

This is the accepted manuscript made available via CHORUS. The article has been published as:

Multiwaveform inference of gravitational waves

Gregory Ashton and Sebastian Khan

Phys. Rev. D **101**, 064037 — Published 18 March 2020

DOI: [10.1103/PhysRevD.101.064037](https://doi.org/10.1103/PhysRevD.101.064037)

Multi-waveform inference of gravitational waves

Gregory Ashton^{1,2} and Sebastian Khan^{3,4,5}

¹*School of Physics and Astronomy, Monash University, Vic 3800, Australia*

²*OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Clayton VIC 3800, Australia**

³*School of Physics and Astronomy, Cardiff University, Queens Buildings, Cardiff, CF24 3AA, United Kingdom†*

⁴*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstr. 38, 30167 Hannover, Germany*

⁵*Leibniz Universität Hannover, D-30167 Hannover, Germany*

Bayesian inference of gravitational wave signals is subject to systematic error due to modelling uncertainty in waveform signal models, coined approximants. A growing collection of approximants are available which use different approaches and make different assumptions to ease the process of model development. We provide a method to marginalize over the uncertainty in a set of waveform approximants by constructing a mixture-model multi-waveform likelihood. This method fits into existing workflows by determining the mixture parameters from the per-waveform evidences, enabling the production of marginalized combined sample sets from independent runs.

I. INTRODUCTION

Numerical relativity simulations of binary black hole mergers solve the full Einstein equations numerically and thus provide the most accurate predictions for the gravitational wave signal from compact binary coalescence (CBC) events [1–4]. These simulations are computationally demanding and the requirement by stochastic parameter estimation methods (see, e.g., [5–10]) to rapidly generate the waveform at an arbitrary point within the prior space makes their direct use impractical, except in grid-based methods [9–11] where the simulations can be pre-computed. To remedy this, a growing collection of rapidly computable *waveform approximants* for CBC signals have been developed, [12–33], some of which are tuned to the numerical relativity simulations.

Typically, inference workflows proceed by first identifying a set of waveforms relevant to the expected signal based on the signal characteristics identified by the search pipelines (see Ref. [34] for an overview of the search process). Then, inference is run for each waveform resulting in a set of *posterior samples* [5]. Differences between the inferred posteriors for each waveforms are understood to be due to the systematic differences in the waveform approximants; to create a set of results which are robust to these systematic waveform uncertainty, the naive-mixing method (used in, e.g., [34, 35]) is to combine equal numbers of samples from the posterior of each waveform into a single combined data set [36].

The choice to combine equal numbers of samples from multiple waveforms constitutes an equal-weighted probability on the waveform approximations. In the absence of additional information, this may appear to be the only choice. However, there exists additional information in the quality of the waveform fit to the data: intuitively the idea presented in this work is to weight the samples by the computed posterior evidence. In Sec. II, by treating

the set of approximants as a mixture model, we show how the fit of the waveforms themselves to the data can be used to infer the appropriate mixing fraction and combine samples. In Sec. III, we discuss the effect of uncertainty on the evidence estimates and in Sec. IV we provide a toy-model to help build intuition. We demonstrate that this method reduces waveform uncertainty by running an injection and recovery simulation in Sec. V and apply the method to GW150914 [37] in Sec. VI. We conclude with a discussion in Sec. VII.

II. METHOD

Given a set of N waveforms $\{w_\ell\}$ with equivalently defined model parameters θ , our goal is to compute $P(\theta|d, \{w_\ell\})$, the posterior distribution conditional on both the data d and the set of waveforms. First, let us associate to each waveform a hypothesis H_ℓ that the data was generated with the ℓ th waveform; the hypothesis includes prior-choices for the model parameters θ .

To obtain the likelihood for some data d , we assume that the hypotheses are exhaustive such that $1 = \sum_\ell P(H_\ell|d, \theta, \{\xi\})$ where $\{\xi\}$ are a set of prior-probability hyperparameters for each hypothesis, $\xi_\ell \equiv P(H_\ell|\{\xi\})$. For unitarity, we require $\xi_N = 1 - \sum_{\ell=0}^{N-1} \xi_\ell$. The likelihood can now be written as a mixture model with mixing parameters ξ_ℓ :

$$P(d|\theta, \{\xi\}) = \sum_\ell P(d|H_\ell, \theta) \xi_\ell. \quad (1)$$

This multi-waveform likelihood can be used in place of the usual likelihood (see, e.g. Veitch *et al.* [5]) to perform multi-waveform inference. (Note that, if used in practise, computing the likelihood for each waveform serially will slow down the per-likelihood compute time; the computation of $P(d|H_\ell, \theta)$ should instead be parallelised to reduce the overall compute time).

While the multi-waveform likelihood is simple to implement, it does not fit into the typical existing workflows described above. The solution is to first run inference for

* gregory.ashton@ligo.org

† sebastian.khan@ligo.org

each waveform independently, then infer the posterior mixing fraction

$$\hat{\xi}_\ell \equiv P(H_\ell|d) = \int d\theta P(H_\ell, \theta|d), \quad (2)$$

where we use a “hat” to distinguish $\hat{\xi}_\ell$ as the posterior mixing fraction for the H_ℓ model. The set $\{\hat{\xi}\}$ sum to unity and can be used to determine the mixing weights which should be applied to posterior samples $P(\theta|H_\ell, d)$. If an equal-weighted prior probability is assigned to each waveform, then

$$\hat{\xi}_\ell = \frac{Z_\ell}{\sum_j Z_j}, \quad (3)$$

where

$$Z_\ell \equiv P(d|H_\ell) = \int d\theta P(d|H_\ell, \theta)P(\theta|H_\ell), \quad (4)$$

is the per-waveform evidence.

For each individual waveform, we run inference and produce posterior samples. Then, instead of combining the samples equally, we combine them with weights given by Eq. (3). This yields a combined set of samples appropriately marginalized over the set of input waveforms. It is worth stating that this is not the same as marginalizing over waveform uncertainty in general: only uncertainty conditional on the input set of waveforms is captured.

This method of combining samples can be used when different waveform-hypotheses imply different priors on the model parameters. This can be seen in Eq. (4), where the model-parameter prior, $P(\theta|H_\ell)$, is conditional on the waveform-hypothesis H_ℓ . Because of this feature, samples from seemingly different waveform types can be combined, provided they refer to the same set of underlying model parameters, but with a differing prior. For example, if w_1 is a waveform including the tidal deformability parameters, λ_1, λ_2 , this can be combined with samples from w_2 , if w_1 and w_2 are equivalent when the tidal deformability parameters tend to zero. In this case, the prior on the tidal deformability parameters for w_2 are Dirac delta functions with peaks at zero.

III. UNCERTAIN MIXING FRACTIONS

The method outlined in Sec. II assumes that the mixing fractions can be calculated exactly. In practise, for gravitational wave signals, we do not have a closed form expression for Eq. (4). Instead, we estimate the evidence and posterior distribution using stochastic sampling methods [5–10]. These typically yield an estimated log-evidence with some uncertainty (usually expressed as an uncertainty on the log-evidence). For a discussion on how this is derived for the *dynesty* sampler used in this work, see Speagle [38].




FIG. 1. A comparison of the mixed posterior for an arbitrary parameter θ from two distributions when neglecting the evidence uncertainty, “certain mix”, and when including it, “uncertain mix”, using the procedure outlines in Sec. III.

If the log-uncertainty is sufficiently small with respect to the differences between log-evidences, it can of course be neglected. However, in cases where this is not true, care must be taken: ignoring the log-evidence uncertainty will result in an overly constrained posterior.

To include the uncertainty on the log-evidence into the mixing process, first we must define the distribution of evidences. For the *dynesty* sampler, we show in Appendix A that $\log(Z) \sim \text{Normal}(\log Z', \sigma_{\log(Z)})$, i.e. the log-evidence is normally distributed with mean given by the estimated log-evidence $\log Z'$ and standard deviation given by the estimated uncertainty $\sigma_{\log(Z)}$.

With an appropriate parametric form of the distribution of evidences, we then combine samples by the following process. (a) Draw a set of evidences for each waveform from their estimated distribution, (b) calculate the weights from Eq. (3), (c) apply the weights when drawing a sample from the combined set of samples. This process is then repeated until a sufficient number of samples are drawn for the mixed posterior. Removing step (a) of course reduces the operation to that defined in Sec. II and is appropriate if the uncertainty on the evidences is sufficiently small.

What effect does uncertainty in the evidence calculations have on the mixed posterior? To investigate this question, we perform a simple numerical experiment: we draw two sets of posterior samples from $P(\theta|A) \sim \text{Normal}(-1, 1)$ and $P(\theta|B) \sim \text{Normal}(1, 1)$ where θ is an arbitrary unknown variable which we wish to estimate and A and B arbitrarily label the posterior from two “waveforms”. We then defined that $\log(Z_A/Z_B) = 1$, but that each the log-evidences themselves have a normally distributed uncertainty with standard deviation $\sigma_{\log(Z)} = 5$. Following the procedure outlined above, we mix the posteriors. The results are given in Fig. 1 for both an “uncertain mix” where we include the uncertainty on the evidences and “certain mix” where we neglect that uncertainty.

Fig. 1 illustrates that including the additional uncer-

tainty on the evidences produces a more conservative estimate. The uncertain mix posterior is less skewed to the B posterior than the certain mix. We note that repeating the numerical experiment, but in a case where $\log(Z_A/Z_B) = 0$, i.e. the Bayes factor between the two posteriors is equally weighted, we find that the uncertain mixture and the certain mixture are indistinguishable for any amount of uncertainty in the evidences.

This numerical study suggests it is prudent to include the effects of uncertainty in the evidence estimates using the procedure outlined above. This adds little computational complexity, provided it is easy to sample from the distribution of evidences.

IV. TOY MODEL

To build intuition about the method, we describe here a simple toy model consisting of a sinusoidal function with a linearly-varying angular frequency

$$y(t) = \sin\left(\omega t + \frac{1}{2}\dot{\omega}t^2\right). \quad (5)$$

We then define two “waveforms” consisting of a choice for the rate of change of angular frequency:

$$w_A \rightarrow \dot{\omega} = 0 \quad (6)$$

$$w_B \rightarrow \dot{\omega} = 0.1 \quad (7)$$

We simulate data consisting of waveform w_B with Gaussian noise of known standard deviation $\sigma = 0.01$. We then then apply inference for w_A and w_B separately, producing two disjoint posteriors (see Fig. 2). The goal of this work is to present a method for combining samples between waveforms. Given the binary choice between the two toy-model waveforms, we have two options to do this. We can mix the posteriors using the posterior mixing fractions, Eq. (3), or we can apply the likelihood of Eq. (1) directly and perform inference on the mixture model itself.

For the simulated data with $\sigma = 0.01$, the Bayes factor between the two waveforms is $Z_B/Z_A \approx 5$: indicating a preference for w_B , but not overwhelmingly so. In Fig. 2, we show the posterior distribution on the only unknown model parameter ω for four cases (see caption). That the inference when using only w_A is biased is expected since the data was simulated using w_B . This case was specifically chosen to illustrate the behaviour when the data is not sufficiently informative to rule out the erroneous waveform, w_B .

The results mixing the posteriors according to Eq. (3) and applying the likelihood Eq. (1), to within the sampling errors, demonstrate equivalent posteriors. This validates that the mixing process is equivalent to direct inference.

The combined posterior in Fig. 2 (from either the mixing or direct methods) is multi-modal. This is a proper




FIG. 2. Posterior distribution on the angular frequency ω as a relative offset compared to the simulated value ω_s for each waveform separately, and the set of both waveforms. The term “mixed” refers to calculating the posterior mixing fraction using Eq. (3); we apply the procedure in Sec. III to include the numerical uncertainty on the estimated evidence. The term “direct” refers to applying Eq. (1) directly. In the title, we give the Bayes factor between the models and the numerical value of $\hat{\xi}$ neglecting the uncertainty.

reflection of the posterior uncertainty on the model parameter: each mode is inherently associated with a different model. This is happening in this special case because the evidences are not especially informative: the Bayes factor only demonstrates a mild preference for w_B .

To gain intuition when the evidences become more information, in Fig. 3, we repeat the toy model of Fig. 2, but decreasing the level of noise. In particular, we reduce the standard deviation σ from 0.01 to 0.009. With this modest decrease in the noise the Bayes factor is now four times more in favour of B than was the case for Fig. 2. Correspondingly, the mixing fraction increases and the posterior is now more weighted to the correct B mode. Further decreasing the noise, the A mode is eventually (for Bayes factor above ~ 50 or so) entirely ruled out. On the other hand, if we were to increase the level of noise, we would see a more equal mixing between the two since the evidence would favour neither one or the other.

V. INJECTION AND RECOVERY

To demonstrate the utility of this method, we now run a simple injection and recovery test. Aligned-spin signals generated by the IMRPhenomD [21, 39] waveform model are added to simulated coloured-Gaussian data from two detectors (Hanford and Livingston) with Advanced LIGO design sensitivity [40, 41]. The data is simulated and analysed using the Bilby [8] Bayesian inference software. The simulated source parameters are generated by random draws from the prior (see Table I); repeated draws are made until the network optimal signal to noise ratio (SNR) exceeds a threshold of 8, a typical




FIG. 3. Repeating the toy simulation of Fig. 2 (see caption for details), but with a smaller level of noise in the data: the reduction in noise leads to more informative evidence estimates and demonstrates the mixed posterior converging to the true value while neglecting the erroneous model A.

Parameter	Prior support	
Chirp mass \mathcal{M}	25 – 100	M_\odot
Mass ratio q	0.125 – 1	
Primary spin $\chi_{1,z}$	-0.9 – 0.9	
Secondary spin $\chi_{2,z}$	-0.9 – 0.9	
Lum. distance d_l	0.1 – 5	Gpc
Inclination θ_{JN}	0 – π	rad.
Right Asc. α	0 – 2π	rad.
Declination δ	$-\pi/2$ – $\pi/2$	rad.
Polarisation angle ψ	0 – π	rad.
Phase ϕ	0 – 2π	rad.
Geocentric-time t_c	-0.1 – 0.1	s

TABLE I. Prior support for the source parameters. $\chi_{1,z}$ and $\chi_{2,z}$ use the “z-prior” (see Eq. (A7) of Lange *et al.* [10]); the luminosity distance prior is $P(d_L) \propto d_L^2$; the inclination angle is Cosine distributed; the declination is Sine distributed; all other parameters are Uniformly distributed. The geocentric time, t_c is given relative to the simulated trigger time.

search threshold.

For each simulated signal, we recover with either IMRPhenomC [20] or IMRPhenomD. To probe the inherent bias, we repeat this process on 500 simulated data sets and perform a percentile-percentile (pp-test) (based on the work of Cook *et al.* [42]). Graphically a pp-test is a plot of the fraction of signals with true parameter recovered to within a credible interval against the credible interval itself (we show an example later in Fig. 4). The pp-test is a useful diagnostic for investigating bias: a pass in a pp-test verifies that the posterior recovery is unbiased with respect to the injections (i.e. the $x\%$ posterior intervals contain the true values $x\%$ of the time). For any pp-test, a summary statistic is obtained by first calculating a p -value from the Kolmogorov-Smirnov test on each parameter separately, then combining these together using Fisher’s method. The resulting set of p -values (for which the null hypothesis is that the results are unbi-




FIG. 4. We simulate 500 IMRPhenomD signals in advanced-LIGO design sensitivity noise [40]. Parameter estimation is performed on all simulated data sets using both the IMRPhenomD and IMRPhenomC waveforms and the resulting samples are mixed using the informed-mixture method, i.e. with mixture parameters calculated from Eq. (3). This figure shown the pp-test diagnostic plot applied to the 500 mixed posteriors.

ased) are given in Table II.

Injection	Recovery	p -value
IMRPhenomD	IMRPhenomD	0.45
IMRPhenomD	IMRPhenomC	< 0.01
IMRPhenomD	informed-mix	0.045
IMRPhenomD	naive-mix	0.0092

TABLE II. Table of combined p -values calculated for each injection-recovery test case. Combinations are made over the full set of parameters in Table I, except the geocentric time which has known systematic shifts between waveforms.

For the case when the injection and recovery are performed using the same waveform, as expected we find a p -value indicating the posteriors are unbiased: this demonstrates that in the absence of systematic differences in the injections and recovery waveforms, the underlying method (i.e. the generation of injection values and posterior sampling) is unbiased.

On the other hand, when the recovery waveform (IMRPhenomC) differs from the injection, the p -value is small, indicating the results are biased. Per-parameter analysis indicates that it is the merger time, mass ratio, and chirp mass parameters which fail the test. The cause for this, systematic differences between waveforms, is well understood and expected [43].

The true signal in the simulated data set is

IMRPhenomD. However, we now consider the case when we have uncertainty about which waveform best approximates the signal. Using the method described in Sec. II, the set of posterior samples conditional on both waveforms is obtained by mixing together samples from the IMRPhenomC- and IMRPhenomD-recovery with a mixing fraction given by the ratio of their evidence to the total evidence, Eq. (3). Repeating this for each simulated data segment, we apply the pp-test to the resulting samples, the pp-test plot itself is given in Fig. 4 and the combined p -value is labelled as “informed-mix” in Table II. The results are biased, but to a substantially lesser extent than for IMRPhenomC alone.

That the p -value indicates a bias for the informed-mixture biased is unsurprising since the pp-test is only expected to pass when the data-generation exactly matches the assumptions of the model-fitting software [42]; in this case we have introduced additional uncertainty into the model-fitting.

Mixing the samples according to Eq. (3) is the better thing to do, given uncertainty about the waveform. The Gravitational Wave Transient Catalogue [34], used the naive-mixing method, combining equal numbers of posterior samples from multiple waveforms. We implement this “naive-mix” method and apply it to the set of samples from each individual waveform. The resulting p -value is smaller than the informed-mix, indicative of a greater degree of bias. Nevertheless, it demonstrates that for cases where the waveforms make highly similar predictions (a statement that can be quantified by a Bayes factor between them), the naive-mixing method is reasonable, but should only be used when the evidence calculations are infeasible.

We have investigated here the typical case for advanced-era CBC detections in which both waveforms perform reasonably well in fitting the data (see the next Section for a demonstration). Had the set of injections (or sensitivity of the simulated instruments) been such that the differences in waveforms were more apparent, the informed mixing fraction would preference the injected waveforms and mix the posterior samples accordingly. As an example, consider the case where two waveforms (w_A and w_B) are applied and each produces a set of 10^4 posterior samples. For the probability of including any samples from w_B to be less than 1, $\hat{\xi} < 10^{-4}$, which implies the Bayes factor between them must be $Z_A/Z_B \lesssim 10^{-4}$. For such a case, the posterior samples would (almost) all be drawn from w_A . If repeated in a pp-test, the informed-mixture method would be unbiased, but the naive-mixture method would not.

VI. APPLICATION TO GW150914

To apply the method in practise, we run Bayesian inference on the first-observed binary black hole coalescence, GW150914 [35, 37, 41]. This system has been well studied and posterior samples are available

[34, 35], but the evidences are not. These original analyses used both precessing and non-precessing waveform approximants. Here, as an illustrative example of the effect of waveform-approximant mixing only, we perform analysis for three non-precessing waveform approximants, IMRPhenomC [20], IMRPhenomD [21, 39], and SEOBNRv4_ROM [24]. The analysis is done using Bilby [8] on data from the Gravitational Wave Open Science Centre [44][45] following the methodology described in Appendix B of Ref. [34].

The evidences computed for each waveform can be used to construct Bayes factors

$$\ln \left(\frac{Z_{\text{IMRPhenomD}}}{Z_{\text{SEOBNRv4_ROM}}} \right) = 0.02 \pm 0.5, \quad (8)$$

and

$$\ln \left(\frac{Z_{\text{IMRPhenomD}}}{Z_{\text{IMRPhenomC}}} \right) = 0.5 \pm 0.5. \quad (9)$$

These results confirm what is known in the literature (see, e.g., [35, 46]): GW150914 and other events seen in the first and second observing runs of LIGO and Virgo are not sufficiently loud to decisively distinguish waveform approximants.

The mixing fractions for these three approximants, applying, Eq. (3), are 0.24, 0.39, and 0.38 for the IMRPhenomC, IMRPhenomD, and SEOBNRv4_ROM waveforms respectively. To illustrate the effect of the mixing, in Fig. 5, we plot the posterior probability density for the detector-frame chirp mass from the three aligned-spin waveform approximants and the mixture. Of the three waveforms, IMRPhenomC has the smaller evidence and only a quarter of the samples are drawn from this posterior: as a result, the mixture is closer to the IMRPhenomD and SEOBNRv4_ROM posteriors.

Because the difference in evidences between the waveforms is small, as demonstrated by the Bayes factors in Eq. (8) and (9), the informed-mixing method will yield results similar to those produced by the naive-mixing method. With future detections, when the data is more informative about features in the waveform approximants, we strongly recommend that the informed-mixing method is applied when combining samples to ensure the results properly reflect the posterior uncertainty.

VII. CONCLUSION AND OUTLOOK

The informed-mixing method presented here provides an improvement on the naive-mixing method to combine samples from multiple waveforms by including information from the estimated evidences about how well the waveforms fit to the data. Ultimately, this should make it easier to include multiple waveforms without concern about the introduction of biases due to sub-optimal combination of posterior samples.




FIG. 5. The posterior probability density for the detector-frame chirp mass, \mathcal{M} , in solar masses for three aligned-spin waveform approximants and their mixture, applying Eq. (3) and the procedure for including uncertainty in the evidences from Sec. III. The same uniform prior on detector-frame chirp mass was applied for all three individual waveform approximants.

To use this optimal method for multi-waveform inference, accurate estimation must be made of the waveform evidence, Eq. (4). As such, the ability to properly handle systematic uncertainty in the waveform is critically underpinned by the ability estimate the evidence. We encourage future analyses of CBC systems to ensure evidence estimates are calculated and reported. In Sec. III, we outline a procedure to include the uncertainty on these evidence estimates into the mixing process.

The mixture-model method presented in this work was discussed in the context of multi-waveform inference. Another systematic uncertainty is in the estimate of the power spectral density (PSD) used to characterise the detector noise [47]. The state of the art method (used in The LIGO Scientific Collaboration and the Virgo Collaboration [34]) involves applying the *BayesLine* algorithm [48, 49] which computes a posterior probability distribution for possible PSDs, then using the median PSD in

PE analyses [50]. However, marginalizing over the uncertainty, rather than making a point estimate of the PSD is preferable. The methodology presented in Sec. II can be applied to this problem. If multiple runs are performed with differing draws from the *BayesLine* posterior, Eq. (3) can be applied to calculate mixing fractions with which to combine posteriors. This method of PSD marginalization is sub-optimal compared to the general method of fitting both the PSD and source model simultaneously [50], however it makes the problem of marginalizing over PSD uncertainty embarrassingly parallel.

ACKNOWLEDGMENTS

We thank Christopher Berry, John Veitch, Michael Pürrer, and members of the LIGO and Virgo Compact Binary Coalescence group for valuable input during the preparation of this manuscript. We also thank the anonymous referee for valuable comments which improved the manuscript during review. G.A. is supported by the Australian Research Council through grants CE170100004, FT150100281, and DP180103155. S.K. acknowledges support by the Max Planck Society’s Independent Research Group Grant. This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gwopenscience.org>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. All analyse performed with *Bilby* in this work makes use of the *dynesty* [38] nested-sampling package. The *scipy* [51] and *matplotlib* [52] packages are used for statistical computations and creating figures.

[1] F. Pretorius, *Phys. Rev. Lett.* **95**, 121101 (2005).
[2] M. Campanelli and others., *Phys. Rev. Lett.* **96**, 111101 (2006).
[3] J. G. Baker *et al.*, *Phys. Rev. Lett.* **96**, 111102 (2006).
[4] B. Brügmann *et al.*, *Phys. Rev. D* **77**, 024027 (2008).
[5] J. Veitch *et al.*, *Phys. Rev. D* **91**, 042003 (2015), [arXiv:1409.7215 \[gr-qc\]](https://arxiv.org/abs/1409.7215).
[6] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite,” free software (GPL) (2018).
[7] C. M. Biwer *et al.*, *Publications of the Astronomical Society of the Pacific* **131**, 024503 (2019), [arXiv:1807.10312 \[astro-ph.IM\]](https://arxiv.org/abs/1807.10312).
[8] G. Ashton *et al.*, *The Astrophysical Journal Supplement*

Series **241**, 27 (2019), [arXiv:1811.02042 \[astro-ph.IM\]](https://arxiv.org/abs/1811.02042).
[9] C. Pankow *et al.*, *Phys. Rev. D* **92**, 023002 (2015), [arXiv:1502.04370 \[gr-qc\]](https://arxiv.org/abs/1502.04370).
[10] J. Lange, R. O’Shaughnessy, and M. Rizzo, *arXiv e-prints*, [arXiv:1805.10457](https://arxiv.org/abs/1805.10457) (2018), [arXiv:1805.10457 \[gr-qc\]](https://arxiv.org/abs/1805.10457).
[11] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **94**, 064035 (2016).
[12] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999).
[13] Y. Pan *et al.*, *Phys. Rev. D* **89**, 084006 (2014), [arXiv:1307.6232 \[gr-qc\]](https://arxiv.org/abs/1307.6232).
[14] A. Buonanno and T. Damour, *Phys. Rev. D* **62**, 064015

- (2000).
- [15] Y. Pan *et al.*, *Phys. Rev. D* **84**, 124052 (2011), [arXiv:1106.1021 \[gr-qc\]](#).
 - [16] A. Taracchini *et al.*, *Phys. Rev. D* **89**, 061502 (2014), [arXiv:1311.2544 \[gr-qc\]](#).
 - [17] L. Blanchet, *Living Reviews in Relativity* **17**, 2 (2014).
 - [18] M. Hannam *et al.*, *Phys. Rev. Lett.* **113**, 151101 (2014), [arXiv:1308.3271 \[gr-qc\]](#).
 - [19] P. Ajith *et al.*, *Gravitational wave data analysis. Proceedings: 11th Workshop, GWDAAW-11, Potsdam, Germany, Dec 18-21, 2006*, *Class. Quant. Grav.* **24**, S689 (2007), [arXiv:0704.3764 \[gr-qc\]](#).
 - [20] L. Santamaría *et al.*, *Phys. Rev. D* **82**, 064016 (2010).
 - [21] S. Khan *et al.*, *Phys. Rev. D* **93**, 044007 (2016), [arXiv:1508.07253 \[gr-qc\]](#).
 - [22] L. London *et al.*, *Phys. Rev. Lett.* **120**, 161102 (2018).
 - [23] R. Cotesta *et al.*, *Phys. Rev. D* **98**, 084028 (2018).
 - [24] A. Bohé *et al.*, *Phys. Rev. D* **95**, 044028 (2017).
 - [25] S. Babak, A. Taracchini, and A. Buonanno, *Phys. Rev. D* **95**, 024010 (2017).
 - [26] J. Blackman *et al.*, *Phys. Rev. D* **96**, 024058 (2017).
 - [27] A. Nagar *et al.*, *Phys. Rev. D* **98**, 104052 (2018).
 - [28] A. Kumar Mehta *et al.*, *Phys. Rev. D* **100**, 024032 (2019), [arXiv:1902.02731 \[gr-qc\]](#).
 - [29] V. Varma *et al.*, *Phys. Rev. D* **99**, 064045 (2019), [arXiv:1812.07865 \[gr-qc\]](#).
 - [30] D. Williams *et al.*, *arXiv e-prints* (2019), [arXiv:1903.09204 \[gr-qc\]](#).
 - [31] Y. E. Setyawati, F. Ohme, and S. Khan, *Phys. Rev. D* **99**, 024010 (2019), [arXiv:1810.07060 \[gr-qc\]](#).
 - [32] S. Khan *et al.*, *Phys. Rev. D* **100**, 024059 (2019), [arXiv:1809.10113 \[gr-qc\]](#).
 - [33] S. Khan, F. Ohme, K. Chatziioannou, and M. Hannam, *arXiv e-prints*, [arXiv:1911.06050](#) (2019), [arXiv:1911.06050 \[gr-qc\]](#).
 - [34] The LIGO Scientific Collaboration and the Virgo Collaboration, *arXiv e-prints* (2018), [arXiv:1811.12907 \[astro-ph.HE\]](#).
 - [35] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **116**, 241102 (2016), [arXiv:1602.03840 \[gr-qc\]](#).
 - [36] <https://dcc.ligo.org/LIGO-T1500597>.
 - [37] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **116**, 061102 (2016), [arXiv:1602.03837 \[gr-qc\]](#).
 - [38] J. S. Speagle, *arXiv e-prints* (2019), [arXiv:1904.02180 \[astro-ph.IM\]](#).
 - [39] S. Husa *et al.*, *Phys. Rev. D* **93**, 044006 (2016).
 - [40] B. P. Abbott *et al.*, *Living Reviews in Relativity* **21**, 3 (2018), [arXiv:1304.0670 \[gr-qc\]](#).
 - [41] J. Aasi *et al.*, *Classical Quantum Gravity* **32**, 074001 (2015).
 - [42] S. R. Cook, A. Gelman, and D. B. Rubin, *Journal of Computational and Graphical Statistics* **15**, 675 (2006).
 - [43] S. Khan, *Numerical modelling of black-hole-binary mergers*, *Ph.D. thesis*, Cardiff University (2016).
 - [44] M. Vallisneri *et al.*, in *Journal of Physics Conference Series*, Journal of Physics Conference Series, Vol. 610 (2015) p. 012021, [arXiv:1410.4839 \[gr-qc\]](#).
 - [45] Gravitational Wave Open Science Center (GWOSC), www.gw-openscience.org/events/GW150914/, DOI:10.7935/K5MW2F23.
 - [46] E. Payne, C. Talbot, and E. Thrane, *arXiv e-prints*, [arXiv:1905.05477](#) (2019), [arXiv:1905.05477 \[astro-ph.IM\]](#).
 - [47] C. Rover, R. Meyer, and N. Christensen, *Class. Quant. Grav.* **28**, 015010 (2011), [arXiv:0804.3853 \[stat.ME\]](#).
 - [48] N. J. Cornish and T. B. Littenberg, *Classical and Quantum Gravity* **32**, 135012 (2015), [arXiv:1410.3835 \[gr-qc\]](#).
 - [49] T. B. Littenberg and N. J. Cornish, *Phys. Rev. D* **91**, 084034 (2015), [arXiv:1410.3852 \[gr-qc\]](#).
 - [50] K. Chatziioannou *et al.*, (2019), [arXiv:1907.06540 \[gr-qc\]](#).
 - [51] E. Jones, T. Oliphant, P. Peterson, *et al.*, “SciPy: Open source scientific tools for Python,” (2001–).
 - [52] J. D. Hunter, *Computing in Science & Engineering* **9**, 90 (2007).
 - [53] S. S. Shapiro and M. B. Wilk, *Biometrika* **52**, 591 (1965).

Appendix A: The distribution of uncertainties from dynesty

In this Appendix, we demonstrate that the distribution of log-evidence errors from the *dynesty* [38] nested sampling package are normally distributed and conservatively bounded by the estimated log-evidence uncertainty. We do this using a one-dimensional problem. We cannot ensure that the same behaviour will occur in higher-dimensional problems, that will require further testing. But, this investigation is designed to give some intuition in a low-dimensional case where repeated numerical calculations are not computationally prohibitive.

To test the sampler, consider the integral

$$Z = \int_{\mu_{\min}}^{\mu_{\max}} \frac{1}{\sqrt{2\pi}} \exp^{-\frac{\mu^2}{2}} \frac{1}{\mu_{\max} - \mu_{\min}} d\mu, \quad (\text{A1})$$

which is the evidence for a Gaussian likelihood with $x = 0$, unknown mean μ , but known variance $\sigma^2 = 1$, and a uniform prior on $[\mu_{\min}, \mu_{\max}]$. If the uniform prior is sufficiently wide with respect to the posterior support, the integral is simply the Gaussian integral and hence the evidence has the approximate closed-form solution

$$Z \approx \frac{1}{\mu_{\max} - \mu_{\min}}. \quad (\text{A2})$$

Alternatively, Eq. (A1) can be estimated using stochastic sampling software.

We use the *dynesty* sampler to compute the evidence and repeat the calculation 500 times whilst varying the number of live-points used by the sampler (a quantity proportional to the expected precision of the evidence estimate). In Tab. III, we report summary statistics for the distribution of

$$f = \log Z + \log(\mu_{\max} - \mu_{\min}), \quad (\text{A3})$$

the difference between the computed log-evidence and the logarithm of Eq. (A2), the approximate closed-form evidence.

That the mean values of Table III are small relative to the typical errors indicates that the evidences themselves are not substantially biased. That the first estimate of the error σ_0 is close to, but always greater than, the

live-points	$\mu(f)$	$\sigma(f)$	σ_0	p
100	0.0032	0.095	0.111	0.97
250	0.0045	0.061	0.073	0.14
500	0.00032	0.047	0.053	0.38

TABLE III. Table of the mean and standard deviation of f , (see, Eq. (A3)) along with σ_0 , the first estimate of the expected evidence error provided by **dynesty**, and the p -value obtained from applying the Shapiro-Wilk normality test [53].

empirically measured value $\sigma(f)$ suggests that the initial reported uncertainty on the evidences are conservative upper estimates. Finally, that the normality test p -values are all greater than a standard threshold of 0.05 indicates that the distribution of f , and hence of $\log Z$ itself is consistent with a normal distribution.