Impact of Bayesian Priors on the Characterization of Binary Black Hole Coalescences

Salvatore Vitale,[1,*] Davide Gerosa,[2,†] Carl-Johan Haster,[3,‡] Katerina Chatziioannou,[3,§] and Aaron Zimmerman[3,‖]
[1]*LIGO, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA*
[2]*TAPIR 350-17, California Institute of Technology, 1200 East California Boulevard, Pasadena, California 91125, USA*
[3]*Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto, Ontario M5S 3H8, Canada*
(Received 24 August 2017; published 20 December 2017)

In a regime where data are only mildly informative, prior choices can play a significant role in Bayesian statistical inference, potentially affecting the inferred physics. We show this is indeed the case for some of the parameters inferred from current gravitational-wave measurements of binary black hole coalescences. We reanalyze the first detections performed by the twin LIGO interferometers using alternative (and astrophysically motivated) prior assumptions. We find different prior distributions can introduce deviations in the resulting posteriors that impact the physical interpretation of these systems. For instance, (i) limits on the 90% credible interval on the effective black hole spin $\chi_{\rm eff}$ are subject to variations of $\sim$10% if a prior with black hole spins mostly aligned to the binary's angular momentum is considered instead of the standard choice of isotropic spin directions, and (ii) under priors motivated by the initial stellar mass function, we infer tighter constraints on the black hole masses, and in particular, we find no support for any of the inferred masses within the putative mass gap $M \lesssim 5\ M_\odot$.

DOI: 10.1103/PhysRevLett.119.251103

Introduction.—Prior distributions are at the cornerstone of Bayesian statistics, where experimental data are used to update prior knowledge into posterior beliefs about observed phenomena. As Bayes himself put it, "there is reason to expect an event with more or less confidence according to the greater or less number of times in which, under given circumstances, it has happened without failing" [1]. Prior assumptions inevitably enter into all types of statistical inference, and the prominent role of priors is one of the main advantages of the Bayesian framework: the data analyst is forced to consider and explicitly specify how priors are incorporated into the analysis, thus avoiding the pitfalls of poorly specified assumptions [2,3].

While choosing priors may in some cases be straightforward (for instance, a mass should always be positive), care must be taken when strong knowledge of the expected distribution is not available. In particular, in a regime where data are weakly informative and a feeble signal needs to be dug out of instrumental noise, different priors may indeed lead to different conclusions.

In this Letter we show that the current statistical inference on some black hole (BH) binary parameters through gravitational-wave (GW) observations does fall in this regime. With four likely observations of binary BH coalescences announced, and many more expected in the coming years, GW astronomy is becoming a reality [4–7]. Bayesian analysis has been used to measure the physical parameters of the binary coalescences in LIGO's observing runs [8,9], the BH merger rate [6,10], as well as to perform tests of general relativity [6,11]. Some of the physical parameters of the sources, e.g., the total mass (for heavier systems [12–14]) or the chirp mass (for lighter binaries [15–17]), strongly affect the phasing evolution of the GW signal. As a consequence, these parameters are usually well measured. Others, such as the mass ratio and spins, have a smaller impact on the waveform and are therefore harder to measure. Measurements of these parameters are more directly affected by the chosen prior (see Refs. [18,19]).

Accurate statistical inference about BH spins is of crucial astrophysical importance. While mass and rates distributions sensibly overlap in many different scenarios (see, e.g., Refs. [20,21]), spins are arguably the best indicator of BH binary formation channels [22–25] (residual eccentricities may also provide a promising avenue; see, e.g., Ref. [26]). Several studies have already demonstrated the potential of spin measurements to discriminate between different pathways of BH binary formation and evolution [17,27–33]. At the same time, precise and accurate estimation of the individual BH masses can be used to reconstruct their mass function [7] and to verify if BHs lighter than 5 $M_\odot$ exist [34]. Quantifying the effect of the prior choice is therefore a crucial step to make solid astrophysical statements using GW data.

Current priors.—Parameter estimations for the first GW events GW150914 [4,7,9,35], GW151226 [5,7], GW170104 [6], and candidate LVT151012 [7,36] were performed using priors (i) uniform in component masses m_1 and m_2, (ii) uniform in dimensionless spin magnitude $\chi_i = |\mathbf{S}_i|/m_i^2$, and (iii) isotropic in spin directions at the reference GW frequency of 20 Hz. While masses and spin magnitudes are constant up to high post-Newtonian (PN) order, spin directions are subject to change due to orbital plane and spin precession (see, e.g., Refs. [37,38]). By the time BH binaries enter the LIGO band, their spin

 © 2017 American Physical Society

misalignments may be different from those at formation, which are needed to discriminate formation channels [39]. Fortunately, isotropic spin distributions are kept isotropic when evolved under the 2PN spin precession equations [39–41], and the impact of this issue on the usual prior can be neglected (but see P_5 below).

Despite being at first sight reasonable, the isotropic spin prior distribution carries important astrophysical consequences. It is generally believed that binaries formed in isolation in the galactic field will have, on average, some tendency towards spins aligned with the orbital angular momentum. On the other hand, spin directions of binaries formed via dynamical interactions in stellar clusters are expected to be isotropically distributed. We are therefore in a risky situation: we may be biasing our astrophysical inference by assuming *a priori* one of the models we try to discriminate.

A similar note can be made regarding the choice of distributing spins uniformly in magnitude. If we were to base our previous knowledge on other observed BH systems, then moderately high spins should be favored as found in most x-ray binaries [42]. On the other hand, core-envelope interactions in massive stars may produce BHs with small spins [43,44] (and it has been suggested that primordial BH spins should also be low [45]). One may therefore want to choose a prior which is peaked at either low or high spins, or perhaps even bimodal. An agnostic approach would be to consider the BH spins as vectors and draw them uniformly in volume, rather than uniformly in magnitude and isotropic in direction. Alternatively, one might naively assume black holes form in situations where a random amount of energy goes into the spin, and draw uniformly in specific rotational energy $E_{\rm rot} \equiv 1 - \sqrt{1+\sqrt{1-\chi^2}}/\sqrt{2}$.

The spin parameter which is measured best (arguably the only spin parameter which is currently measured at all [6,7]) is the so-called effective spin [46],

$$\chi_{\rm eff} = \frac{\mathbf{S}_1/m_1 + \mathbf{S}_2/m_2}{m_1+m_2}\cdot\hat{\mathbf{L}} = \frac{\chi_1\cos\theta_1 + q\chi_2\cos\theta_2}{1+q}, \quad (1)$$

where $q = m_2/m_1 \leq 1$ is the mass ratio, and $\theta_i = \arccos(\hat{\mathbf{S}}_i\cdot\hat{\mathbf{L}})$ are the angles between the spins $\mathbf{S}_i$ and the binary's orbital angular momentum $\hat{\mathbf{L}}$. It is clear from Eq. (1) that the mass ratio, spin magnitude, and spin direction priors are all entangled in determining the prior distribution of $\chi_{\rm eff}$.

Prior choices.—In order to explore some of these issues and gauge the impact of priors on Bayesian inference, we have reanalyzed the BH coalescences detected by LIGO during its first observing run (O1) using a variety of alternative prior distributions. Results have been obtained using the nested sampling algorithm implemented in LALINFERENCE [8] and a reduced-order quadrature [48] implementation of the IMRPHENOMPV2 waveform model, which partially accounts for spin precession effects through a single parameter χ_p [49]. We restrict our study to (detector-frame [50]) chirp masses $8\ M_\odot \leq M_c \leq 45\ M_\odot$, mass ratios $q \geq 1/8$, and dimensionless spin magnitudes $\chi_i \leq 0.89$ [48]. These restrictions are not a problem for our study since for none of the runs we performed the posterior distributions had support near these boundaries. We have analyzed the 32-sec data frames publicly released by the LIGO/Virgo Collaboration [51], using the BAYESWAVE algorithm [52,53] to estimate the on-source power spectral density needed for the likelihood evaluations [8], marginalizing over calibration uncertainties as in Refs. [6,7].

Each event was analyzed multiple times, using one of the following priors: uniform in individual masses and spin magnitudes, isotropic in spin direction (P_1, the default choice used in LIGO analyses) [54]; uniform in individual masses and rotational energy of the BHs, isotropic in spin direction (P_2); uniform in individual masses, spin vectors uniform in volume (P_3, volumetric); uniform in individual masses, bimodal in spin magnitudes, $P(\chi_i) \propto \exp[-(\chi_i-\mu_1)^2/(2\sigma^2)] + \exp[-(\chi_i-\mu_2)^2/(2\sigma^2)]$ with $\mu_1 = 0$, $\mu_2 = 0.89$, $\sigma = 0.1$, isotropic in spin directions (P_4) [55–57]; uniform in individual masses and spin magnitude, peaked around alignment for spin direction $P(\theta_i) \propto \exp[-(\cos\theta_i - \mu)^2/2\sigma^2]$ where $\mu = 1$, $\sigma = 1 - \cos(10°)$ [P_5; cf. Ref. [29], which suggested these parameters could be inferred with $\mathcal{O}(50)$ observations]; power law in primary's mass $P(m_1) \propto m^\alpha$ with $\alpha = -2.3$ (as in Kroupa's initial mass function for massive stars [58], cf. also Ref. [59] and references therein), uniform in secondary's mass, uniform in spin magnitude, isotropic in spin direction (P_6); power law in m_1 with $\alpha = -2.3$, logistic prior in the mass ratio $P(q) \propto 1/\{1 + \exp[-k(q-q_0)]\}$ with $k = 20$, $q_0 = 0.8$ (this is meant to mimic numerical results of BH mergers in globular clusters [60]), uniform in spin magnitude, isotropic in spin direction (P_7); uniform in individual masses, Gaussian around zero for dimensionless spin magnitude $P(\chi_i) \propto \exp[-(\chi_i - \mu)^2/(2\sigma^2)]$ with $\mu = 0$, $\sigma = 0.1$ (P_8). In what follows we occasionally refer to the evidence of the priors. What is meant is the evidence of a Bayesian *model* where the spins and masses are distributed as in that prior.

Spins.—Marginalized prior and posterior distributions in $\chi_{\rm eff}$ are shown in Fig. 1. Table I shows the corresponding values of the Kullback-Leibler (KL) divergence [61] (a measurement of the information gain between prior and posterior, or equivalently the relative entropy between two distributions) for both spin parameters $\chi_{\rm eff}$ and χ_p, together with the odds ratio $\log_{10}\mathcal{O}$ between the posteriors obtained with each of our choice P_n and the default analysis P_1.

For GW150914, all results are compatible with zero $\chi_{\rm eff}$, which might come from a combination of small spins or in-plane spins; see Eq. (1). Small spin magnitudes are preferred for some prior choices. The bimodal prior P_4 is particularly interesting: when restricted to choose

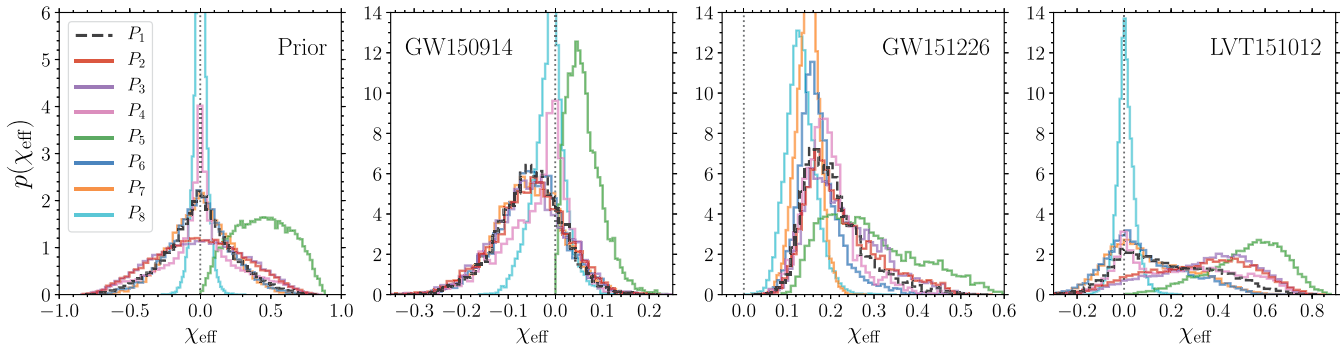

FIG. 1. Marginalized prior and posterior distributions on $\chi_{\rm eff}$ for each of the current GW events. Dashed black line shows results obtained with canonical prior choice P_1 (uniform in masses, spin magnitude, and isotropic in spin directions), while lighter colored lines shows our alternative prior assumptions P_2–P_8. In particular, the event GW150914 is compatible with nonspinning BHs, while $\chi_{\rm eff} < 0$ is excluded at the 99% credible level for GW151226 for all prior choices.

between high and low spins only, the data clearly favor the low spin mode for both objects, with $p(\chi_1 < 0.3) = 0.78$, $p(\chi_2 < 0.3) = 0.68$. If aligned spins are assumed (P_5), the posterior distribution still peaks close to the nonspinning configuration $\chi_{\rm eff} = 0$ with high information gain ($D_{\rm KL}^{\chi_{\rm eff}} \sim 4.2$, $D_{\rm KL}^{\chi_p} \sim 1.7$) but low odds ($\mathcal{O} \sim 0.02$), thus suggesting the algorithm cannot model the data equally well if both tilts are low. Together, these observations indicate that the conclusion that $\chi_i \lesssim 0.3$ for GW150914 is robust to changes in the prior.

Conversely, for GW151226 all priors exclude $\chi_{\rm eff} \leq 0$ at the 99% credible level, thus confirming with extremely high significance that at least one of the two BHs was spinning [5]. The bimodal prior P_4 in this case favors the high spin mode for the spin of the heavier BH, while both modes are equally likely for the less massive object: $p(\chi_1 > 0.445) = 0.83$; $p(\chi_2 > 0.445) = 0.59$ [but $p(\chi_1 < 0.445 \cup \chi_2 < 0.445) < 0.01$]. The case of the aligned-spin prior P_5 presents important astrophysical consequences. With odds very similar to the P_1 run, P_5 allows for posterior values of $\chi_{\rm eff}$ as large as 0.49 at the 90% credible level, thus allowing for moderately large spin magnitudes. Interestingly, priors P_6 and P_7 lead to narrower posteriors compared to P_1, as they both place more support in the $q \lesssim 1$ region, thus partly breaking the $q - \chi_{\rm eff}$ degeneracy [7]. This effect is less important for GW150914 because of its higher total mass.

Finally, prior effects are even more pronounced for LVT151012. This is not surprising, as its lower signal-to-noise ratio indicates the data are less informative. In particular, two modes appear to be present in the marginalized posterior of the effective spin, located at $\chi_{\rm eff} \sim 0$ and $\chi_{\rm eff} \sim 0.5$, respectively. Which of the two modes is preferred depends on the prior distribution: P_4, P_6, P_7, and P_8 prefer the low-$\chi_{\rm eff}$ mode, while P_2, P_3, and especially P_5 favor higher values of $\chi_{\rm eff}$.

As shown in Table I, the values of $D_{\rm KL}^{\chi_p}$ are close to zero for most priors and all events, thus indicating the χ_p prior distribution is typically returned as a posterior almost unchanged by the data. As already mentioned, an exception is P_5 for GW150914, which however is disfavored at $\mathcal{O} \sim 0.02$. For GW151226, the P_4 prior yields a posterior on χ_p which is quite different from the prior, $D_{\rm KL}^{\chi_p} = 0.7$ bits, and models that data comparably well. This happens because the prior of the spins magnitude is bimodal, while the posterior prefer the high spin mode for the primary.

TABLE I. For each of the three O1 events, we show the KL divergence $D_{\rm KL}$ on (i) $\chi_{\rm eff}$ and (ii) χ_p measuring the information gain between prior to posterior in bits, and (iii) the Bayesian odds ratio $\log_{10}\mathcal{O}$ of each single analysis P_n compared to the standard one P_1. P_n (P_1) is preferred if $\log_{10}\mathcal{O}$ is positive (negative). For comparison, the KL divergence between the $\chi_{\rm eff}$ P_1 prior and a uniform distribution over the same range is 0.82 bits. The log odds have an uncertainty of ± 0.04.

	GW150914			GW151226			LVT151012		
	$D_{\rm KL}^{\chi_{\rm eff}}$	$D_{\rm KL}^{\chi_p}$	$\log_{10}\mathcal{O}$	$D_{\rm KL}^{\chi_{\rm eff}}$	$D_{\rm KL}^{\chi_p}$	$\log_{10}\mathcal{O}$	$D_{\rm KL}^{\chi_{\rm eff}}$	$D_{\rm KL}^{\chi_p}$	$\log_{10}\mathcal{O}$
P_1	1.02	0.03	⋯	1.93	0.21	⋯	0.53	0.03	⋯
P_2	1.36	0.06	−0.3	1.78	0.04	0.0	0.89	0.05	−0.1
P_3	1.52	0.09	−0.4	1.76	0.02	0.0	0.95	0.04	0.0
P_4	0.88	0.12	0.0	2.56	0.70	−0.1	0.61	0.12	−0.1
P_5	4.21	1.75	−1.7	0.82	0.21	0.0	0.22	0.07	0.5
P_6	0.96	0.01	0.1	2.12	0.08	0.4	0.24	0.00	0.4
P_7	0.93	0.06	0.4	2.63	0.02	0.4	0.26	0.01	0.5
P_8	0.14	0.07	0.3	4.82	0.70	−1.7	0.03	0.02	−0.1

Since values of both the (detector-frame) masses and of $\chi_{\rm eff}$ are similar, we expect the most recent event GW170104 to present the same trends as GW150914 (perhaps with broader posterior distributions due to the lower signal-to-noise ratio). We verified this by constructing a software replica of GW170104 with parameters consistent to those presented in Ref. [6] and simulating its noise power spectral density using the same method as Ref. [62]. Variations in the 90% credible interval to $\chi_{\rm eff}$ of up to ~30% are observed: while priors which include both high and misaligned spins all return 90% credible intervals of $-0.4 \lesssim \chi_{\rm eff} \lesssim 0.1$, inference with the low-spin prior P_8 returns $-0.07 < \chi_{\rm eff} < 0.05$, while $0.03 < \chi_{\rm eff} < 0.23$ if aligned spins are assumed (P_5).

Component masses.—Different choices for the component mass priors also carry important astrophysical implications. Figure 2 shows posterior distributions of the two BH masses for the three O1 events using the default P_1 prior [63], as well as P_6 and P_7. Both P_6 and P_7 have larger prior support for binary mass ratio close to unity compared to P_1. This additional weight at comparable masses has a visible effect on the posteriors. While presenting odds similar to P_1, the resulting posterior distributions for P_6 and P_7 now prefer the region closer to the $m_1 = m_2$ line in Fig. 2. In particular, when GW151226 is analyzed with any of these two priors, the 99% credible interval for the source-frame mass of the secondary object is above 5 $M_\odot$.

This point has serious consequences in the astrophysical interpretation of the results. Electromagnetic observations of neutron stars and stellar-mass BHs (cf. Ref. [42] and references therein) hint at a putative mass gap between the highest neutron star masses ($m \lesssim 3\ M_\odot$) and the lowest BH masses ($m \gtrsim 5\ M_\odot$). Current measurements, however, are not conclusive since the lack of BHs at lower masses could be entirely due to selection effects [64]. The confirmation or exclusion of the mass gap by GW observations is expected to provide unique insights on stellar collapse and compact-object formation [65]. Figure 2 shows that, when analyzed with priors motivated by stellar physics like Kroupa's initial mass function, GW data for GW151226 are fully consistent with the existence of a mass gap. Both P_6 and P_7 are slightly favored over the default prior, with $\mathcal{O} \approx 2.5$. A careful considerations of priors may thus be important to securely discriminate between BHs and neutron stars [66–71] and thus establish the presence of the mass gap between the two classes of sources.

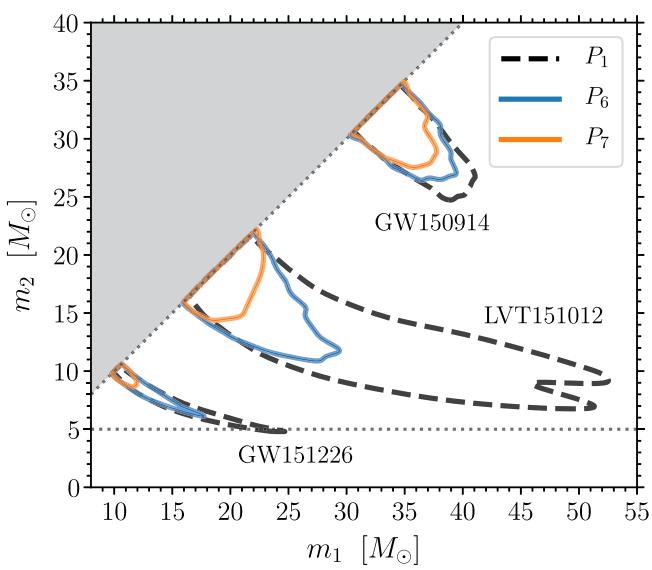

FIG. 2. Posterior distribution for the individual masses m_1 and $m2$ (with $m_2 < m_1$) derived using the canonical prior P_1 and two other choices motivated by stellar physics (P_6 and P_7). For GW151226, the region $m_2 < 5\ M_\odot$ is excluded at $> 99\%$ probability for both P_6 and P_7.

Finally, we have verified that the marginalized chirp mass posteriors are stable over the change of priors. For GW150914, all cases except P_5 (which, however, presents lower odds) yield posterior medians within a ~0.5 $M_\odot$ interval. The median for P_5 is 1 $M_\odot$ larger than that of P_1 (for comparison, the 90% credible interval for the P_1 run is ~3 $M_\odot$). For GW151226, all runs yield posterior medians within a ~0.03 $M_\odot$ interval, compared with a 90% credible interval of ~0.6 $M_\odot$ for the P_1 run.

Conclusions and future prospects.—In this Letter we have shown how different prior choices can influence the statistical inference on the parameters of binary BHs. We have estimated the parameters of the first GW events detected under different prior assumptions and verified that both the component masses and the effective spin are impacted by the choice of the prior. For example, for GW150914 and the software replica of GW170104, the effective spin can be made close to zero using a prior that prefers small spin magnitudes, without significant loss of evidence. Conversely, for GW151226 all priors we used give a posterior density for $\chi_{\rm eff}$ which is positive at the 99% credible interval. Using the default priors for GW151226 results in a source-frame mass posterior for the lighter BH which has support in the suggested mass gap between BHs and neutron stars. We have shown that this conclusion does not hold if a different, yet reasonable, prior on the masses (i.e., a power law modeled on the initial stellar mass function with and without a logistic function on the mass ratio) is used.

As exemplified by the case of LVT151012, the effect of different prior choices is more severe for weak GW signals. Data will be more informative for future loud events and, eventually, more and more physical conclusions will become robust with respect to the details of the prior choice. This point will be specifically addressed in future work, together with a wider variety of prior distributions more carefully modeled on BH formation pathways, along with the effect of the PN spin evolution [72]. Other investigations of prior effects in GW data analysis are also underway [73].

We hope our work may spark new efforts at incorporating a range of prior choices into inferences about the progenitors

of binary BHs. Instead of attempting to explain constraints obtained from a single prior choice, we argue different astrophysical models should be used as priors in the data analysis process. Model selection should then be applied to assess which model better explains the observations. For instance, posteriors calculated under different priors naturally enter in hierarchical analyses [29,30,33,74–76], or when one wants to calculate the branching ratio between populations [27]. More simply, one can just use the evidences calculated for various priors, which we report in Table I, to calculate the odds ratios between models. In this scenario, a cumulative odds given multiple detections can be easily produced (see, e.g., Ref. [77]).

Ultimately, statistical inference consists of the continual update of one's current (prior) knowledge in the face of new observational evidence. The dependence of one's beliefs on prior knowledge should be viewed as a strength rather than a weakness as we approach new observations, since priors allow us to account for all past evidence on the subject.

We thank Thomas Dent, Jim Fuller, Chris Pankow, Harald Pfeiffer, Vivien Raymond, and Leo Stein for various stimulating discussions. We thank Christopher Berry, James Clark, Will Farr, Heather Fong, Prayush Kumar, Richard O'Shaughnessy, Harald Pfeiffer, and John Veitch for discussions on the use of the volumetric prior and making their implementation public. We thank Lisa Barsotti for providing a simulated noise spectral density for the analysis of the GW170104 software replica. S. V. acknowledges support of MIT through the Solomon Buchsbaum Research Fund, the National Science Foundation, and the LIGO Laboratory. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. 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. The authors acknowledge the LIGO Data Grid clusters. This project started during the workshop "Strong Gravity and Binary Dynamics with Gravitational Wave Observations," funded through NSF CAREER Grant No. PHY-1055103. This paper carries LIGO Document No. P1700176.

*salvatore.vitale@ligo.org
†dgerosa@caltech.edu
‡haster@cita.utoronto.ca
§kchatziioannou@cita.utoronto.ca
‖azimmer@cita.utoronto.ca

[1] T. Bayes, Phil. Trans. R. Soc. London **53**, 370 (1763).
[2] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, England, 2003).
[3] G. D'Agostini, arXiv:hep-ph/9512295.
[4] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. Lett. **116**, 061102 (2016).
[5] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. Lett. **116**, 241103 (2016).
[6] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaboration), Phys. Rev. Lett. **118**, 221101 (2017).
[7] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. X **6**, 041015 (2016).
[8] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, S. Vitale *et al.*, Phys. Rev. D **91**, 042003 (2015).
[9] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. Lett. **116**, 241102 (2016).
[10] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaboration), Astrophys. J. Lett. **833**, L1 (2016).
[11] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **116**, 221101 (2016).
[12] P. B. Graff, A. Buonanno, and B. S. Sathyaprakash, Phys. Rev. D **92**, 022002 (2015).
[13] J. Veitch, M. Pürrer, and I. Mandel, Phys. Rev. Lett. **115**, 141101 (2015).
[14] C.-J. Haster, Z. Wang, C. P. L. Berry, S. Stevenson, J. Veitch, and I. Mandel, Mon. Not. R. Astron. Soc. **457**, 4499 (2016).
[15] C. Cutler and É. E. Flanagan, Phys. Rev. D **49**, 2658 (1994).
[16] S. Vitale, R. Lynch, V. Raymond, R. Sturani, J. Veitch, and P. Graff, Phys. Rev. D **95**, 064053 (2017).
[17] S. Vitale, R. Lynch, J. Veitch, V. Raymond, and R. Sturani, Phys. Rev. Lett. **112**, 251101 (2014).
[18] H. Fong, H. Pfeiffer, P. Kumar, and W. Farr (private communication).
[19] J. Veitch and R. O'Shaughnessy (private communication).
[20] K. Belczynski, D. E. Holz, T. Bulik, and R. O'Shaughnessy, Nature (London) **534**, 512 (2016).
[21] C. L. Rodriguez, C.-J. Haster, S. Chatterjee, V. Kalogera, and F. A. Rasio, Astrophys. J. Lett. **824**, L8 (2016).
[22] D. Gerosa, M. Kesden, E. Berti, R. O'Shaughnessy, and U. Sperhake, Phys. Rev. D **87**, 104028 (2013).
[23] R. O'Shaughnessy, D. Gerosa, and D. Wysocki, Phys. Rev. Lett. **119**, 011101 (2017).
[24] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, Astrophys. J. Lett. **832**, L2 (2016).
[25] D. Kushnir, M. Zaldarriaga, J. A. Kollmeier, and R. Waldman, Mon. Not. R. Astron. Soc. **462**, 844 (2016).
[26] A. Nishizawa, A. Sesana, E. Berti, and A. Klein, Mon. Not. R. Astron. Soc. **465**, 4375 (2017).
[27] S. Vitale, R. Lynch, R. Sturani, and P. Graff, Classical Quantum Gravity **34**, 03LT01 (2017).
[28] S. Stevenson, C. P. L. Berry, and I. Mandel, Mon. Not. R. Astron. Soc. **471**, 2801 (2017).
[29] C. Talbot and E. Thrane, Phys. Rev. D **96**, 023012 (2017).
[30] W. M. Farr, S. Stevenson, M. C. Miller, I. Mandel, B. Farr, and A. Vecchio, Nature (London) **548**, 426 (2017).
[31] D. Gerosa, R. O'Shaughnessy, M. Kesden, E. Berti, and U. Sperhake, Phys. Rev. D **89**, 124025 (2014).
[32] D. Gerosa and E. Berti, Phys. Rev. D **95**, 124046 (2017).
[33] M. Fishbach, D. E. Holz, and B. Farr, Astrophys. J. Lett. **840**, L24 (2017).
[34] I. Mandel, W. M. Farr, A. Colonna, S. Stevenson, P. Tiňo, and J. Veitch, Mon. Not. R. Astron. Soc. **465**, 3254 (2017).

[35] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. X **6**, 041014 (2016).
[36] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. D **93**, 122003 (2016).
[37] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Phys. Rev. D **49**, 6274 (1994).
[38] L. E. Kidder, Phys. Rev. D **52**, 821 (1995).
[39] D. Gerosa, M. Kesden, U. Sperhake, E. Berti, and R. O'Shaughnessy, Phys. Rev. D **92**, 064016 (2015).
[40] J. D. Schnittman, Phys. Rev. D **70**, 124020 (2004).
[41] T. Bogdanović, C. S. Reynolds, and M. C. Miller, Astrophys. J. Lett. **661**, L147 (2007).
[42] M. C. Miller and J. M. Miller, Phys. Rep. **548**, 1 (2015).
[43] H. C. Spruit, Astron. Astrophys. **381**, 923 (2002).
[44] J. Fuller, M. Cantiello, D. Lecoanet, and E. Quataert, Astrophys. J. **810**, 101 (2015).
[45] T. Chiba and S. Yokoyama, Prog. Theor. Exp. Phys. (**2017**) 083E01.
[46] $\chi_{\rm eff}$ is a constant of motion at 2PN [47] and is therefore largely unaffected by the aforementioned issue on the reference frequency.
[47] É. Racine, Phys. Rev. D **78**, 044021 (2008).
[48] R. Smith, S. E. Field, K. Blackburn, C.-J. Haster, M. Pürrer, V. Raymond, and P. Schmidt, Phys. Rev. D **94**, 044031 (2016).
[49] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Phys. Rev. Lett. **113**, 151101 (2014).
[50] S. Vitale and M. Evans, Phys. Rev. D **95**, 064052 (2017).
[51] https://losc.ligo.org.
[52] T. B. Littenberg and N. J. Cornish, Phys. Rev. D **91**, 084034 (2015).
[53] N. J. Cornish and T. B. Littenberg, Classical Quantum Gravity **32**, 135012 (2015).
[54] This is essentially the same prior used in [4–7,9,35] apart from the M_c, q, and χ_i limitations required by the reduced-order quadrature implementation.
[55] M. Zaldarriaga, D. Kushnir, and J. A. Kollmeier, arXiv: 1702.00885.
[56] K. Hotokezaka and T. Piran, Astrophys. J. **842**, 111 (2017).
[57] K. Belczynski, J. Klencki, G. Meynet, C. L. Fryer, D. A. Brown, M. Chruslinska, W. Gladysz, R. O'Shaughnessy, T. Bulik, E. Berti *et al.*, arXiv:1706.07053.
[58] P. Kroupa, Mon. Not. R. Astron. Soc. **322**, 231 (2001).
[59] N. Bastian, K. R. Covey, and M. R. Meyer, Annu. Rev. Astron. Astrophys. **48**, 339 (2010).
[60] C. L. Rodriguez, S. Chatterjee, and F. A. Rasio, Phys. Rev. D **93**, 084029 (2016).
[61] S. Kullback and R. A. Leibler, Ann. Math. Stat. **22**, 79 (1951).
[62] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaboration), Living Rev. Relativity **19**, 1 (2016).
[63] We notice that the P_1 posteriors on LVT151012 are wider than those presented in Ref. [7]. This is due to differences in the way the power spectral density was estimated compared to this analysis.
[64] L. Kreidberg, C. D. Bailyn, W. M. Farr, and V. Kalogera, Astrophys. J. **757**, 36 (2012).
[65] K. Belczynski, V. Kalogera, and T. Bulik, Astrophys. J. **572**, 407 (2002).
[66] M. Hannam, D. A. Brown, S. Fairhurst, C. L. Fryer, and I. W. Harry, Astrophys. J. Lett. **766**, L14 (2013).
[67] K. Chatziioannou, N. Cornish, A. Klein, and N. Yunes, Astrophys. J. Lett. **798**, L17 (2015).
[68] K. Chatziioannou, N. Cornish, A. Klein, and N. Yunes, Phys. Rev. D **89**, 104023 (2014).
[69] I. Mandel, C.-J. Haster, M. Dominik, and K. Belczynski, Mon. Not. R. Astron. Soc. **450**, L85 (2015).
[70] T. B. Littenberg, B. Farr, S. Coughlin, V. Kalogera, and D. E. Holz, Astrophys. J. Lett. **807**, L24 (2015).
[71] B. Farr, C. P. L. Berry, W. M. Farr, C.-J. Haster, H. Middleton *et al.*, Astrophys. J. **825**, 116 (2016).
[72] K. Chatziioannou, C.-J. Haster, S. Vitale, A. Zimmerman, and D. Gerosa (to be published).
[73] A. R. Williamson, J. Lange, R. O'Shaughnessy, J. A. Clark, P. Kumar, J. Calderón Bustillo, and J. Veitch, arXiv: 1709.03095.
[74] D. W. Hogg, A. D. Myers, and J. Bovy, Astrophys. J. **725**, 2166 (2010).
[75] D. Wysocki, D. Gerosa, R. O'Shaughnessy, K. Belczynski, W. Gladysz, E. Berti, M. Kesden, and D. Holz, arXiv: 1709.01943.
[76] M. Zevin, C. Pankow, C. L. Rodriguez, L. Sampson, E. Chase, V. Kalogera, and F. A. Rasio, Astrophys. J. **846**, 82 (2017).
[77] M. Agathos, W. Del Pozzo, T. G. F. Li, C. Van Den Broeck, J. Veitch, and S. Vitale, Phys. Rev. D **89**, 082001 (2014).