s)	$D_{KL}(M)$	$D_{KL}(m_1, m_2)$
M10	\emptyset	0.02	0.21
M18	25	0.006	0.094
M17	50	0	0
M16	70	0.016	0.28
M15	130	0.1	1.26
M14	200	0.17	1.56
M13	265	0.40	2.1

We assume the progenitor stellar binary is comprised of stars whose spin axes are aligned with the orbital angular momentum, reflecting natal or tidal [64,65] alignment (but cf. [66]). After the first supernova, several processes could realign the stellar or BH spin with the orbital plane, including mass accretion onto the BH and tidal dissipation in the star. Following Gerosa *et al.* [30], we consider two possibilities. In our default scenario (“no tides”), spin-orbit alignment is only influenced by BH natal kicks. In the other scenario (“tides”), tidal dissipation will cause the stellar spin in stellar-BH binaries to align parallel to the orbital plane. In the “tides” scenario, the second-born stellar spin is aligned with the orbital angular momentum prior to the second SN. Following [30], the “tides” scenario assumes alignment always occurs for merging BH-BH binaries, independent of the specific evolutionary trajectory involved (e.g., binary separation); cf. the discussion in [42]. In both formation scenarios, we do not allow mass accretion onto the BH to change the BH’s spin direction. Given the extremely small amount of mass accreted during either conventional or common-envelope mass transfer, even disk warps and the Bardeen-Petterson effect should not allow the BH spin direction to evolve [67–70]. For coalescing BH-BH binaries the second SN often occurs when the binary is in a tight orbit, with high orbital speed, and thus

TABLE II. The most significant formation scenarios and fraction of detected binaries formed from that channel, for the M15 model. While many of the coalescing BH-BH binaries form via a BH-star binary undergoing some form of stellar mass transfer or interaction, a significant fraction of binaries form without any Roche lobe overflow mass transfer after the first SN. In this example, in the second channel alone more than **23%** of binaries form without interaction after the first SN. (The remaining formation channels account for **16%** of the probability.) In this notation, integers in braces characterize the types of the stellar system in the binary; the prefix refers to different phases of stellar interaction (e.g., MT denotes “mass transfer,” SN denotes “supernova,” and CE denotes “common envelope evolution”); and the last integer SN x indicates whether the initial primary star (1) or initial secondary star (2) has collapsed and/or exploded to form a BH. [Some of our BHs are formed without luminous explosions; we use SN to denote the death of a massive star and the formation of a compact object.] A detailed description of these formation channels and stellar types notation is provided in [20,57]; in this shorthand, 1 denotes a main sequence star; 2 denotes a Hertzsprung gap star; 4 denotes a core helium burning star; 7 denotes a main sequence naked helium star; and 14 denotes a black hole.

Formation mechanism	Fraction
MT1(2-1) MT1(4-1) SN1 CE2(14-4;14-7) SN2	0.261
MT1(4-4) CE2(7-4;7-7) SN1 SN2	0.234
MT1(4-1) SN1 CE2(14-4;14-7) SN2	0.140
MT1(2-1) SN1 CE2(14-4;14-7) SN2	0.075
MT1(4-4) CE2(4-4;7-7) SN1 SN2	0.071
MT1(2-1) SN1 MT2(14-2) SN2	0.037
CE1(4-1;7-1) SN1 MT2(14-2) SN2	0.028
CE1(4-1;7-1) SN1 CE2(14-4;14-7) SN2	0.020
CE1(4-1;7-1) CE2(7-4;7-7) SN1 SN2	0.014
MT1(4-4) CE12(4-4;7-7) SN2 SN1	0.014
SN1 CE2(14-4;14-7) SN2	0.014
Other channels	0.16

less effect on spin-orbit misalignment [15,31]. Therefore, in the “tides” scenario, the second-born BH’s spin is more frequently nearly aligned with the final orbital plane, even for large BH natal kicks.

B. Evolving from birth until merger

The procedure above produces a synthetic universe of binary BHs, providing binary masses, spins, and orbits just after the second BH is born. Millions to billions of years must pass before these binaries coalesce, during which time the orbital and BH spin angular momenta precess substantially [45,46]. We use precession-averaged 2PN precessional dynamics, as implemented in PRECESSION [71], to evolve the spins from birth until the binary BH orbital frequency is 10 Hz (i.e., until the GW frequency is 20 Hz); see [55] for details. When identifying initial conditions, we assume the binary has already efficiently circularized. When identifying the final separation, we only use the Newtonian-order relationship between separation

and orbital frequency. The PRECESSION code is publicly available at github.com/dgerosa/precession.

C. Characterizing the observed distribution of binaries

At the fiducial reference frequency adopted in this work (20 Hz), a binary BH is characterized by its component masses and its (instantaneous) BH spins $\mathbf{S}_{1,2}$. For the heavy BHs of current interest to LIGO, the principal effect of BH spin on the orbit and emitted radiation occurs through the spin combination

$$\begin{aligned} \chi_{\text{eff}} &= (\mathbf{S}_1/m_1 + \mathbf{S}_2/m_2) \cdot \hat{\mathbf{L}} / (m_1 + m_2) \\ &= (\chi_1 m_1 \cos \theta_1 + \chi_2 m_2 \cos \theta_2) / (m_1 + m_2), \end{aligned} \quad (1)$$

where $\theta_{1,2}$ denote the angles between the orbital angular momentum and the component BH spins. That said, depending on the duration and complexity of the source responsible, GW measurements may also provide additional constraints on the underlying spin directions themselves [50], including on the spin-orbit misalignment angles $\theta_{1,2}$. For the purposes of this work, we will be interested in the (source-frame) binary masses m_1, m_2 and the spin parameters $\chi_{\text{eff}}, \theta_1, \theta_2$, as an approximate characterization of the most observationally accessible degrees of freedom; cf. Stevenson *et al.* [53], which used $\theta_{1,2}$, and Trifirò *et al.* [50], which used $\theta_{1,2}$ and the angle $\Delta\Phi$ between the spins’ projection onto the orbital plane. In particular, $\Delta\Phi$ is well-known to contain valuable information [30] and be observationally accessible [50]. At present, the preferred model adopted for parameter inference, known as IMRPhenomP, does not incorporate the necessary degree of freedom [72], so we cannot incorporate its effect here. With additional and more informative binary black hole observations, however, our method should be extended to employ all of the spin degrees of freedom, particularly $\Delta\Phi$. As input, this extension will require inference results that incorporate the effect of two precessing spins, either by using semianalytical models [73–75] or by using numerical relativity [49].

We adopt a conventional model for LIGO’s sensitivity to a population of binary BHs [3,25,76]. In this approach, LIGO’s sensitivity is limited by the second-most-sensitive interferometer, using a detection threshold signal-to-noise ratio $\rho = 8$ and the fiducial detector sensitivity reported for O1 [4]. This sensitivity model is a good approximation to the performance reported for both in O1 and early in O2 [5]. Following [25,39], we use the quantity r_i (Eq. (8) in [39]) to account for the contribution of this binary to LIGO’s detection rate in our synthetic universe, accounting for the size of the universe at the time the binary coalesces and LIGO’s orientation-dependent sensitivity. For simplicity and following previous work [3,25], we estimate the detection probability without accounting for the effects of BH spin. Previous studies have used this detection-weighted procedure to evaluate and report on the expected

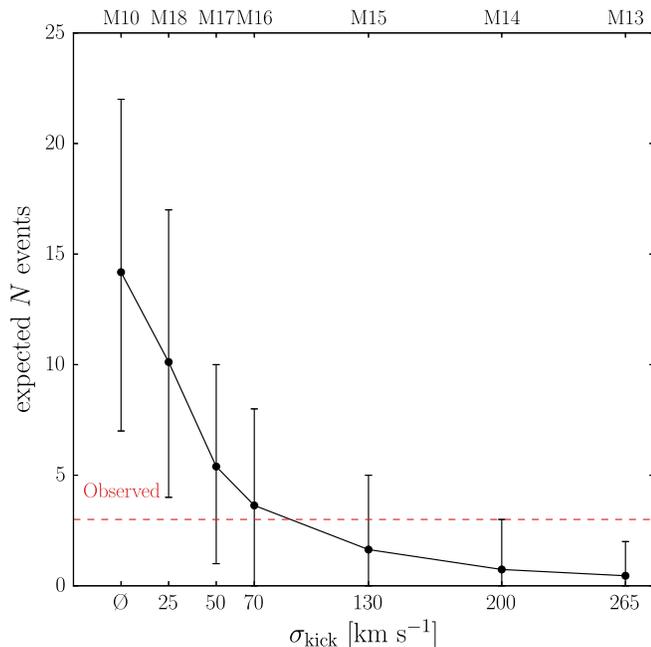


FIG. 1. *Expected number of events versus kick strength:* Expected number of BH merger detections predicted at LIGO's O1 sensitivity and for the duration of O1 by our formation scenarios. The predicted number of events decreases rapidly as a function of the BH natal kick. Also shown is the 95% confidence interval, assuming Poisson distribution with mean predicted by our model. This purely statistical error bar does not account for any model systematics (e.g., in the overall star formation rate and metallicity history of the universe). The horizontal red dashed line corresponds to the number (3) of observations reported in O1 [4].

distribution of binary BHs detected by LIGO [26,42,59]. Since the same binary evolution A occurs many times in our synthetic universe, we simplify our results by computing one overall detection rate $r_A = \sum_{i \in A} r_i$ for each evolution. When this procedure is performed, relatively few distinct binary evolutions A have significant weight. While our synthetic universe contains millions of binaries, only $O(10^4)$ distinct BH-BH binaries are significant in our final results for each of the formation scenarios listed in Table I. Figure 1 illustrates the expected detected number versus assumed BH natal kick strength.

The significant BH natal kicks adopted in all of our formation scenarios (except M10) frequently produce significant spin-orbit misalignment. Figure 2 shows that strong misalignment occurs ubiquitously, even for small BH natal kicks; see [55] for more details. This strong spin-orbit misalignment distribution produces an array of observationally accessible signatures, most notably via an invariably wide distribution of χ_{eff} . In [55] the distribution was constructed for all of our models, finding that (except for M10) considerable support exists for $\chi_{\text{eff}} < 0$. Our calculation is fully consistent with the limited initial exploration reported in Rodriguez *et al.* [14], which claimed $\chi_{\text{eff}} < 0$ was implausible except for extremely large natal kicks. Their collection of calculations explored

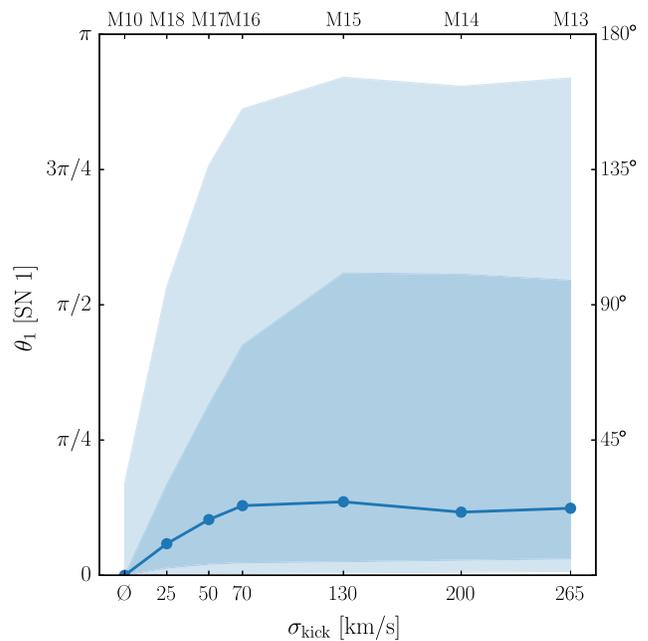


FIG. 2. *Spin-orbit misalignment versus kick strength:* The misalignment $\theta_{1,\text{SN1}}$ after the first SN event, as a function of the characteristic BH natal kick σ . (Note $\theta_{1,\text{SN1}}$ should be distinguished from θ_1 described in the text: θ_1 is the angle between the more massive BH and the orbital angular momentum, at 20 Hz.) The solid line shows the median value; shaded region shows the 68% and 95% confidence intervals.

fallback-suppressed kicks (e.g., equivalent to our model M10); adopted natal kicks larger than we explored here; or adopted mass-dependent natal kicks. We show that significant spin-orbit misalignment is plausible if all BHs—even massive ones—receive a modest natal kick. BH natal kicks therefore provide a robust mechanism to explain the observed χ_{eff} and spin-orbit misalignments reported by LIGO for its first few detections.

The procedure described above *samples* a synthetic universe and synthetic observations by LIGO. However, to compare to LIGO's observations, we need to be able to assess the likelihood of *generic* binaries according to our formation scenario, extrapolating between what we have simulated. We therefore estimate the merger rate distribution as a function of binary masses, spins, and spin-orbit misalignments. Our estimate uses a carefully calibrated Gaussian mixture model, with special tuning as needed to replicate sharp features in our mass and misalignment distribution; see Appendix A for details.

III. COMPARISON WITH GRAVITATIONAL WAVE OBSERVATIONS

A. Gravitational wave observations of binary black holes

During its first observing run of $T_1 = 48.6$ days, LIGO has reported the observation of three BH-BH mergers: GW150914, LVT151012, and GW151226 [1,4,77]. In an

analysis of $T_2 = 11$ days of data from its second observing run, at comparable sensitivity, LIGO has since reported the observation of another binary BH: GW170104 [5]. To draw out more insight from each observation, rather than use the coarse summary statistics LIGO provides in tabular form, we employ the underlying posterior parameter distribution estimates provided by the LIGO Scientific Collaboration for the three O1 events [4,47,48]. For GW170104, we instead adopt an approximate posterior distribution described in Appendix C based solely on reported tabular results; that said, we are confident that this approximation makes no difference to our conclusions. For each event, for brevity indexed by an integer $n = 1, 2, 3, \dots, N$, these estimates are generated by comparing a proposed gravitational wave source x with the corresponding stretch of gravitational wave data d using a (Gaussian) likelihood function $p(d|x)$ that accounts for the frequency-dependent sensitivity of the detector (see, e.g., [47–49] and references therein). In this expression x is shorthand for the 15 parameters needed to fully specify a quasicircular BH-BH binary in space and time, relative to our instrument; and d denotes all the gravitational wave data from all of LIGO’s instruments. This analysis adopts prior assumptions about the relative likelihood of different progenitor binaries $p_{\text{ref}}(x)$: equally likely to have any pair of component masses, any spin direction, any spin magnitude, any orientation, and any point in spacetime (i.e., uniform in comoving volume). Then, using standard Bayesian tools [47,48], the LIGO analysis produced a sequence of independent, identically distributed samples $x_{n,s}$ ($s = 1, 2, \dots, S$) from the posterior distribution for each event n ; that is, each $x_{n,s}$ is drawn from a distribution proportional to $p(d_n|x_n)p_{\text{ref}}(x_n)$. This approach captures degeneracies in the posterior not previously elaborated in detail, most notably the well-known strong correlations between the inferred binary’s component masses and spins (e.g., between χ_{eff} and m_2/m_1).² Equivalently, this approach gives us direct access to properties of the posterior distribution that were not reported in published tables [4], most notably for the relative posterior probabilities of different choices for binary BH spins (e.g., the data underlying Fig. 3).

B. Comparing models to observations

The overall likelihood of GW data $\{d\}$ using a model parametrized by Λ is [78]

²Different properties of the binary, like the masses and spins, influence the inspiral, and thus the radiation $h(t)$, in generally different ways; however, sometimes, several parameters can influence the radiation in a similar or degenerate way. For example, both the binary mass ratio and (aligned) binary spin can extend the duration of the inspiral. Similarly, both the binary masses and spins—8 parameters—determine the final complex frequency of the BH—at leading order, only set by two parameters. Due in part to degeneracies like these, LIGO’s inferences about the parameters x for each merging BH lead to a highly correlated likelihood $p(d|x)$ and hence posterior distribution; see, e.g., [47–49] and references therein.

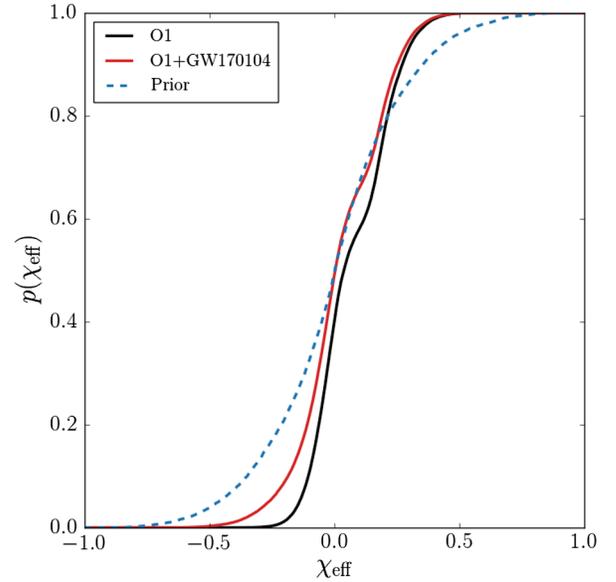


FIG. 3. *Empirical cumulative distribution function for χ_{eff}* : The solid blue line shows the conventional prior distribution for χ_{eff} , generated by selecting masses uniformly in $m_{1,2} \geq 1M_{\odot}$, $m_{1,2} \leq 100M_{\odot}$, $m_1 + m_2 < 100M_{\odot}$, and isotropic spins generated independently and uniformly in magnitude. This prior was adopted when analyzing all LIGO events. The solid black line shows the empirical cumulative distribution for χ_{eff} , derived from the collection of events $\alpha = \text{GW150914, GW151226, and LVT151012}$ via their posterior cumulative distributions $P_{\alpha}(\chi_{\text{eff}})$ via $P(\chi_{\text{eff}}) = \sum_{\alpha} P_{\alpha}(\chi_{\text{eff}})/3$. In this curve, the posterior distributions are provided by LIGO’s full O1 analysis results [4], as described in the text. The solid red line shows the corresponding result when GW170104 is included. The approximate posterior distribution for GW170104 is based on published results, as described in Appendix C.

$$p(\{d\}|\Lambda) \propto e^{-\mu} \prod_n \int dx_n p(d_n|x_n) \mathcal{R}p(x_n|\Lambda) \quad (2)$$

where x_n denote candidate intrinsic and extrinsic parameters for the n th observation, μ is the expected number of detections according to the formation scenario Λ , $p(d_n|x_n)$ is the likelihood for event n ; $p(\{d\}|\Lambda)$ is the marginalized likelihood; $p(x_n|\Lambda)$ is the prior evaluated at event n ; and \mathcal{R} (implicitly depending on Λ as well) is the average number of merger events per unit time and volume in the Universe. In this expression, we have subdivided the data $\{d\}$ into data with confident detections d_1, d_2, \dots, d_N and the remaining data; the Poisson prefactor $\exp(-\mu)$ accounts for the absence of detections in the remaining data; and the last product accounts for each independent observation d_n . Combined, the factors $e^{-\mu} \prod_n \mathcal{R}p(x_n)$ are the distribution function for an inhomogeneous Poisson process used to characterize the formation and detection of coalescing BH binaries [79,80]. As described in Appendix B, the probability density functions $p(x|\Lambda)$ are estimated from the weighted samples that define each synthetic universe Λ , and the integrals $\int p(d|x)p(x|\Lambda)$ are performed efficiently

via Monte Carlo integration. Similarly, the expected number of detections μ at O1 sensitivity—a known constant for each model Λ —is already provided by the detailed cosmological integration performed in prior work; see Sec. II and Fig. 1. Since the marginal likelihood can always be evaluated, the model inference on our discrete set of models becomes an application of Bayesian statistics. In this work, we report the Bayes factor or likelihood ratio $K_{ij} = p(\{d\}|\Lambda_i)/p(\{d\}|\Lambda_j)$ between two different sets of assumptions. To fix the zero point for the log Bayes factor, we adopt the M16 model with $\chi_1 = \chi_2 = 0.5$, henceforth denoted collectively as J , and henceforth use $\ln K$ as shorthand for $\ln K_{ij}$.

In what follows, we will mainly discuss comparisons of our models to all of LIGO's reported detection candidates in O1: GW150914, GW151226, and LVT151012 [4]. We do this because LIGO's O1 observational time and survey results are well-defined and comprehensively reported [4]; because we can employ detailed inference results for all O1 events; and because, as we show below, adding GW170104 to our analysis produces little change to our results. Using the approximate posterior described in Appendix C for GW170104, we will also compare all reported LIGO observations (O1 and GW170104) to our models.

Critically, for clarity and to emphasize the information content of the data, in several of our figures we will illustrate the marginal likelihood of the data $p(\{d\}|\Lambda)$ evaluated assuming all binaries are formed with *identical* natal spins. These strong assumptions in our illustrations show just how much the data informs our understanding of BH natal spins. With only four observations, assumptions about the spin distribution are critical to make progress. As described in Appendix B, we can alternatively evaluate the marginal likelihood accounting for any concrete spin distribution, or even all possible spin distributions—in our context, all possible mixture combinations of the 100 different choices for χ_1 and χ_2 that we explored. In the latter case, as we show below, just as one expects *a priori*, observations cannot significantly inform this 100-dimensional posterior spin distribution. As suggested in previous studies [e.g., [2,42,49,54]], LIGO's observations in O1 and O2 can be fit by models that includes a wide range of progenitor spins, so long as sufficient probability exists for small natal spin and/or significant misalignment. As a balance between complete generality on the one hand (a 100-dimensional distribution of natal spin distributions) and implausibly rigid assumptions on the other (fixed natal spins), we emphasize a simple one-parameter model, where BH natal spins χ are drawn from the piecewise constant distribution

$$p(\chi) = \begin{cases} \lambda_A/0.6 & \chi \leq 0.6 \\ (1 - \lambda_A)/0.4 & 0.6 < \chi < 1 \end{cases} \quad (3)$$

where λ_A is the probability of a natal spin ≤ 0.6 and the choice of cutoff 0.6 is motivated by our results below.

IV. RESULTS

In this section we calculate the Bayes factor $\ln K$ for each of the binary evolution models described above. Unless otherwise noted, we compare our models to LIGO's O1 observations (i.e., the observation of GW150914, GW151226, and LVT151012), using each model's correlated predictions for the event rate, joint mass distribution (m_1, m_2) , χ_{eff} distribution, and the distribution of θ_1, θ_2 . For numerical context, a Bayes factor of $\ln 10 \approx 2.3$ is by definition equivalent to 10:1 odds in favor of some model over our reference model. Bayes factors that are more than 5 below the largest Bayes factor observed are in effect implausible (e.g., more than 148:1 odds against), whereas anything within 2 of the peak are reasonably likely.

A. Standard scenario and limits on BH natal spins (O1)

The M10 model allows us to examine the implications of binary evolution with effectively zero natal kicks. The M10 model adopts fiducial assumptions about binary evolution and BH natal kicks, as described in prior work [26,59]. In this model, BH kicks are suppressed by fallback; as a result, the heaviest BHs receive nearly or exactly zero natal kicks and hence have nearly or exactly zero spin-orbit misalignment.

If heavy BH binaries have negligible spin-orbit misalignment, then natal BH spins are directly constrained from LIGO's measurements (e.g., of χ_{eff}). For example, LIGO's observations of GW150914 severely constrain its component spins to be small, if the spins must be strictly and positively aligned [48,49]. Conversely, however, LIGO's observations for GW151226 require some nonzero spin. Combined, if we assume all BHs have spins drawn from the same, mass-independent distribution and have negligible spin-orbit misalignments, then we conclude BH natal spins should be preferentially small. [We will return to this statement in Sec. IV D.]

Figure 4 shows one way to quantify this effect within the context of our calculations. The left panel shows the Bayes factor for all of our formation models (including M10) as a function of BH natal spin, assuming all BHs have the same (fixed) natal spin $\chi = \chi_1 = \chi_2$. As expected from LIGO's data, large natal BH spins cannot be adopted with M10 and remain consistent with LIGO's observations. The right panel shows the Bayes factor for M10 as function of both BH natal spins, allowing the more massive and less massive BHs to receive different (fixed) natal spins. [The blue line on the left panel uses precisely the same data as the diagonal $\chi_1 = \chi_2$ on the right.] The colorscale graphically illustrates the same conclusion: though marginally greater freedom exists for natal BH spin on the smaller of the two BHs, we can rule out that all BHs, independent of their mass, have significant natal spin if M10 is true. Conversely, if M10 is true and all BHs have the same natal spins, then this natal spin is likely small.

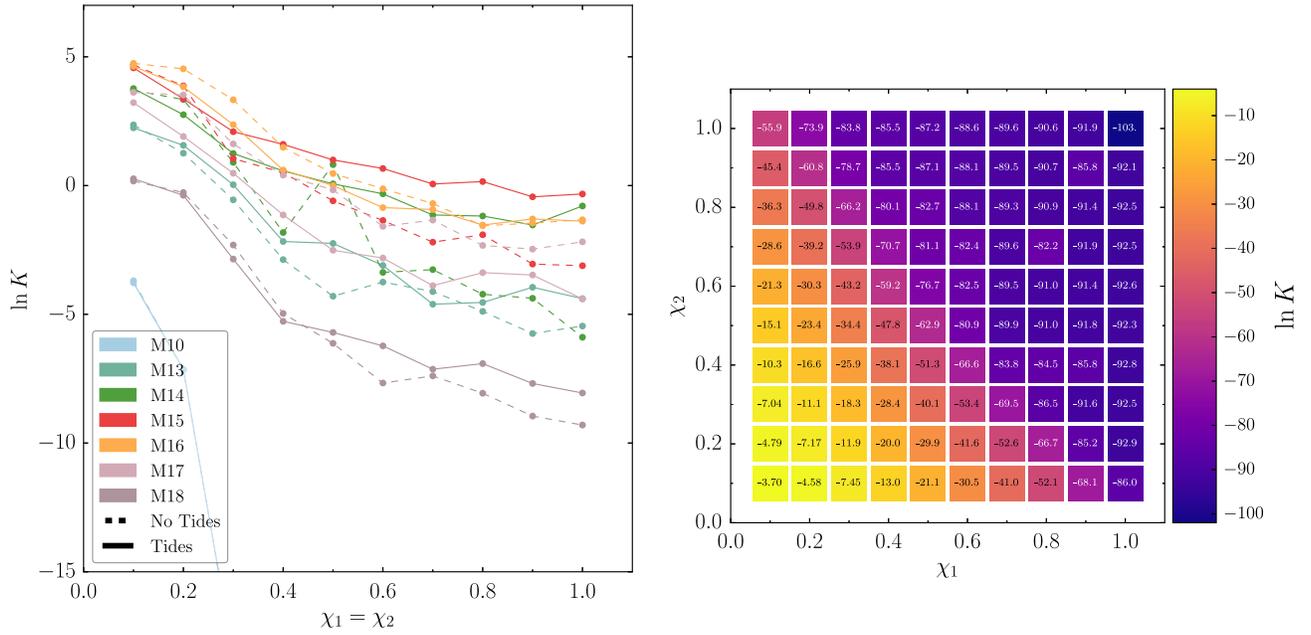


FIG. 4. *Standard small-kick scenario (M10) requires small natal BH spin:* Left panel: A plot of the Bayes factor K derived by comparing GW151226, GW150914, and LVT151012 to the M10 (blue) formation scenario, versus the magnitude of assumed BH natal spin $\chi_1 = \chi_2$. All other models are shown for comparison. Colors and numbers indicate the Bayes factor; dark colors denote particularly unlikely configurations. Right panel: As before (i.e., for M10), but in two dimensions, allowing the BH natal spins for the primary and secondary BH to be independently selected (but fixed); color indicates the Bayes factor. As this scenario predicts strictly aligned spins for the heaviest BH-BH binaries, only small BH natal spins are consistent with LIGO’s constraints on the (aligned) BH spin parameter χ_{eff} in O1 (and GW170104); see Abbott, B. P. *et al.* (LIGO and Virgo Scientific Collaboration) [5,49,54].

B. BH natal kicks and misalignment (O1)

In the absence of BH natal kicks, the preponderance of observed BH-BH binaries consistent with $\chi_{\text{eff}} \approx 0$ (e.g., GW150914 and GW170104) provided *conditional* evidence in favor of small BH natal spins. But even small BH natal kicks can frequently produce significant spin-orbit misalignment. Once one incorporates models that permit nonzero BH natal kicks, then even binary BHs with large BH natal spins could be easily reconciled with every one of LIGO’s observations. Figures 4 and 5 provide a quantitative illustration of just how much more easily models with even modest BH natal kicks can explain the data, for a wide range of BH spins. When natal kicks greater than 25 km/s are included, the BH natal spin is nearly unconstrained. As is particularly apparent in Fig. 5, some natal BH spin is required to reproduce the nonzero spin seen in GW151226.

Larger kicks produce frequent, large spin-orbit misalignments and therefore greater consistency with the properties of all of LIGO’s observed binary BHs. Spin-orbit misalignment is consistent with the spin distribution of GW151226, and helpful to explain the distribution of χ_{eff} for LIGO’s other observations. However, larger kicks also disrupt more binaries, substantially decreasing the overall event rate (see Fig. 1). Figure 5 illustrates the tradeoff between spin-orbit misalignment and event rate.

C. Tides and realignment (O1)

All other things being equal, our “no tides” scenarios most frequently produce significant spin-orbit misalignment. As a result, even for large BH natal spins, these models have a greater ability to explain LIGO’s observations, which are largely consistent with $\chi_{\text{eff}} \approx 0$. The “tides” scenario produces smaller misalignments for the second-born BH. Figure 5 quantitatively illustrates how the “no tides” scenario marginally fits the data better. In order to reproduce the inferred distribution of spin-orbit misalignments (in GW151226) and low χ_{eff} (for all events so far), the “tides” models likely have (a) larger BH natal kicks ≈ 200 km/s and (b) low BH natal spins $\chi_{1,2} \lesssim 0.2$. Conversely, when “no tides” act to realign the second BH spin, *small* natal kicks ≈ 50 km/s are favored. Figure 5 illustrates the two distinct conclusions about BH natal kick strength drawn, depending on whether stellar tidal realignment is efficient or inefficient. Based on this figure (and hence on the assumption of fixed natal spins), we estimate that massive BHs should receive a natal kick of ~ 50 km/s if no processes act to realign stellar spins. Significantly larger natal kicks, with one-dimensional velocity dispersion ≈ 200 km/s, will be required if stellar spins efficiently realign prior to the second BH’s birth.

Tides also introduce an asymmetry between the spin-orbit misalignment of the first-born (generally more

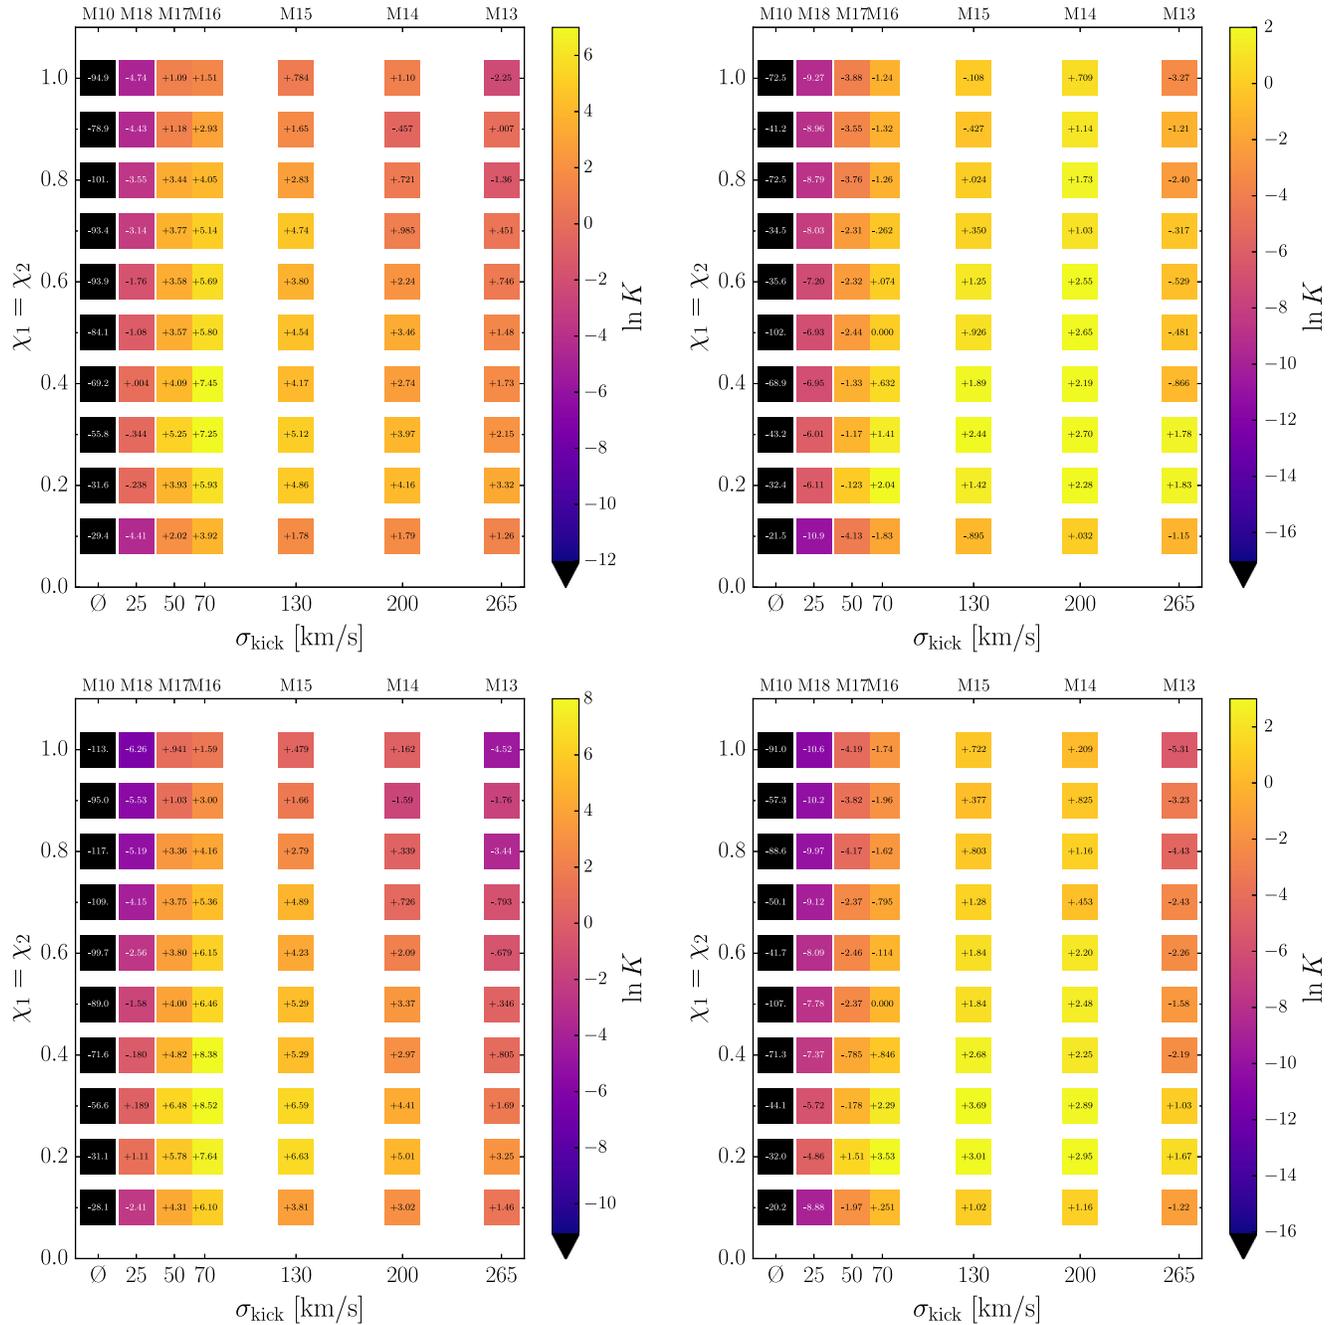


FIG. 5. *Bayes factor versus spin and kicks, with and without tides:* A plot of the Bayes factor versus BH natal spin ($\chi = \chi_1 = \chi_2$) and natal kick (σ_{kick}). The left and right panels correspond to “no tides” and “tides,” respectively. The top two panels use only the O1 events; the bottom two panels account for the events and network sensitivity updates reported in the GW170104 discovery paper. In each panel, the zero point of the Bayes factor is normalized to the BH-BH formation scenario with $\chi_1 = \chi_2 = 0.5$ and $\sigma = 70$ km/s and “tides.”

massive) and second-born (generally less massive) BH [30]. As a result, when we consider general prescriptions for BH natal spins $\chi_1 \neq \chi_2$, we find that scenarios without tides produce largely symmetric constraints on $\chi_{1,2}$. When we assume tidal alignment, we can draw stronger constraints about the second-born spin rather than the first. Paradoxically, large natal spin on the *first* born BH is consistent with observations. The second born BH cannot

significantly misalign its spin through a natal kick; therefore, for comparable mass binaries like GW150914, we know that the second-born BH spin must be small, if it is strongly aligned. More broadly, since observations rule out large χ_{eff} , binary formation scenarios with tides and with $\chi_1 > \chi_2$ fit the data substantially better than scenarios with tides and $\chi_2 > \chi_1$. Because tides act to realign the second spin, only when $\chi_2 \leq \chi_1$ will we have a chance at producing

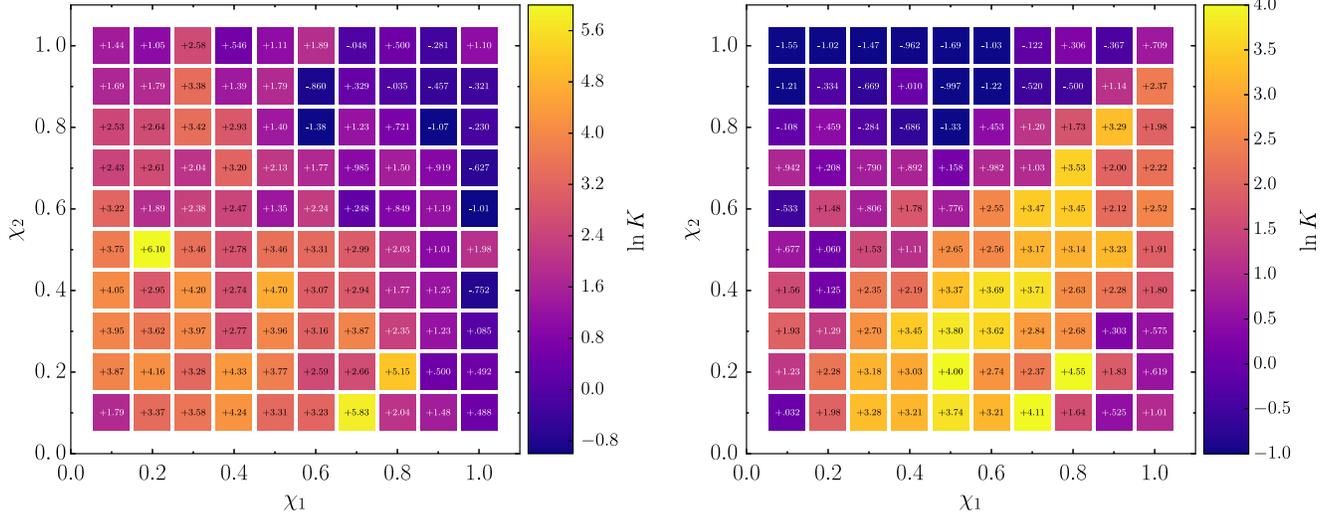


FIG. 6. *Bayes factor versus spin, with and without tides (O1)*: For the M14 model ($\sigma = 200$ km/s), a plot of the Bayes factor versus $\chi_{1,2}$. Colors and numbers indicate the Bayes factor; dark colors denote particularly unlikely configurations. The left panel assumes no spin realignment (“no tides”); the right panel assumes the second-born BH’s progenitor had its spin aligned with its orbit just prior to birth (“tides”). Spin-orbit realignment and the high orbital velocity just prior to the second SN ensures the second spin is at best weakly misaligned; therefore, χ_2 would need to be small for these models to be consistent with LIGO’s observations to date.

small $|\chi_{\text{eff}}|$, as LIGO’s O1 observations suggest. Figure 6 illustrates this asymmetry.

The illustrative results described in this section follow from our strong prior assumptions: fixed BH natal spins. As described below, if we instead adopt some broad distribution of BH natal spins, the substantially greater freedom to reproduce LIGO’s observations reduces our ability to draw other distinctions, in direct proportion to the complexity of the prior hypotheses explored. We describe results with more generic spin distributions below.

D. BH natal spins, given misalignment (O1)

So far, to emphasize the information content in the data, we have adopted the simplifying assumption that each pair of BHs has the same natal spins χ_1, χ_2 . This extremely strong family of assumptions allows us to leverage all four observations, producing large changes in Bayes factor as we change our assumptions about (all) BH natal spins. Conversely, if the BH natal spins are nondeterministic, drawn from a distribution with support for any spin between 0 and 1, then manifestly only four observations cannot hope to constrain the BH natal spin distribution, even were LIGO’s measurements to be perfectly informative about each BH’s properties. Astrophysically-motivated or data-driven prior assumptions must be adopted in order to draw stronger conclusions about BH spins (cf. [81]).

As a concrete example, we consider the simple two-bin BH natal spin model described in Eq. (3), with probability λ_A that any BH has natal spin $\chi_i \leq 0.6$ and probability $1 - \lambda_A$ that any BH natal spin is larger than 0.6. The choice of 0.6 is motivated by our previous results in Fig. 5, as well as by the empirical χ_{eff} distribution shown in Fig. 3. Using

the techniques described in Appendix B, we can evaluate the posterior probability for λ_A given LIGO’s O1 observations, within the context of each of our binary evolution models. Figure 7 shows the result: LIGO’s observations weakly favor low BH natal spins. For models like M10 and M13, with minimal BH natal kicks and hence spin-orbit misalignment, low BH natal spin is necessary to reconcile models with the fact that LIGO hasn’t seen BH-BH binaries with large, aligned spins and thus large χ_{eff} . Conversely, LIGO’s observations will modestly less strongly disfavor models that frequently predict large BH natal spins (e.g., $\lambda \lesssim 0.6$).

As we increase the complexity of our prior assumptions, our ability to draw conclusions from only four observations rapidly decreases. For example, we can construct the posterior distribution for a generic BH natal spin distribution (i.e., our mixture coefficients λ_α for each spin combination can take on any value whatsoever). The mean spin distribution can be evaluated using closed-form expressions provided in Appendix B. In this extreme case, the posterior distribution closely resembles the prior for almost all models, except M10.

To facilitate exploration of alternative assumptions about natal spins and kicks, we have made publicly available all of the marginalized likelihoods evaluated in this work, as Supplemental Material [56].

E. Information provided by GW170104

The observation of GW170104 enables us to modestly sharpen all of the conclusions drawn above, due to the reported limits on χ_{eff} : between -0.42 and 0.09 [5]. Of course, the reported limits for all events must always be

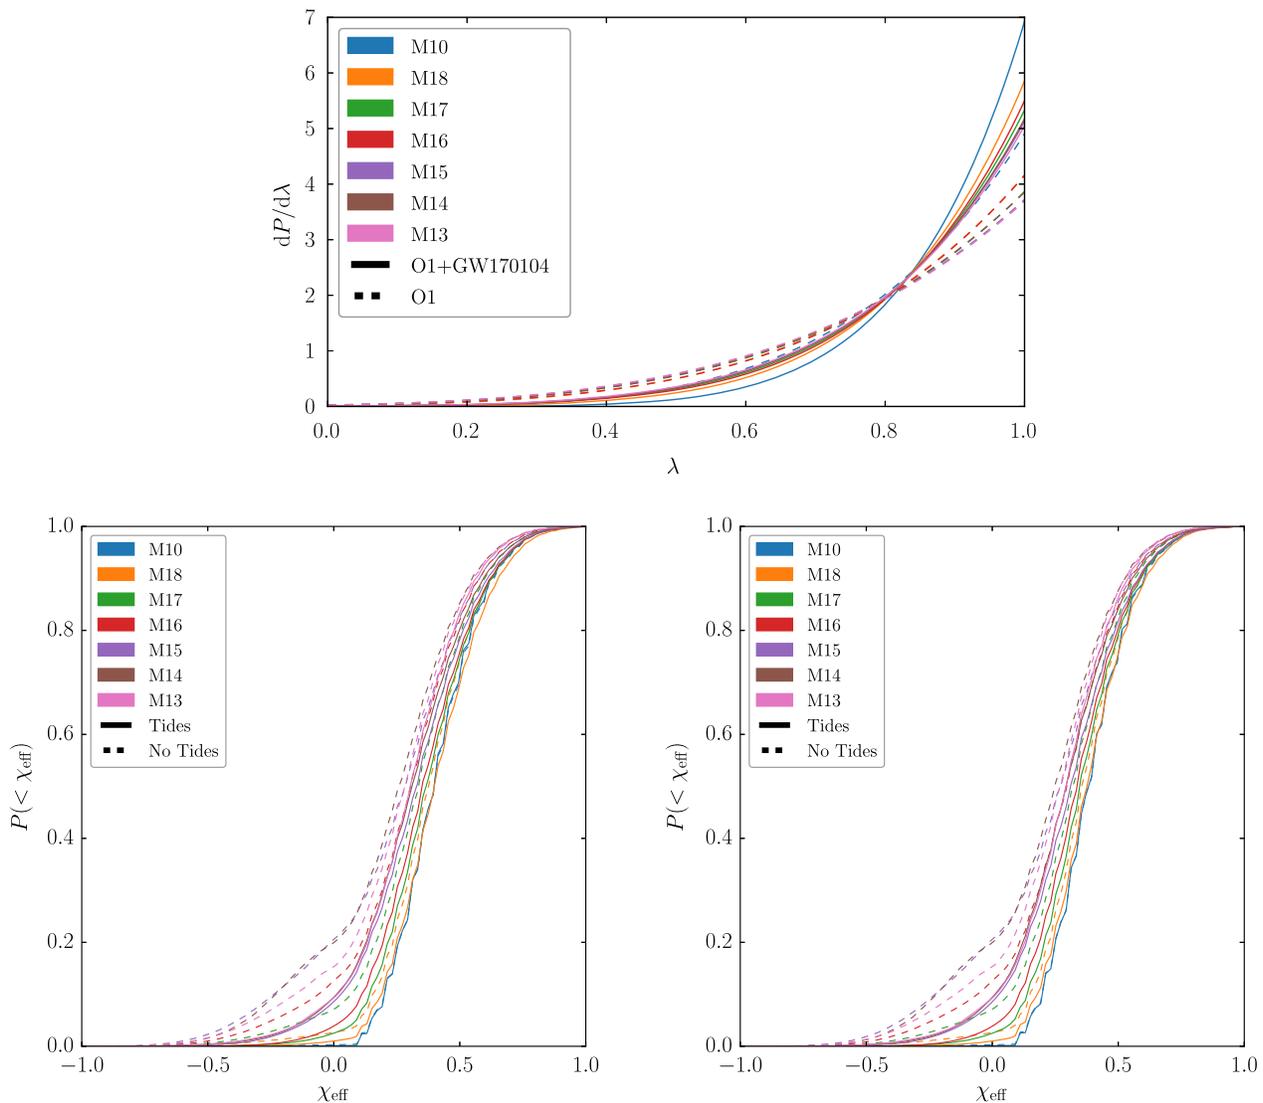


FIG. 7. *High or low natal spin?* Top panel: Posterior distribution on λ_A , the fraction of BHs with natal spins ≤ 0.6 [Eq. (3)], based on O1 (dotted) or on O1 with GW170104 (solid), compared with our binary evolution models (colors), assuming “no tides.” Unlike Fig. 5, which illustrates Bayes factors calculated assuming fixed BH natal spins, this calculation assumes each BH natal spin is drawn at random from a mass- and formation-scenario-independent distribution that is piecewise constant above and below $\chi = 0.6$. With only four observations, LIGO’s observations consistently but weakly favor low BH natal spins. Left panel: Posterior distribution for χ_{eff} implied by the distribution of λ_A shown in the top panel [i.e., by comparing our models to LIGO’s O1 observations, under the assumptions made in Eq. (3)]. Right panel: As in the left panel, but including GW170104. Adding this event does not appreciably or qualitatively change our conclusions relative to O1.

taken in context, as they are inferred using very specific assumptions—*a priori* uniform spin magnitudes, isotropically oriented. Necessarily, inference performed in the context of any astrophysical model for natal BH spins and kicks will draw different conclusions about the allowed range, since the choice of prior influences the posterior spin distribution (see, e.g., [81,82]). Even taking these limits at face value, however, this one observation can easily be explained using some combination of two effects: a significant probability for small natal BH spins, or some BH natal kicks. First and most self evidently, if all BHs have similar natal spins, then binary evolution models that

assume alignment at birth; do not include processes that can misalign heavy BH spins, like M10; and which adopt a common natal BH spin for all BHs are difficult to reconcile with LIGO’s observations. On the one hand, GW170104 would require extremely small natal spins in this scenario; on the other, GW151226 requires nonzero spin. Of course, a probabilistic (mixture) model allowing for a wide range of mass-independent BH natal spins can easily reproduce LIGO’s observations, even without permitting any alignment; see also [42], which adopts a deterministic model that also matches these two events. Second, binary evolution models with significant BH natal kicks can also explain

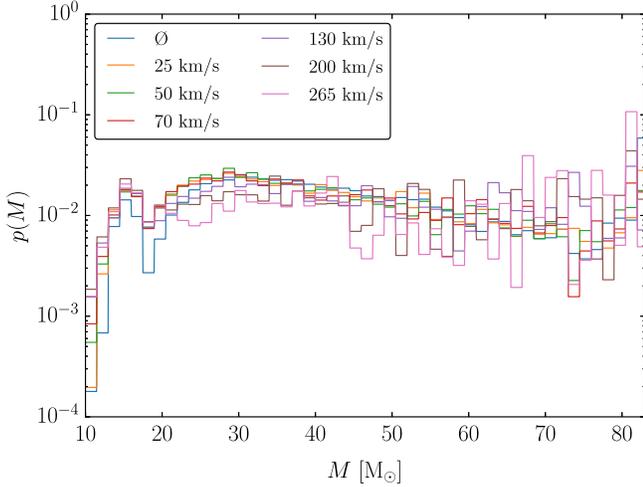


FIG. 8. Detection-weighted total mass distributions of our models, labeled by their σ_{kick} values, without accounting for LIGO measurement error. The overall mass distributions are very similar, particularly for low kicks.

LIGO’s observations. As seen in the bottom left panel of Fig. 5, large BH natal spins are harder to reconcile with LIGO’s observations, if we assume BH spin alignments are only influenced by isotropic BH natal kicks. This conclusion follows from the modest χ_{eff} seen so far for all events. Conversely, if we assume efficient alignment of the second-born BH, then the observed distribution of χ_{eff} (and θ_1 , mostly for GW151226) suggest large BH natal kicks, as illustrated by the bottom right panel of Fig. 5.

F. Information provided by the mass distribution

The underlying mass distributions predicted by our formation models do depend on our assumptions about BH natal kicks, as shown concretely in Fig. 8. These modest differences accumulate as BH natal kicks increasingly disrupt and deplete all BH-BH binaries. To quantify the similarity between our distributions, Table I reports an information-theory-based metric (the KL divergence) that attempts to quantify the information rate or “channel capacity” by which the universe communicates information about the mass distribution to us. If $p(x)$, $q(x)$ are two probability distributions over a parameter x , then in general the KL divergence has the form

$$D_{KL}(p|q) = \int dx p(x) \ln[p(x)/q(x)]. \quad (4)$$

Except for the strongest BH natal kicks, we find our mass distributions are nearly identical. Even with perfect mass measurement accuracy, we would need $\mathcal{O}(1/D_{KL})$ fair draws from our distribution to confidently distinguish between them. As demonstrated by previous studies [79,83], LIGO will be relatively inefficient at discriminating between the different detected mass distributions. LIGO

is most sensitive to the heaviest BHs, which dominate the astrophysically observed population, but has extremely large measurement uncertainty in this regime. Thus, accounting for selection bias and smoothing using estimated measurement error, the mass distributions considered here look fairly similar [79]. For constraints on BH natal kicks, the information provided by the mass distribution is far less informative than the insights implied by constraints on χ_{eff} and $\theta_{1,2}$.

As a measure of the information LIGO can extract per event about the mass distribution from each detection, we enumerate how many different BH-BH binaries LIGO can distinguish, which are consistent with the expected stellar-mass BH-BH population (i.e., motivated by LIGO’s reported observations to date, limiting to $m_2/m_1 > 0.5$, $m_1 + m_2 < 75M_{\odot}$, $m_2 > 3M_{\odot}$, and $m_1 < 40M_{\odot}$). Counting up the distinct waveforms used by gravitational wave searches in O2 [84], including spin, there are only **236** templates with chirp masses above LVT151012 (i.e., $\mathcal{M}_c > 15M_{\odot}$), and only $\approx 1,200$ with chirp masses above GW151226 (i.e., $\mathcal{M}_c > 8.88M_{\odot}$). This estimate is highly optimistic, because it neglects distance and hence redshift uncertainty, which decreases our ability to resolve the smallest masses (i.e., the uncertainty in chirp mass for GW151226), and it also uses both mass and spin information. Judging from the reported mass distributions alone (e.g., the top left panel of Fig. 4 in [4]), LIGO may efficiently isolate BHs to only a few tens of distinct mass bins, de facto limiting the resolution of any mass distribution which can be nonparametrically resolved with small-number statistics; see, e.g., the discussion in [83].

V. PREDICTIONS AND PROJECTIONS

Using the Bayes factors derived above for our binary evolution models and BH natal spin assumptions (collectively indexed by Λ), we can make predictions about future LIGO observations, characterized by a probability distribution $p_{\text{future}}(x) = \sum_{\Lambda} p(x|\Lambda)p(\Lambda|d)$ for a candidate future binary with parameters x . We can then account for LIGO’s mass-dependent sensitivity to generate the relative probability of observing binaries with those parameters. In the context of the infrastructure described above, we evaluate this detection-weighted posterior probability using a mixture of synthetic universes, with relative probabilities $p(\Lambda|d)$ and relative weight r_i of detecting an individual binary drawn from it.

Using our fiducial assumptions about BH spin realignment (“no tides”), our posterior probabilities point to nonzero BH natal kicks, with BH natal spins that can neither be too large nor too small (Figs. 5 and 7). In turn, because each of our individual formation scenarios Λ preferentially forms binaries with $\chi_{\text{eff}} > 0$ [55], with a strong preference for the largest χ_{eff} allowed, we predict future LIGO observations will frequently include binaries with the largest χ_{eff} allowed by the BH natal spin distribution. These measurements will

self-evidently allow us to constrain the natal spin distribution (e.g., the maximum natal BH spin). For example, if future observations continue to prefer small χ_{eff} , then the data would increasingly require smaller and smaller natal BH spins, within the context of our models. For example, this future scenario would let us rule out models with large kicks and large spins, as then LIGO should nonetheless frequently detect binaries with large χ_{eff} .

As previously noted, with only four GW observations, the data does not strongly favor any spin magnitude distribution. Strongly modeled approaches which assume specific relationships between the relative prior probability of different natal spins can draw sharper constraints, as in [54]. If we allow the spin distribution to take on any form [53,85], many observations would be required to draw conclusions about the spin distribution. Conversely, as described previously and illustrated by Fig. 7, if we adopt a weak (piecewise-constant) model, we can draw some weak conclusions about the BH natal spin distributions that are implied by our binary evolution calculations and LIGO's observations.

Neither the expected number of events nor their mass distribution merits extensive discussion. The large Poisson error implied by only four observations leads to a wide range of probable event rates, previously shown to be consistent with all the binary evolution models presented here [26,59]. Conversely, due to the limited size of our model space—the discrete model set and single model parameter (BH natal kicks) explored—these posterior distributions by no means fully encompass all of our prior uncertainty in binary evolution and all we can learn by comparing GW observations with the data. While our calculations illuminate how GW measurements will inform our understanding of BH formation, our calculations are not comprehensive enough to provide authoritative constraints except for the most robust features.

Finally, all of our calculations and projections have been performed in the context of one family of formation scenarios—isolated binary evolution. Our calculations within this framework do not allow for one possible variant of this channel: homogeneous chemical evolution, where close binaries become tidally locked and rapidly rotating, leading to a distinctively different evolutionary trajectory that produces massive BH binaries while circumventing the common envelope phase [6,7]. Globular clusters could also produce a population of merging compact binaries [8], with random spin-orbit misalignments [86]. Several previous studies have described or demonstrated how to identify whether either model contributes to the detected population, and by how much, using constraints on merging BH-BH spins [2,14,15,52,53,85].

VI. CONCLUSIONS

By comparing binary evolution models with different assumptions about BH natal kicks to LIGO observations of

binary BHs, we estimate that heavy BHs should receive a natal kick of order 50 km/s if no processes act to realign stellar spins. Significantly larger natal kicks, with one-dimensional velocity dispersion ≈ 200 km/s, will be required if stellar spins efficiently realign prior to the second BH's birth. These estimates are consistent with observations of galactic X-ray binary misalignment [32–35] and recoil velocity [35–41]. Our estimate is driven by two simple factors. The natal kick dispersion σ is bounded from above because large kicks disrupt too many binaries (reducing the merger rate below the observed value). Conversely, the natal kick distribution is bounded from below because modest kicks are needed to produce a range of spin-orbit misalignments. A distribution of misalignments increases our models' compatibility with LIGO's observations, if all BHs are likely to have natal spins.

Closely related work by Belczynski *et al.* [42] uses similar evolutionary models but with a fixed physically-motivated BH natal spin model that depends on BH mass. They predict a distribution of χ_{eff} with substantial support at large values, in increasing tension with observations reported to date. They conclude that more efficient angular momentum transport needs to be adopted in evolutionary calculations to revise their BH natal spin model and to match LIGO/Virgo observations. In this work, by contrast, we explore a wide range of possible spin distributions. We consistently find that distributions which favor low BH natal spins can more easily reproduce current observations.

Given limited statistics, we have for simplicity (and modulo M10) assumed all binary BHs receive natal kicks and spins drawn from the same formation-channel-independent distributions. This strong assumption about BH natal spins allows us to draw sharp inferences about BH natal spins and kicks by combining complementary information provided by GW151226 (i.e., nonzero spins required, with a suggestion of misalignment) and the remaining LIGO observations (i.e., strong limits on χ_{eff}). Future observations will allow us to directly test more complicated models not explored here, where the natal spin and kick distribution depends on the binary BH mass as in Belczynski *et al.* [42]. Necessarily, if BH natal spins are small for massive BHs and large for small BHs, as proposed in Belczynski *et al.* [42], then measurements of low-mass BH binaries like GW151226 will provide our primary channel into constraining BH natal spins and kicks. At present, however, inferences about BH natal spins and spin-orbit misalignment are strongly model or equivalently prior driven, with sharp conclusions only possible with strong assumptions. We strongly recommend results about future BH-BH observations be reported or interpreted using multiple and astrophysically motivated priors, to minimize confusion about their astrophysical implications (e.g., drawn from the distribution of χ_{eff}).

For simplicity, we have also only adjusted one assumption (BH natal kicks) in our fiducial model for how compact

binaries form. A few other pieces of unknown and currently parametrized physics, notably the physics of common envelope evolution, should play a substantial role in how compact binaries form and, potentially, on BH spin misalignment. Other assumptions have much smaller impact on the event rate and particularly on BH spin misalignment. Adding additional sources of uncertainty will generally diminish the sharpness of our conclusions. For example, the net event rate depends on the assumed initial mass function as well as the star formation history and metallicity distribution throughout the universe; once all systematic uncertainties in these inputs are included, the relationship between our models and the expected number of events is likely to include significant systematic as well as statistical uncertainty. Thus, after marginalizing over all sources of uncertainty, the event rate may not be as strongly discriminating between formation scenarios. By employing several independent observables (rate, masses, spins and misalignments), each providing weak constraints about BH natal kicks, we protect our conclusions against systematic errors in the event rate. Further investigations are needed to more fully assess sources of systematic error and enable more precise constraints.

Due to the limited size of our model space—the discrete model set and single model parameter (BH natal kicks) explored—these posterior distributions by no means fully encompass all of our prior uncertainty in binary evolution and all we can learn by comparing GW observations with the data. As in previous early work [87–90], a fair comparison must broadly explore many more elements of uncertain physics in binary evolution, like mass transfer and stellar winds. Nonetheless, this nontrivial example of astrophysical inference shows how we can learn about astrophysical models via simultaneously comparing GW measurements of several parameters of several detected binary BHs to predictions of any model(s). While we have applied our statistical techniques to isolated binary evolution, these tools can be applied to generic formation scenarios, including homogeneous chemical evolution; dynamical formation in globular clusters or AGN disks; or even primordial binary BHs.

Forthcoming high-precision astrometry and radial velocity from GAIA will enable higher-precision constraints on existing X-ray binary proper motions and distances [91,92], as well as increasing the sample size of available BH binaries. These forthcoming improved constraints on BH binary velocities will provide a complementary avenue to constrain BH natal kicks using binaries in our own galaxy.

ACKNOWLEDGMENTS

We thank Christopher Berry, Simon Stevenson, and Will Farr for helpful comments on the draft. D. W. is supported by the Rochester Institute of Technology (RIT) through the Frontiers in Gravitational Wave Astrophysics (FGWA) Signature Interdisciplinary Research Areas (SIRA) initiative and College of Science (COS). R. O. is supported by NSF

Grants No. AST-1412449, PHY-1505629, and PHY-1607520. D. G. is supported by NASA through Einstein Postdoctoral Fellowship Grant No. PF6-170152 awarded by the Chandra X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for NASA under Contract NAS8-03060. E. B. is supported by NSF Grants No. PHY-1607130 and AST-1716715, and by FCT contract IF/00797/2014/CP1214/CT0012 under the IF2014 Programme. M. K. is supported by the Alfred P. Sloan Foundation Grant No. FG-2015-65299 and NSF Grant No. PHY-1607031. R. O. and E. B. acknowledge the hospitality of the Aspen Center for Physics, supported by NSF PHY-1607611, where this work was completed. K. B. acknowledges support from the Polish National Science Center (NCN) grant: Sonata Bis 2 (DEC-2012/07/E/ST9/01360). D. E. H. was partially supported by NSF CAREER Grant No. PHY-1151836 and NSF Grant No. PHY-1708081. He was also supported by the Kavli Institute for Cosmological Physics at the University of Chicago through NSF Grant No. PHY-1125897 and an endowment from the Kavli Foundation. Computations were performed on the Caltech computer cluster “Wheeler,” supported by the Sherman Fairchild Foundation and Caltech. Partial support is acknowledged by NSF CAREER Award PHY-1151197. The authors thank to the LIGO Scientific Collaboration for access to the data and gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, and the Max-Planck-Society (MPS) for support of the construction of Advanced LIGO. Additional support for Advanced LIGO was provided by the Australian Research Council.

APPENDIX A: APPROXIMATING PARAMETER DISTRIBUTIONS FROM FINITE SAMPLES

Our population synthesis techniques allow us to generate an arbitrarily high number of distinct binary evolutions from each formation scenario, henceforth indexed by Λ . Instead of generating individual binary evolution histories, we weigh each one by an occurrence rate, allowing it to represent multiple binaries. For our calculations, however, we instead require the relative probability of different binaries, not just samples from the distribution. We estimate this distribution from the large but finite sample of binaries available in each synthetic universe. We do not simply use an occurrence rate-weighted histogram of all the samples. Histograms work reliably for any single parameter (e.g., $p(m_1|\Lambda)$), where many samples are available per potential histogram bin, but for high-dimensional joint distributions (e.g., $p(m_1, m_2, \theta_1, \theta_2, \chi_{\text{eff}}|\Lambda)$), many histogram bins will be empty simply due to the curse of dimensionality.

In all our calculations, we instead approximate the density as a mixture of Gaussians, labeled $k = 1, 2, \dots, K$, with means and covariances (μ_k, Σ_k) to be estimated, along

with weighting coefficients w_k , which must sum to unity. The density can therefore be written as

$$p(\mathbf{x}) \approx \sum_{k=1}^K w_k \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (\text{A1})$$

where $\mathcal{N}(\cdot)$ represents the (multivariate) Gaussian distribution. We select the number of Gaussians K by using the Bayesian information criterion.

To estimate the means and covariances of our mixture of Gaussians, we used the expectation maximization algorithm [93]; see, e.g., [94] for a pedagogical introduction. Specifically, we used a small modification to an implementation in `scikit-learn` [95], to allow for weighted samples in the update equation (e.g., adding weights to Eq. (11.27) in [94]).

Ideally we would simply approximate each formation scenario Λ 's intrinsic predictions $p(\mathbf{x} | \Lambda)$ with a mixture of Gaussians, using the merger rate for each sample binary as its weighting factor. However, all astrophysical indications suggest that more massive progenitors form more rarely, implying this procedure would result in a distribution that is strongly skewed in favor of the much more intrinsically frequent low mass systems; our fitting algorithm might end up effectively neglecting the samples with small weights. This would risk losing information about the most observationally pertinent samples, which due to LIGO's mass-dependent sensitivity are concentrated at the highest observationally accessible masses. Alternatively, for every choice of detection network, we can approximate each formation scenario's predictions *for that network*. If $TV(\mathbf{x})$ is the average sensitive 4-volume for the network, according to this procedure we approximate $V(\mathbf{x})p(\mathbf{x})$ by a Gaussian mixture, then divide by $V(\mathbf{x})$ to estimate $p(\mathbf{x})$. To minimize duplication of effort involved in regenerating our approximation for each detector network, we instead adopt a *fiducial* (approximate) network sensitivity model $V_{\text{ref}}(\mathbf{x})$ for the purposes of density estimation. We adopt the simplest (albeit ad-hoc) network sensitivity model: the functional form for $V(\mathbf{x})$ that arises by using a single detector network and ignoring cosmology (i.e., $EV \propto \mathcal{M}_c^{15/6}$) [76]. The overall, nominally network- and run-dependent normalization constant in this ad-hoc model V_{ref} scales out of all final results.

APPENDIX B: HIERARCHICAL COMPARISONS OF OBSERVATIONS WITH DATA

As described in Sec. III B, the population of binary mergers accessible to our light cone can be described as an inhomogeneous Poisson process, characterized by a probability density $e^{-\mu} \prod_k \mathcal{R}p(x_k)$ where $x_k = x_1 \dots x_N$ are the distinct binaries in our observationally accessible parameter volume \mathcal{V} . In this expression, the expected number of events and parameter distribution are related by $\mu = \int dx \sqrt{g} \mathcal{R}p(x)$; the multidimensional integral $\int dx \sqrt{g}$ is

shorthand for a suitable integration over a manifold with metric; and the probability density $p(x)$ is expressed relative to the fiducial (metric) volume element, but normalized on a larger volume than \mathcal{V} . Accounting for data selection [78], the likelihood of all of our observations is therefore given by Eq. (2).

To insure we fully capture the effects of precessing spins, we work not with the full likelihood—a difficult function to approximate in 8 dimensions—but instead with a fiducial posterior distribution $p_{\text{post}} = Z^{-1} p(d_k | x) p_{\text{ref}}(x_k)$, as would be provided by a Bayesian calculation using a reference prior $p_{\text{ref}}(x_k)$. Rewriting the integrals $\int dx_k p(d_k | x) p(x_k | \Lambda)$ appearing in Eq. (2) using the reference prior we find integrals appearing in this expression can be calculated by Monte Carlo, using some sampling distribution $p_{s,k}(x_k)$ for each event (see, e.g., [96]):

$$\begin{aligned} & \int dx_k p(d_k | x_k) p(x_k | \Lambda) \\ &= \frac{1}{N_k} \sum_s \frac{[p(d_k | x_k) p_{\text{ref}}(x_{k,s})] p(x_{k,s} | \Lambda)}{p_{s,k}(x_{k,s}) p_{\text{ref}}(x_{k,s})}, \end{aligned} \quad (\text{B1})$$

where $s = 1, \dots, N_k$ indexes the Monte Carlo samples used. One way to evaluate this integral is to adopt a sampling distribution $p_{s,k}$ equal to the posterior distribution evaluated using the reference prior, and thus proportional to $p(d_k | x_k) p_{\text{ref}}(x_k | \Lambda)$. If for this event k we have samples $x_{k,s}$ from the posterior distribution—for example, as provided by a Bayesian Markov chain Monte Carlo code—the integrals appearing in Eq. (2) can be estimated by

$$\int dx_k p(d_k | x_k) p(x_k | \Lambda) \approx \frac{Z}{N_k} \sum_s \frac{p(x_{k,s} | \Lambda)}{p_{\text{ref}}(x_{k,s})}. \quad (\text{B2})$$

We use this expression to evaluate the necessary marginal likelihoods, for any proposed observed population $p(x | \Lambda)$.

In the expression above, we need only consider *some* of the degrees of freedom in the problem. Notably, the probability distributions for extrinsic parameters like the source orientation, sky location, and distance will always be in common between our models and our reference prior. So will any Jacobians associated with changes of coordinate. Moreover, these assumptions are independent of one another and of the intrinsic parameter distributions. Therefore, the ratio of probability densities $p(x | \Lambda) / p_{\text{ref}}(x)$ usually has product form, canceling term by term. We therefore truncate the ratio to only account for *some* of the degrees of freedom.

To verify and better understand our results, we can also *approximate* the likelihood function, using suitable summary statistics. As an example, B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaboration) [49] reproduce parameter estimates of GW150914 using a Gaussian approximation to the likelihood and the assumption of perfect spin-orbit alignment. Using this approximation, and a similar approximation for

GW151226, we can alternatively approximate each integral appearing in the likelihood by using the (weighted) binary evolution samples $x_{k,A}$ and their weights w_A :

$$\int dx_k p(d_k|x_k) p(x_k|\Lambda) \approx \frac{\sum_A w_A \hat{p}(d_k|x_{k,A})}{\sum_A w_A} \quad (\text{B3})$$

where \hat{p} refers to our approximate likelihood for the k th event. Even though these likelihood approximations neglect degrees of freedom associated with spin precession, we can reproduce the observed mass and χ_{eff} distributions reported in B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaboration) [49]. We used this approximate likelihood approach to validate and test our procedure. We also use this approach to incorporate information about GW170104, which was not available at the same level of detail as the other events.

As an example, we describe how to evaluate this integral in the case where $p(x_k|\lambda)$ is a mixture model $p(x|\lambda) = \sum_\alpha \lambda_\alpha p_\alpha(x)$, for λ an array of parameters. In this case, all the integrals can be carried out via

$$\begin{aligned} & \prod_k \int dx_k p(d_k|x_k) p(x_k|\lambda) \\ &= \prod_k \left[\sum_\alpha \lambda_\alpha \int dx_k p(d_k|x_k) p_\alpha(x_k) \right] = \prod_k \sum_\alpha \lambda_\alpha c_{\alpha,k} \end{aligned} \quad (\text{B4})$$

where $c_{\alpha,k}$ are integrals we can compute once and for all for each event, using for example the posterior samples from some fiducial analysis. As a result, the observation-dependent factor in likelihood for a mixture model always reduces to a homogeneous N th-degree polynomial in the mixture parameters λ_α . Bayes theorem can be applied to λ to infer the distribution over mixture parameters. Depending on the mixture used, this calculation could incorporate a physically-motivated prior on λ .

We use a mixture model approach to hierarchically constrain the spin magnitude distribution implied by our data. In our approach, we first consider models where both spin magnitudes are fully constrained. In the notation of the mixture model discussion above, we adopt some specific prior $p_\alpha(\chi_1, \chi_2|\sigma) = \delta(\chi_1 - x_\alpha) \delta(\chi_2 - y_\alpha)$ where x_α, y_α are the spin λ_α . A mixture model allowing generic λ and thus including all such components allows both component spins to take arbitrary (discrete) values. [We could similarly extend our mixture model to include kicks.] The posterior distribution over all possible spin distributions $p(\lambda|d) = p(d|\lambda)p(\lambda)/p(d)$ follows from Bayes' theorem and the concrete likelihood given in Eq. (B4). In practice, however, we don't generally compute or report the full posterior distribution, as it contains far more information than we need (e.g., the extent of the ensemble of possible spin distributions that fit the data). Instead, we compute the *expected spin distribution*

$$p_{\text{post}}(x) = \sum_\alpha \langle \lambda_\alpha \rangle p_\alpha(x) \quad (\text{B5})$$

and the *variance in each* λ_α . For the modest number of mixture components of interest here (≈ 100 possible choices of both spin magnitudes) and the modest degree of the polynomial ($\approx 4-5$), all necessary averages can be computed by direct symbolic quadrature of a polynomial in λ_α . The integral can be expressed as a sum of terms of homogeneous degree in λ , and integrals of each of these terms can be carried out via the following general formula:

$$\begin{aligned} & n! \int_{\sum_i x_i \leq 1} dx_1 \dots dx_n x_{i_1}^{\alpha_1} \dots x_{i_Z}^{\alpha_Z} \\ &= \frac{n!}{(n-Z)!} \prod_{k=1}^Z B\left(\alpha_k + 1, n + 1 - k + \sum_{q>k} \alpha_q\right) \end{aligned} \quad (\text{B6})$$

where the integral is over the region $x_i \geq 0$ and $\sum_i x_i \leq 1$. We can also find the maximum likelihood estimate of λ_α , for example by using the expectation-maximization algorithm [93]. In this work, however, we have many more basis models $\alpha = 1, 2, \dots$ used in our (spin) mixture than observations. Normally, we would reduce the effective dimension, for example by adopting prior assumptions in how the mixture coefficients can change as a function of spins χ_1, χ_2 . To minimize additional formal overhead, we instead simply treat the spins hierarchically in blocks [Eq. (3)], considering lower-dimensional models where (for example) λ_A denotes the *a priori* probability for $\chi_i \leq 0.6$ and $1 - \lambda_A$ denotes the *a priori* probability for $\chi_i > 0.6$, so for example the prior probability for $(\chi_1, \chi_2) = 0.1$ is $\lambda_A^2/36$. In this four-block and one-parameter model, we can compute the average value of λ_A in terms of the net weights associated with each block: $C_{AA,k} = \sum_{\chi_1, \chi_2 \in A} c_{\alpha,k}$, $C_{A\bar{A}k}$, $C_{\bar{A}Ak}$ and $C_{\bar{A}\bar{A}k}$. For example, if for each of three synthetic observations, $C_{AA} = 1$ and all other weights are negligible, then we would conclude *a posteriori* that $\langle \lambda_A \rangle = 0.875$ and $\sigma_{\lambda_A} = 0.11$. This approach was adopted in Fig. 7, in contrast to the preceding figures which adopted fixed natal spins for all BHs.

APPENDIX C: APPROXIMATE POSTERIOR DISTRIBUTION FOR GW170104

For most events examined in this study, we made use of posterior samples provided and performed by the LIGO Scientific Collaboration, generated by comparing each event to the IMRPv2 approximation [72]. Because we cannot employ the same level of detail for GW170104, we instead resort to an approximate posterior distribution, derived from the reported GW170104 results [5] and our understanding of gravitational wave parameter estimation, as approximated using a Fisher matrix [97].

For GW170104 we construct an approximate (truncated) Gaussian posterior distribution in only three correlated

binary parameters: $\mathcal{M}_c, \eta, \chi_{\text{eff}}$. The shape of this Gaussian (i.e., its inverse covariance matrix) was constructed via a Fisher matrix approximation, derived using the median detector-frame parameters reported for GW170104 (i.e., $m_1 \approx 37.1M_\odot$, $m_2 \approx 22.6M_\odot$, and—breaking degeneracy with an ad-hoc choice— $\chi_{1,z}\chi_{2,z} \approx \chi_{\text{eff}} \approx -0.12$); the reported network SNR of GW170104 (i.e., $\rho \approx 13.0$); and a suitable single-detector noise power spectrum. Our effective Fisher matrix estimate for the inverse covariance matrix Γ [98] adopted the noise power spectrum at GW150914, using a minimum frequency $f_{\text{min}} = 30$ Hz; employed the (nonprecessing) SEOBNRv4 approximation [99], evaluated on a grid in $\mathcal{M}_c, \eta, \chi_{1,z}, \chi_{2,z}$; and fit as a quadratic function of $\mathcal{M}_c, \eta, \chi_{\text{eff}}$. We adopt a nonprecessing model and lower-dimensional Fisher matrix approximation because the posterior of this event, like GW150914, is consistent with nonprecessing spins and is very well approximated, in these parameters, by a nonprecessing model; see, e.g., [49]. This simple approximation captures

important correlations between \mathcal{M}_c, η and χ_{eff} , and the diagonal terms of $\Gamma^{-1}\rho^2$ roughly reproduce the width of the posterior distribution reported for GW170104. To obtain better agreement with the reported one-dimensional credible intervals, we scaled the terms $\Gamma_{\mathcal{M}_c, x}$ for $x = \mathcal{M}_c, \eta, \chi_{\text{eff}}$ by a common scale factor 0.29 and the term $\Gamma_{\chi_{\text{eff}}, \chi_{\text{eff}}}$ by 0.9. For similar reasons, we likewise hand-tuned the center of the Gaussian distribution to the (unphysical) parameter location to $\mathcal{M}_c = 22.9, \eta = 0.32, \chi_{\text{eff}} = 0.013$. Using this ansatz, we generate GW170104-like posterior samples in $\mathcal{M}_c, \eta, \chi_{\text{eff}}$ from this Gaussian distribution, truncating any unphysical samples (i.e., with $\eta > 1/4$). For our tuned posterior, the median and 90% credible regions on the synthetic posteriors approximate the values and ranges reported. According to our highly simplified and purely synthetic approach, the resulting 90% credible regions are $M_{\text{tot}} = 51.2_{-6.8}^{+7.6} M_\odot$, $q = 0.62_{-0.24}^{+0.25}$, $\chi_{\text{eff}} = -0.12_{-0.27}^{+0.28}$.

-
- [1] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Astrophys. J.* **818**, L22 (2016).
- [3] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Astrophys. J.* **833**, L1 (2016).
- [4] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Phys. Rev. X* **6**, 041015 (2016).
- [5] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Phys. Rev. Lett.* **118**, 221101 (2017).
- [6] I. Mandel and S. E. de Mink, *Mon. Not. R. Astron. Soc.* **458**, 2634 (2016).
- [7] P. Marchant, N. Langer, P. Podsiadlowski, T. M. Tauris, and T. J. Moriya, *Astron. Astrophys.* **588**, A50 (2016).
- [8] C. L. Rodriguez, C.-J. Haster, S. Chatterjee, V. Kalogera, and F. A. Rasio, *Astrophys. J.* **824**, L8 (2016).
- [9] S. Bird, I. Cholis, J. B. Muñoz, Y. Ali-Haïmoud, M. Kamionkowski, E. D. Kovetz, A. Raccañelli, and A. G. Riess, *Phys. Rev. Lett.* **116**, 201301 (2016).
- [10] D. Kushnir, M. Zaldarriaga, J. A. Kollmeier, and R. Waldman, *Mon. Not. R. Astron. Soc.* **462**, 844 (2016).
- [11] A. Lamberts, S. Garrison-Kimmel, D. R. Clausen, and P. F. Hopkins, *Mon. Not. R. Astron. Soc.* **463**, L31 (2016).
- [12] I. Dvorkin, E. Vangioni, J. Silk, J.-P. Uzan, and K. A. Olive, *Mon. Not. R. Astron. Soc.* **461**, 3877 (2016).
- [13] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Phys. Rev. Lett.* **116**, 131102 (2016).
- [14] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, *Astrophys. J.* **832**, L2 (2016).
- [15] R. O’Shaughnessy, D. Gerosa, and D. Wysocki, *Phys. Rev. Lett.* **119**, 011101 (2017).
- [16] K. Breivik, C. L. Rodriguez, S. L. Larson, V. Kalogera, and F. A. Rasio, *Astrophys. J.* **830**, L18 (2016).
- [17] A. Nishizawa, E. Berti, A. Klein, and A. Sesana, *Phys. Rev. D* **94**, 064020 (2016).
- [18] A. V. Tutukov and L. R. Yungelson, *Mon. Not. R. Astron. Soc.* **260**, 675 (1993).
- [19] V. M. Lipunov, K. A. Postnov, and M. E. Prokhorov, *Mon. Not. R. Astron. Soc.* **288**, 245 (1997).
- [20] K. Belczynski, V. Kalogera, and T. Bulik, *Astrophys. J.* **572**, 407 (2002).
- [21] G. Nelemans, L. R. Yungelson, and S. F. P. Zwart, *Astron. Astrophys.* **375**, 890 (2001).
- [22] R. Voss and T. M. Tauris, *Mon. Not. R. Astron. Soc.* **342**, 1169 (2003).
- [23] M. Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, T. Bulik, I. Mandel, and R. O’Shaughnessy, *Astrophys. J.* **759**, 52 (2012).
- [24] M. Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, T. Bulik, I. Mandel, and R. O’Shaughnessy, *Astrophys. J.* **779**, 72 (2013).
- [25] M. Dominik, E. Berti, R. O’Shaughnessy, I. Mandel, K. Belczynski, C. Fryer, D. E. Holz, T. Bulik, and F. Pannarale, *Astrophys. J.* **806**, 263 (2015).
- [26] K. Belczynski, D. E. Holz, T. Bulik, and R. O’Shaughnessy, *Nature (London)* **534**, 512 (2016).
- [27] J. J. Eldridge and E. R. Stanway, *Mon. Not. R. Astron. Soc.* **462**, 3302 (2016).
- [28] S. Stevenson, A. Vigna-Gómez, I. Mandel, J. W. Barrett, C. J. Neijssel, D. Perkins, and S. E. de Mink, *Nat. Commun.* **8**, 14906 (2017).
- [29] G. Hobbs, D. R. Lorimer, A. G. Lyne, and M. Kramer, *Mon. Not. R. Astron. Soc.* **360**, 974 (2005).
- [30] D. Gerosa, M. Kesden, E. Berti, R. O’Shaughnessy, and U. Sperhake, *Phys. Rev. D* **87**, 104028 (2013).
- [31] V. Kalogera, *Astrophys. J.* **541**, 319 (2000).

- [32] R. M. Hjellming and M. P. Rupen, *Nature (London)* **375**, 464 (1995).
- [33] J. A. Orosz and C. D. Bailyn, *Astrophys. J.* **477**, 876 (1997).
- [34] J. A. Orosz, E. Kuulkers, M. van der Klis, J. E. McClintock, M. R. García, P. J. Callanan, C. D. Bailyn, R. K. Jain, and R. A. Remillard, *Astrophys. J.* **555**, 489 (2001).
- [35] R. G. Martin, C. A. Tout, and J. E. Pringle, *Mon. Not. R. Astron. Soc.* **401**, 1514 (2010).
- [36] S. Repetto, M. B. Davies, and S. Sigurdsson, *Mon. Not. R. Astron. Soc.* **425**, 2799 (2012).
- [37] S. Repetto and G. Nelemans, *Mon. Not. R. Astron. Soc.* **453**, 3341 (2015).
- [38] S. Repetto, A. P. Igoshev, and G. Nelemans, *Mon. Not. R. Astron. Soc.* **467**, 298 (2017).
- [39] K. Belczynski, S. Repetto, D. E. Holz, R. O’Shaughnessy, T. Bulik, E. Berti, C. Fryer, and M. Dominik, *Astrophys. J.* **819**, 108 (2016).
- [40] T.-W. Wong, F. Valsecchi, T. Fragos, and V. Kalogera, *Astrophys. J.* **747**, 111 (2012).
- [41] T.-W. Wong, F. Valsecchi, A. Ansari, T. Fragos, E. Glebbeek, V. Kalogera, and J. McClintock, *Astrophys. J.* **790**, 119 (2014).
- [42] K. Belczynski, J. Klencki, G. Meynet, C. L. Fryer, D. A. Brown, M. Chruslinska, W. Gladysz, R. O’Shaughnessy, T. Bulik, E. Berti, D. E. Holz, D. Gerosa, M. Giersz, S. Ekstrom, C. Georgy, A. Askar, J.-P. Lasota, and D. Wysocki, [arXiv:1706.07053](https://arxiv.org/abs/1706.07053).
- [43] K. Belczyński and T. Bulik, *Astron. Astrophys.* **346**, 91 (1999).
- [44] É. Racine, *Phys. Rev. D* **78**, 044021 (2008).
- [45] M. Kesden, D. Gerosa, R. O’Shaughnessy, E. Berti, and U. Sperhake, *Phys. Rev. Lett.* **114**, 081103 (2015).
- [46] D. Gerosa, M. Kesden, U. Sperhake, E. Berti, and R. O’Shaughnessy, *Phys. Rev. D* **92**, 064016 (2015).
- [47] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin, W. D. Pozzo, F. Feroz, J. Gair, C.-J. Haster, V. Kalogera, T. Littenberg, I. Mandel, R. O’Shaughnessy, M. Pitkin, C. Rodriguez, C. Röver, T. Sidery, R. Smith, M. Van Der Sluys, A. Vecchio, W. Voursden, and L. Wade, *Phys. Rev. D* **91**, 042003 (2015).
- [48] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Phys. Rev. Lett.* **116**, 241102 (2016).
- [49] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Phys. Rev. D* **94**, 064035 (2016).
- [50] D. Trifirò, R. O’Shaughnessy, D. Gerosa, E. Berti, M. Kesden, T. Littenberg, and U. Sperhake, *Phys. Rev. D* **93**, 044071 (2016).
- [51] D. Gerosa, R. O’Shaughnessy, M. Kesden, E. Berti, and U. Sperhake, *Phys. Rev. D* **89**, 124025 (2014).
- [52] S. Vitale, R. Lynch, R. Sturani, and P. Graff, *Classical Quantum Gravity* **34**, 03LT01 (2017).
- [53] S. Stevenson, C. P. L. Berry, and I. Mandel, *Mon. Not. R. Astron. Soc.* **471**, 2801 (2017).
- [54] W. M. Farr, S. Stevenson, M. C. Miller, I. Mandel, B. Farr, and A. Vecchio, *Nature (London)* **548**, 426 (2017).
- [55] D. Gerosa, R. O’Shaughnessy, K. Belczynski, W. Gladerys, D. Wysocki, E. Berti, M. Kesden *et al.* (to be published).
- [56] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevD.97.043014>, which contains ASCII tables providing the means, rates and integrals in Equation (2), which may be used to reproduce many of our analyses, and explore combinations of models not investigated in this paper.
- [57] K. Belczynski, V. Kalogera, F. A. Rasio, R. E. Taam, A. Zezas, T. Bulik, T. J. Maccarone, and N. Ivanova, *Astrophys. J. Suppl. Ser.* **174**, 223 (2008).
- [58] C. L. Fryer, K. Belczynski, G. Wiktorowicz, M. Dominik, V. Kalogera, and D. E. Holz, *Astrophys. J.* **749**, 91 (2012).
- [59] K. Belczynski, A. Heger, W. Gladysz, A. J. Ruiter, S. Woosley, G. Wiktorowicz, H.-Y. Chen, T. Bulik, R. O’Shaughnessy, D. E. Holz, C. L. Fryer, and E. Berti, *Astron. Astrophys.* **594**, A97 (2016).
- [60] G. E. Brown, C.-H. Lee, R. A. M. J. Wijers, H. K. Lee, G. Israelian, and H. A. Bethe, *New Astron.* **5**, 191 (2000).
- [61] J. M. Bardeen, *Nature (London)* **226**, 64 (1970).
- [62] K. S. Thorne, *Astrophys. J.* **191**, 507 (1974).
- [63] R. O’Shaughnessy, J. Kaplan, V. Kalogera, and K. Belczynski, *Astrophys. J.* **632**, 1035 (2005).
- [64] P. Hut, *Astron. Astrophys.* **99**, 126 (1981).
- [65] T. Mazeh, *EAS Publ. Ser.* **29**, 1 (2008).
- [66] S. Albrecht, J. Setiawan, G. Torres, D. C. Fabrycky, and J. N. Winn, *Astrophys. J.* **767**, 32 (2013).
- [67] K. Belczynski, R. E. Taam, V. Kalogera, F. A. Rasio, and T. Bulik, *Astrophys. J.* **662**, 504 (2007).
- [68] K. Belczynski, R. E. Taam, E. Rantsiou, and M. van der Sluys, *Astrophys. J.* **682**, 474 (2008).
- [69] A. R. King and U. Kolb, *Mon. Not. R. Astron. Soc.* **305**, 654 (1999).
- [70] P. Natarajan and J. E. Pringle, *Astrophys. J.* **506**, L97 (1998).
- [71] D. Gerosa and M. Kesden, *Phys. Rev. D* **93**, 124066 (2016).
- [72] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [73] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, N. W. Taylor, and A. Zenginoglu, *Phys. Rev. D* **89**, 061502 (2014).
- [74] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, C. D. Ott, M. Boyle, L. E. Kidder, H. P. Pfeiffer, and B. Szilágyi, *Phys. Rev. D* **96**, 024058 (2017).
- [75] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, D. A. Hemberger, P. Schmidt, and R. Smith, *Phys. Rev. D* **95**, 104023 (2017).
- [76] R. O’Shaughnessy, V. Kalogera, and K. Belczynski, *Astrophys. J.* **716**, 615 (2010).
- [77] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), *Phys. Rev. Lett.* **116**, 241103 (2016).
- [78] T. J. Loredo, *AIP Conf. Proc.* **735**, 195 (2004).
- [79] S. Stevenson, F. Ohme, and S. Fairhurst, *Astrophys. J.* **810**, 58 (2015).
- [80] R. O’Shaughnessy, *Phys. Rev. D* **88**, 084061 (2013).
- [81] S. Vitale, D. Gerosa, C.-J. Haster, K. Chatzioannou, and A. Zimmerman, *Phys. Rev. Lett.* **119**, 251103 (2017).
- [82] A. R. Williamson, J. Lange, R. O’Shaughnessy, J. A. Clark, P. Kumar, J. C. Bustillo, and J. Veitch, *Phys. Rev. D* **96**, 124041 (2017).

- [83] M. Zevin, C. Pankow, C. L. Rodriguez, L. Sampson, E. Chase, V. Kalogera, and F. A. Rasio, *Astrophys. J.* **846**, 82 (2017).
- [84] C. Capano, I. Harry, S. Privitera, and A. Buonanno, *Phys. Rev. D* **93**, 124007 (2016).
- [85] C. Talbot and E. Thrane, *Phys. Rev. D* **96**, 023012 (2017).
- [86] R. O'Shaughnessy, Presentation at Syracuse University (2009).
- [87] R. O'Shaughnessy, C. Kim, T. Fragos, V. Kalogera, and K. Belczynski, *Astrophys. J.* **633**, 1076 (2005).
- [88] R. O'Shaughnessy, C. Kim, V. Kalogera, and K. Belczynski, *Astrophys. J.* **672**, 479 (2008).
- [89] R. O'Shaughnessy, K. Belczynski, and V. Kalogera, *Astrophys. J.* **675**, 566 (2008).
- [90] I. Mandel and R. O'Shaughnessy, *Classical Quantum Gravity* **27**, 114007 (2010).
- [91] T. Prusti, J. H. J. de Bruijne, A. G. A. Brown, A. Vallenari, C. Babusiaux, C. A. L. Bailer-Jones, U. Bastian, M. Biermann, D. W. Evans *et al.* (Gaia Collaboration), *Astron. Astrophys.* **595**, A1 (2016).
- [92] A. G. A. Brown, A. Vallenari, T. Prusti, J. H. J. de Bruijne, F. Mignard, R. Drimmel, C. Babusiaux, C. A. L. Bailer-Jones, U. Bastian *et al.* (Gaia Collaboration), *Astron. Astrophys.* **595**, A2 (2016).
- [93] A. Dempster, N. Laird, and D. Rubin, *J. R. Stat. Soc. Ser. B* **39**, 1 (1977).
- [94] K. P. Murphy, Machine learning: a probabilistic perspective (2012), <https://mitpress.mit.edu/books/machine-learning-0>.
- [95] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Journal of Machine Learning Research* **12**, 2825 (2011).
- [96] D. W. Hogg, A. D. Myers, and J. Bovy, *Astrophys. J.* **725**, 2166 (2010).
- [97] C. Cutler and É. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [98] H.-S. Cho, E. Ochsner, R. O'Shaughnessy, C. Kim, and C.-H. Lee, *Phys. Rev. D* **87**, 024004 (2013).
- [99] A. Bohé, L. Shao, A. Taracchini, A. Buonanno, S. Babak, I. W. Harry, I. Hinder, S. Ossokine, M. Pürrer, V. Raymond, T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **95**, 044028 (2017).