

Digging the population of compact binary mergers out of the noise

Sebastian M. Gaebel

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

John Veitch

*Institute for Gravitational Research, School of Physics and Astronomy,
University of Glasgow, Glasgow, G12 8QQ, United Kingdom*

Thomas Dent

Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-30167 Hannover, Germany

Will M. Farr

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

Center for Computational Astrophysics, Flatiron Institute,

162 Fifth Avenue, New York NY 10010, USA and

Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY, 11794, USA

(Dated: September 12, 2018)

Coalescing compact binaries emitting gravitational wave (GW) signals, as recently detected by the Advanced LIGO-Virgo network, constitute a population over the multi-dimensional space of component masses and spins, redshift, and other parameters. Characterizing this population is a major goal of GW observations and may be approached via parametric models. We demonstrate hierarchical inference for such models with a method that accounts for uncertainties in each binary merger's individual parameters, for mass-dependent selection effects, and also for the presence of a second population of candidate events caused by detector noise. Thus, the method is robust to potential biases from a contaminated sample and allows us to extract information from events that have a relatively small probability of astrophysical origin.

I. INTRODUCTION

Since the first detection of gravitational waves in 2015 [1], the Advanced LIGO and Advanced Virgo detectors have observed the coalescence of multiple compact binary systems, and have begun to reveal the population of coalescing compact objects [2]. This population is enabling studies in fields from probing alternative theories of gravity to constraining models of stellar evolution. These tend to be interested either in individual, preferably loud, signals, or in the population of sources as a whole. The latter type of population analysis tries to estimate the parameters governing the distribution of sources in the universe, their masses and spins, and the value of the astrophysical merger rate [3] are of particular interest. As the sensitivity of detectors improves over the coming years, the detected number of sources is expected to grow at an accelerated pace, rapidly increasing the amount of information available for population studies [4].

When undertaking population analyses, one has to consider that real detectors may produce noise transients which cannot in all cases be distinguished from astrophysical GW sources. To avoid population inferences being biased by noise events one might consider only events with a much higher probability to be of astrophysical origin than to be caused by noise artefacts. In templated searches for compact binaries, the relative probability of astrophysical vs. noise origin for a candidate event is a function of a detection statistic calculated for each event by a search analysis pipeline (see e.g. [5–10]). Typically a candidate event is generated as a local maximum in matched filter signal-to-noise ratio (SNR) above a search threshold; the detection statistic value then incorporates the matched filter SNR, as well as other goodness-of-fit tests to reject non-Gaussian instrumental noise transients [11, 12].

At low SNR, the population of events is dominated by the noise ‘background’, whereas at high SNR (or in general for events assigned high statistic values) the astrophysical ‘foreground’ dominates.¹ To limit possible pollution of the sample used for population inference, one may place a minimum threshold on the detection statistic; any event above threshold is then assumed to be astrophysical, whereas all other events are discarded as potential noise transients. Note that the choice of threshold value requires an empirical estimate of the rate and distribution of background events [13], since the rate, strength and morphologies of detector noise artefacts are not known *a priori* [14].

A simple strategy of thresholding is sub-optimal for two reasons. First, discarding events below the threshold will almost certainly discard information from some number of quiet but still identifiable signals; second, there is still a finite chance that the resulting ‘signal’ set is nevertheless contaminated by noise, leading to potentially biased inferences. The choice of SNR threshold requires a trade-off between these two considerations and depends on the intended use. One also has to take into consideration any bias in the observed population produced by the effect of source parameters on the loudness of a signal, and thus its chance of exceeding a SNR threshold [15]; we expect potentially major observation selection effects for binary mass(es) [16] and component spins [17].

Here, we propose a method which alleviates the issues associated with simple thresholding by applying a hierarchical mixture model under which each event is considered to originate from either a foreground or a background (noise) population. For each event, the probability of either case naturally defines a weight for its contribution to inferences on population parameters.

This method combines the processes of estimating the expected number of events in either class [18] (the number of foreground events being a proxy for the astrophysical merger rate density) and estimating parameters of the underlying populations, which have previously been performed separately. It avoids being biased through the inclusion of background events, while being able to make use of events with a non-negligible probability of noise origin, which would be discarded by thresholding. In theory this method allows the SNR threshold to be reduced to an arbitrarily low value, though in practice we are still limited by the computational resources required to extract the source parameters from each event under consideration.²

Our method is applicable to any hierarchical model of a source population, examples of which have been explored in the literature. This includes analyses which combine information from multiple events to infer a parameter common to all, such as deviations from general relativity [20] or a parameterised neutron star equation of state [21]. The use of a mixture model with astrophysical and noise populations is particularly useful when the model of interest has a strong effect on the detectability of sources, i.e. the detected events are unrepresentative of the underlying population. A good example is the mass distribution of sources, which we consider below.

Messenger and Veitch [22] considered the problem of selection effects in mass distribution inference by dividing the observing time into discrete chunks, which each contain zero or one sources, and computing population likelihoods while accounting for false alarms and false dismissals from an idealised noise distribution. Farr *et al.* [18] derived

¹ We will loosely refer to the detection statistic as ‘SNR’ when discussing the distinction between instrumental noise events and astrophysical events.

² An analysis that effectively removes all SNR thresholds, applying Bayesian analysis to the entirety of the gravitational-wave data set rather than restricting to data close to events triggered on SNR maxima, is proposed in [19]; its application appears at present to be still more limited by computational cost.

an equivalent formalism for rate inference that allows a population shape function to be estimated alongside. Our derivation in Section II follows similar lines.

The selection function for masses was important in estimating the astrophysical event rates in the first Advanced LIGO observing run (O1), which inferred rates using a mixture model, for fixed choices of population shape (i.e. mass distribution) [3, 23]. A separate analysis also estimated the slope of a power law model of the mass distribution function considering the detected events, described in [2] (and updated in [24]).

The selection function in the form of a sensitivity-weighted measure of the space-time volume VT is also important when considering searches which do not make a clear detection. There, the loudest background event, or a nominal detection threshold, is used to set an upper limit on the astrophysical rate of a fiducial source population, for example limits on the rate of mergers of binary neutron stars and neutron star–black hole binaries in O1 [25]. The sensitive-volume approach has been in use since the initial detector era [26, 27], and continues to be refined to incorporate mass- and spin-dependent selection effects [17, 28, 29] and cosmological effects, as well as to improve the accuracy of measurement [16].

As the number of detections increases, determination of the population of coalescing compact binaries is expected to provide insight into the astrophysics of black hole and neutron star binary formation, as reviewed in [30–32]). Population synthesis models can describe the masses and spins of coalescing compact binaries under a variety of formation scenarios (see e.g. [33–36]). Comparison of these predictions to the observed distribution can be used to constrain the uncertainties in parametrised models of source populations [37, 38]. This has motivated the development of methods to determine the mass-dependent coalescence rate in the absence of false alarms, using both specific parameterised models [39–42] and non-parametric methods [43]. The alignment of black hole spins is expected to be a key distinguishing feature between binaries formed in the field or through dynamical interactions (see e.g. [44–52]), which also requires an understanding of the spin selection function [17, 28, 29]. The work presented here is complementary to these studies, as it aims to incorporate an astrophysical distribution model as part of a mixture with a noise component. As the observed population is limited by the sensitivity of Advanced ground-based detectors, the population of candidate sources at the greatest distances (lowest detection significance) will be contaminated with background events. We expect our method, using information from such sources, to improve both the precision and accuracy of merger rate and population parameter estimates; though as we will see, the degree of improvement depends on how easily the foreground and background populations can be separated by existing analyses.

We start by defining our notation and deriving the general form of the model in section II. Section III describes its application to a toy model of mass distribution inference in the presence of noise, and shows its application to a range of simple analytic population models. In section V we consider a more realistic simulated data set derived from an engineering run prior to the start of Advanced LIGO observations in 2015. We conclude in section VI.

II. DERIVATION OF THE GENERIC MODEL

We consider a mixture of two populations, the astrophysical ‘foreground’ and terrestrial noise ‘background’: quantities defined analogously for both populations will be distinguished by the subscripts F or B respectively. Quantities without subscript will then refer to the total population which is the union of foreground and background.

The model is also hierarchical: each event, if assumed astrophysical, has a set of intrinsic properties such as component masses and spins, which we collectively denote γ . The distribution of these properties over each population is assumed to have a form described by a set of hyper-parameters. We do not have access to the ‘true’ values of properties for each event, only to a set of samples from a (typically Bayesian) estimate based on data around the event. These samples are derived under the assumption that the event is astrophysical, thus events that are in fact background will also be assigned parameter estimates.³

We then define the core quantities used in the following derivation as

- $\rho_i, \{\rho\}$: ranking statistic for one, resp. for all events in a given data set.
- $N_{\text{obs}}, N_{\text{F,obs}}, N_{\text{B,obs}}$: observed number of events above a threshold $\rho_i > \rho_{\text{thr}}$
- $N_{\text{exp}}, N_{\text{F,exp}}, N_{\text{B,exp}}$: expected number of events with $\rho_i > \rho_{\text{thr}}$, when modelling these as a Poisson process
- $\theta_{\text{F}}, \theta_{\text{B}}$: hyper-parameters which describe the population distributions
- $\eta_i, \{\eta\}$: flag indicating whether any given event, resp. all events, belong(s) to the astrophysical ($\eta = F$) or to the noise population ($\eta = B$)

³ We do not, of course, know with certainty that any given event is background.

- $\vec{\gamma}_i, \{\vec{\gamma}\}$: vector of samples representing the parameter estimates (masses, spins, etc.) of one, resp. all events, under the assumption that events are astrophysical.

We wish to infer the joint posterior probability distribution of rates and population parameters for the two populations, given some events for which $\{\rho\}$ and $\{\vec{\gamma}\}$ have been determined by the search and parameter estimation stages of data analysis:

$$p(N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}} | \{\rho\}, \{\vec{\gamma}\}, N_{\text{obs}}). \quad (1)$$

Using Bayes' theorem we can express the posterior distribution (1) in terms of prior and likelihood functions,

$$p(N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}} | \{\rho\}, \{\vec{\gamma}\}, N_{\text{obs}}) = \frac{p(\{\rho\}, \{\vec{\gamma}\}, N_{\text{obs}} | N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}}) p(N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}})}{p(\{\rho\}, \{\vec{\gamma}\}, N_{\text{obs}})} \quad (2)$$

We drop the normalisation constant $p(\{\rho\}, \{\vec{\gamma}\}, N_{\text{obs}})$ and factor out the likelihood for N_{obs} as being independent of the population shape parameters,

$$\begin{aligned} p(\{\rho\}, \{\vec{\gamma}\}, N_{\text{obs}} | N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}}) &= p(N_{\text{obs}} | N_{\text{F,exp}}, N_{\text{B,exp}}) p(\{\rho\}, \{\vec{\gamma}\} | N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}}) \\ &= \frac{N_{\text{exp}}^{N_{\text{obs}}} e^{-N_{\text{exp}}}}{N_{\text{obs}}!} p(\{\rho\}, \{\vec{\gamma}\} | N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}}), \end{aligned} \quad (3)$$

where we use a Poisson likelihood for N_{obs} with a total expected number of events $N_{\text{exp}} = N_{\text{F,exp}} + N_{\text{B,exp}}$. The second term, $p(\{\rho\}, \{\vec{\gamma}\} | N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}})$, is the likelihood for the observed SNRs and parameter estimates, for the mixture model. We assume each event is conditionally independent given the population parameters, and so the joint likelihood is just the product of the likelihood for each one,

$$p(\{\rho\}, \{\vec{\gamma}\} | N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}}) = \prod_i p(\rho_i, \vec{\gamma}_i | N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}}). \quad (4)$$

Now, we can split each of these into terms for the astrophysical and noise sub-models by introducing an indicator variable $\eta_i \in (F, B)$, whose probability will depend on the rate parameters $N_{\text{F,exp}}$ and $N_{\text{B,exp}}$,

$$\begin{aligned} p(\rho_i, \vec{\gamma}_i | N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}}) &= p(\rho_i, \vec{\gamma}_i | \theta_{\text{F}}, \eta_i = \text{F}) p(\eta_i = \text{F} | N_{\text{F,exp}}, N_{\text{B,exp}}) \\ &\quad + p(\rho_i, \vec{\gamma}_i | \theta_{\text{B}}, \eta_i = \text{B}) p(\eta_i = \text{B} | N_{\text{F,exp}}, N_{\text{B,exp}}) \\ &= p(\rho_i, \vec{\gamma}_i | \theta_{\text{F}}, \eta_i = \text{F}) \frac{N_{\text{F,exp}}}{N_{\text{exp}}} + p(\rho_i, \vec{\gamma}_i | \theta_{\text{B}}, \eta_i = \text{B}) \frac{N_{\text{B,exp}}}{N_{\text{exp}}}, \end{aligned} \quad (5)$$

where the probability of each class is just the expected fraction of the total number. Since this is a sum of probability densities, special care must be taken to ensure all terms are properly normalised, such that

$$\int_{\rho_{\text{thr}}}^{\infty} d\rho \int d\gamma p(\rho, \gamma | \theta_{\eta}, \eta_i) = 1, \quad (6)$$

for $\eta = \text{F}$ and $\eta = \text{B}$, where ρ_{thr} is a minimum SNR value for which events are considered, either as a result of the event generation method or as a choice to limit computational costs. Neglecting this normalization would introduce an artificial preference for one component over the other. An extension to further sub-populations is simply achieved by including additional classes with their own rate and shape parameters.

Recombining the pieces, we can write the desired posterior in Eq. (1) as

$$p(N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}} | \{\rho\}, \{\vec{\gamma}\}, N_{\text{obs}}) \propto p(N_{\text{F,exp}}, N_{\text{B,exp}}, \theta_{\text{F}}, \theta_{\text{B}}) e^{-N_{\text{exp}}} \prod_i [p(\rho_i, \vec{\gamma}_i | \theta_{\text{F}}, \eta_i = \text{F}) N_{\text{F,exp}} + p(\rho_i, \vec{\gamma}_i | \theta_{\text{B}}, \eta_i = \text{B}) N_{\text{B,exp}}]. \quad (7)$$

This expression is similar to Eq.(21) from Farr *et al.* [18] with an explicitly added dependence on source parameter estimates. This implies that our formalism reduces to the Farr *et al.* [18] result as used by LVC to estimate binary black hole merger rates [2, 3], if the event distribution over mass or similar parameters is not free to vary.

The dependence on event parameters arises through the use of samples $\vec{\gamma}_i$, $i = 1 \dots n$, drawn from the likelihood function of the data d for a given point in parameter space $p(d|\gamma)$. These allow us evaluate the population likelihood

function via marginalisation over the unknown true parameters, using the n samples to perform a Monte Carlo integral as in [53],

$$\begin{aligned} p(d|\theta_F) &= \int p(d|\gamma)p(\gamma|\theta_F)d\gamma \\ &= \langle p(\vec{\gamma}_i|\theta_F) \rangle_{p(d|\vec{\gamma}_i)} \\ &= n^{-1} \sum_j p(\gamma_j|\theta_F). \end{aligned} \quad (8)$$

Samples from the likelihood therefore serve as a useful intermediate representation of the raw interferometer data d .

To obtain a quantity directly relevant for an astrophysical interpretation, the expected number can be transformed into the local merger rate R using the observing time T and the sensitive volume $V(\gamma)$:

$$R = \frac{N_{F,\text{exp}}}{T \int d\gamma V(\gamma)p(\gamma|\theta_F)}, \quad (9)$$

where the integral marginalises over the space of source parameters γ . In practice $V(\gamma)$ is estimated for a particular data-set by a Monte Carlo campaign, adding (‘injecting’) a large number of simulated signals to the data and counting the resulting number of events above threshold.

An additional quantity which is not directly used in the derivation of our model, but is important for an astrophysical interpretation, is the probability of any given event originating from the astrophysical foreground p_{astro} :

$$p(\eta_i = \text{F}, N_{F,\text{exp}}, N_{B,\text{exp}}, \theta_F, \theta_B | \rho_i, \vec{\gamma}_i) = \frac{p(\rho_i, \vec{\gamma}_i | \theta_F, \eta_i = \text{F}) N_{F,\text{exp}} \times p(N_{F,\text{exp}}, N_{B,\text{exp}}, \theta_F, \theta_B)}{p(\rho_i, \vec{\gamma}_i | \theta_F, \eta_i = \text{F}) N_{F,\text{exp}} + p(\rho_i, \vec{\gamma}_i | \theta_B, \eta_i = \text{B}) N_{B,\text{exp}}}, \quad (10)$$

$$p_{\text{astro},i} = \int_0^\infty dN_{F,\text{exp}} \int_0^\infty dN_{B,\text{exp}} \int d\theta_F \int d\theta_B p(\eta_i = \text{F}, N_{F,\text{exp}}, N_{B,\text{exp}}, \theta_F, \theta_B | \rho_i, \vec{\gamma}_i), \quad (11)$$

where we marginalised over the population parameters.

III. TOY MODEL

We construct a simple toy model of the universe to test our inference framework in various ways. The toy model allows us to generate a large number of realisations from the same underlying parameters, and to be certain that we use the correct model when analysing these realisations. For simplicity we consider a static, flat, and finite universe. Events are characterised completely by the distance r to the source and a single mass parameter m , which takes the place of the γ used in the previously derived expressions. This mass parameter can be thought of as similar to the chirp mass. Additional effects such as inclination, spins, mass ratio, or antenna patterns are ignored.

For our detection statistic ρ we simply use (a simplified proxy for) the SNR, whose expected value ρ_{true} is determined by r and m as

$$\rho_{\text{true}} = K \frac{m}{r}. \quad (12)$$

where K is an arbitrary constant which quantifies the detector sensitivity. We also model the uncertainty in the estimation of the mass parameter as

$$\sigma_{\text{PE}} \propto \frac{m}{\rho} \quad (13)$$

which is a simplification of the relation given in [54].

To apply the generic form found in Eq. (7) to a specific problem, we need to evaluate the terms $p(\rho_i, \vec{m}_i | \theta_F, \eta_i = \text{F})$ and $p(\rho_i, \vec{m}_i | \theta_B, \eta_i = \text{B})$. This involves finding a functional form for the selection effects. For sources distributed uniformly in our static universe we can derive the needed expression directly by manipulating the joint distribution of masses and observed (detection) SNRs ρ_{obs} . We use the SNR relation defined above in Eq. (12), which defines a mass dependent lower cut-off $\rho_{\text{cutoff}}(m) = K m r_U^{-1}$ to the ρ_{true} possible in our toy universe. Additionally we use the fact

that the euclidean distances $r = |\vec{x}|$ of sources distributed uniformly in volume follow a r^2 -distribution. Therefore

$$\begin{aligned}
p(\rho_{\text{obs}}, m) &= \int_{V_U} d\vec{x} p(\rho_{\text{obs}}, m|\vec{x})p(\vec{x}) \\
&= \frac{1}{V_U} \int_0^{r_U} dr 4\pi r^2 p(\rho_{\text{obs}}, m|r) \\
&= \frac{1}{V_U} \int_0^{r_U} dr \int_{\rho_{\text{cutoff}}(m)}^{\infty} d\rho_{\text{true}} 4\pi r^2 p(\rho_{\text{obs}}|\rho_{\text{true}})p(\rho_{\text{true}}|m, r)p(m) \\
&= \frac{4\pi p(m)}{V_U} \int_{\rho_{\text{cutoff}}(m)}^{\infty} d\hat{\rho} \int_{\rho_{\text{cutoff}}(m)}^{\infty} d\rho_{\text{true}} p(\rho_{\text{obs}}|\rho_{\text{true}}) \frac{(Km)^3}{\hat{\rho}^4} \delta(\hat{\rho} - \rho_{\text{true}}) \\
&= \frac{4\pi K^3 m^3 p(m)}{V_U} \int_{\rho_{\text{cutoff}}(m)}^{\infty} d\rho_{\text{true}} \rho_{\text{true}}^{-4} p(\rho_{\text{obs}}|\rho_{\text{true}}) \\
&\propto m^3 p(m) \int_{\rho_{\text{cutoff}}(m)}^{\infty} d\rho_{\text{true}} \rho_{\text{true}}^{-4} p(\rho_{\text{obs}}|\rho_{\text{true}}), \tag{14}
\end{aligned}$$

where r_U and V_U are the radius and volume of our universe, respectively, and $\rho_{\text{cutoff}}(m) = \frac{Km}{r_U}$ is the lower SNR cut-off defined by a source of mass m being placed at the maximum allowed distance r_U . Thus, the SNR distribution for astrophysical sources is $p(\rho) \propto \rho^{-4}$, and we expect the observed mass distribution to be biased by a factor of m^3 . The mass and SNR components of Eq. (14) are generally connected via the mass dependent SNR cut-off. The term $p(\rho_{\text{obs}}|\rho_{\text{true}})$ accounts for the shift in search SNR relative to the expected value due to detector noise:

$$p(\rho_{\text{obs}}|\rho_{\text{true}}) = \chi_{\text{NC}}(\rho_{\text{obs}}; \lambda = \rho_{\text{true}}, k = 2), \tag{15}$$

where χ_{NC} is the non-central chi distribution with a non-centrality of $\lambda = \rho_{\text{true}}$ and $k = 2$ degrees of freedom.

For real binary merger events we can pursue an analogous derivation, though the resulting relation differs as the SNR is a more complex function of event parameters than Eq. (12). Additional complications arise if the detector is sensitive to events at cosmological distances, causing the observed masses to be redshifted by a distance-dependent amount.

The search and parameter estimation analyses that produce our events can only cover a finite range of masses, of which we denote the limits as m_{min} , m_{max} . We will assume that all astrophysical foreground events have masses lying within these limits; in practice one should take sufficiently wide limits that the density of foreground events at these limits becomes vanishingly small.

For the mass distribution of the astrophysical foreground we consider two types of population distribution: a truncated power law

$$p(m|\theta_F) \equiv p(m|\alpha, m_{\text{low}}, m_{\text{high}}) \propto \begin{cases} m^\alpha & \text{if } m_{\text{low}} < m < m_{\text{high}} \\ 0 & \text{else} \end{cases} \tag{16}$$

with three free parameters, the slope α , lower mass cut-off m_{low} , and high mass cut-off m_{high} . The two mass cut-offs are constrained by the mass range considered as $m_{\text{min}} \leq m_{\text{low}} < m_{\text{high}} \leq m_{\text{max}}$. The second population is a Gaussian

$$p(m|\theta_F) = p(m|\mu, \sigma) \propto \begin{cases} \mathcal{N}(m; \mu, \sigma) & \text{if } m_{\text{min}} < m < m_{\text{max}} \\ 0 & \text{else} \end{cases} \tag{17}$$

with two free parameters, the mean μ and standard deviation σ . Strictly this distribution is a truncated Gaussian, however in practice we consider parameter ranges such that $p(m|\theta_F) \ll 1$ at the boundaries. In contrast to the explicit differentiation between the true and observed SNR, Bayesian parameter estimation provides us with samples from the probability distribution of the true mass which are used directly as in Eq. (8), which eliminates the need to introduce a variable representing an observed mass.

In the background case, there are no selection effects, and we assume the noise characteristics are such that there is no correlation between the mass distribution and the SNR distribution. As a result, $p(\rho_i, \vec{m}_i|\theta_B, \eta_i = B)$ decomposes as simply

$$p(\rho_i, \vec{m}_i|\theta_B, \eta_i = B) = p(\rho_i|\theta_B, \eta_i = B)p(\vec{m}_i|\theta_B, \eta_i = B). \tag{18}$$

Note that in realistic data the SNR distribution of background events may be strongly dependent on the mass (and other template parameters) [5] so this decomposition is not necessarily valid.

The expected rate and distribution of background events caused by instrumental noise can, in practice, be measured to high precision using techniques such as time-shifted analyses [6, 13] (see also [55]). For our artificial universe we have the freedom to choose the SNR and mass distributions, though this choice was informed by observed distributions in real data. We choose a power law with slope -12 in SNR; the mass posteriors are of constant width with their central values distributed uniformly between m_{\min} and m_{\max} ,

$$p(\rho|\eta=\text{B}) \propto \rho^{-12}, \quad (19)$$

$$p(m|\eta=\text{B}) \propto 1. \quad (20)$$

More realistic choices would include the effect of template bank density [56] and transient noise glitches [10, 12] on the distribution of noise triggers over mass space. Note that our inference of the foreground mass distribution is expected to become more precise the more distinct the foreground and background are, especially in SNR. Here, both SNR distributions are falling power laws, however background drops off much more rapidly than foreground.

Finally we combine the mass distribution with Eqs. (14)-(15). Using (16) for the truncated power law we obtain:

$$p(\rho, m|\alpha, m_{\text{low}}, m_{\text{high}}, \eta=\text{F}) \propto \int_{\rho_{\text{cutoff}}(m)}^{\infty} d\rho_{\text{true}} \rho_{\text{true}}^{-4} \chi_{\text{NC}}(\rho; \lambda=\rho_{\text{true}}, k=2) \times \begin{cases} m^{\alpha+3} & \text{if } m_{\text{low}} < m < m_{\text{high}} \\ 0 & \text{else} \end{cases} \quad (21)$$

Using (17) for the Gaussian:

$$p(\rho, m|\mu, \sigma, \eta=\text{F}) \propto \int_{\rho_{\text{cutoff}}(m)}^{\infty} d\rho_{\text{true}} \rho_{\text{true}}^{-4} \chi_{\text{NC}}(\rho; \lambda=\rho_{\text{true}}, k=2) \times \begin{cases} m^3 \mathcal{N}(m; \mu, \sigma) & \text{if } m_{\min} < m < m_{\max} \\ 0 & \text{else} \end{cases} \quad (22)$$

The background model does not involve selection effects and yields

$$p(\rho, m|\eta=\text{B}) \propto \rho^{-12}. \quad (23)$$

In general normalising these expressions requires an integral over ρ and m which can be difficult or computationally expensive. In our model this simplifies somewhat as the integrand $\rho_{\text{true}}^{-4} \chi_{\text{NC}}(\rho; \lambda=\rho_{\text{true}}, k=2)$ happens to assume values very close to zero for the ρ_{cutoff} values of $[0.25, 4]$ allowed by our prior mass range of $[5, 80]$, therefore we are able to approximate $\rho_{\text{cutoff}} = 1$.

To generate the artificial datasets we draw a total number of foreground and background events from a Poisson distribution around the true values determined by the intrinsic rate. Each of those events corresponds necessarily to a local maximum of signal likelihood over time, mass, and, in general, other parameters - we generally approximate this local maximum as a multivariate Gaussian distribution. For each foreground event we then draw the true mass and distance, from which we can uniquely determine the intrinsic SNR. We then simulate the impact of noise on the measurement of both SNR and mass by drawing a value from a non-central chi distribution around the true SNR to obtain the observed SNR, and drawing the maximum likelihood value from a normal distribution around the true mass with a width as determined by Eq. (13). We use a uniform in mass prior, therefore the posterior samples for the mass estimate are drawn from a Gaussian with the same width around the maximum likelihood value. For background events the observed SNR is drawn directly from a power-law as the background SNR distribution is determined empirically from the observed SNR values, while the posterior samples for the mass are drawn from a constant-width Gaussian around a central value drawn from a uniform distribution between m_{\min} and m_{\max} for each event.

Our method does not require us to make strong assumptions about the shape or width of the mass likelihood, however it is important to generate these artificial results carefully as negligence can have unexpected consequences. In practice, we have found the scaling and width of the posteriors to have little effect on our results when the posteriors are of smaller scale than the population.

IV. TOY MODEL RESULTS

We applied our method to a large number of realisations for each choice of foreground distribution, though the figures in the following section only show results for a single realisation. The results across realisations will be given in

text only. The mass limits chosen for all toy model results⁴ were $m_{\min} = 5$, $m_{\max} = 80$. The total expected number of events above an SNR of 8, the lowest threshold considered, is 1600, with 95% contamination due to background events. The chosen slope of -12 for the background SNR distribution is less steep than in typical LIGO-Virgo analyses for stellar-mass compact binary mergers; our choice exaggerates the transition region in which the chances of an event belonging to either foreground or background are comparable.

To simulate the limitations due to the computational costs of the analysis we impose a SNR threshold on events, assessing its influence on our inferences by varying its value between 8 and 15. These SNR values would typically be considered as sub-threshold, since an SNR of ≈ 13.7 is required for an event to have a p_{astro} value of 50%, and to reach $p_{\text{astro}} = 90\%$ an SNR of ≈ 18 is needed. Lastly, the free parameters in Eq. (12) and Eq. (13) are chosen such that an event of true mass $m = 30$ at a notional distance of 400 Mpc has a mass posterior with width $\sigma_{\text{PE}} \approx 1$, and has an SNR of $\rho \approx 50$. The width of the mass posteriors of background events is set to the constant value of 3.2, typical of foreground events at the lowest SNR considered.

The priors chosen are flat in all hyper-parameters, with two exceptions: the width of Gaussian populations, where the prior was flat-in-log, and the expected number of astrophysical foreground events, where we used a Jeffreys prior:

$$\text{prior}(\sigma) \propto \frac{1}{\sigma}, \quad (24)$$

$$\text{prior}(N_{\text{F,exp}}) \propto \frac{1}{\sqrt{N_{\text{F,exp}}}}. \quad (25)$$

Parameter estimation was performed using the `emcee` [57] implementation of an Affine Invariant Markov chain Monte Carlo Ensemble sampler [58].

A. Power law distribution

The first population considered was the truncated power law, which was inspired by the idea that black hole masses may be distributed analogously to the initial mass function of their progenitor stars. We add parameters m_{\min} and m_{\max} to define the lower and upper limits of the power law distribution. This is motivated by the desire to determine whether there are gaps in the astrophysical black hole mass distribution: at the low end to compare with the apparent lower limit of black hole mass in X-ray binaries [59], and at the high end to determine the maximum mass above which a pair-instability supernova completely disrupts the star [60]. In our simulation we chose the power law slope to be -2.4 , and the cut-off values to be 12 and 64.

Our primary results, the estimates of model parameters and their correlations, are shown in figure 1. These results use a SNR threshold 8, the lowest value for which we run our analysis, as we would expect this to yield the best possible parameter estimates. Notable features are the large spread of possible merger rate densities (abbreviated as “Rate”) and their correlation with the lower mass cut-off. This is a consequence of the fact that the power law slope is effectively increased by 3 due to selection effects, thus detected events are described by a positive slope. The detection bias towards high mass means that fewer events are available to define the lower cut-off value, and low mass events which are observed tend to have lower SNR values. As the total rate is still dominated by low mass (and low-amplitude) events, the large uncertainty of the low mass cut-off yields a high uncertainty on the rate. The estimated fraction of foreground events in the sample of observed events couples linearly to the merger rate density, but is less significant than the lower mass cut-off.

To assess our method in the light of its main goal of avoiding bias while lowering the SNR threshold, a single analysis result is insufficient. Therefore, we analyse the same data with a range of different SNR thresholds to observe the change in the hyper-parameter estimates. Figure 2 shows the marginalised posteriors for the rate and power law slope as a function of SNR threshold. We can observe the posteriors growing slightly wider as the SNR threshold is increased and information from fewer events is considered. The result from one single realisation is, however, not necessarily representative of the general behaviour. Combining the results from multiple realisation shows there is no noticeable bias regardless of the threshold chosen, and estimates improve slightly as the threshold is lowered. The width of the 90% credible intervals decreases by factors of ≈ 1.5 for the power-law slope, ≈ 1.1 and ≈ 1.05 for the lower and upper mass cut-offs respectively, and ≈ 1.2 for the log of the inferred merger rate density.

Given the estimates of population parameters we can also compute an estimate of the underlying mass distribution, which we show in figure 3. Here we show the 50% and 90% confidence bands, as defined by computing the percentiles of $p(m|\theta)$ across all samples θ from the posterior for any given mass m . We observe that the true distribution is

⁴ Since we do not claim a specific link to astrophysics in the toy model, the mass units are arbitrary.




FIG. 1. Parameter estimates for a single realisation of the toy model. The foreground population model is a truncated power law with slope α , and cut-offs m_{\min} and m_{\max} . The expected number of events above the SNR threshold of 8 are 1620, 5% of which are expected to be foreground events. The black lines show the kernel density estimate of the posterior, the thin blue lines the corresponding histogram. The red dash-dotted line indicates the true value for the underlying population.

contained well within the credible interval and deviations are generally caused by an underestimated lower cut-off. In general there is a trade-off between expanding the bounds of the mass distribution to include additional events, and shrinking it to increase the PDF for highly significant events. The lower cut-off tends to have more freedom of movement as there are fewer high SNR events at low mass to constrain it.

As a final result for this population, figure 4 shows the estimated probability of any given event to have an astrophysical origin p_{astro} , and how it compares to a SNR-only estimate indicated by the black dash-dotted line. While this figure does not show quantitative results, we do observe that foreground events are largely located above the dash-dotted black line, indicating that our confidence in them being real has increased, while background events tend to be located below and are often on the $p_{\text{astro}} = 0$ line when their masses are outside the hard cut-offs of the truncated power law population. Thus we find, as expected, that the discrimination between signal and noise populations is improved with the incorporation of information about their mass distributions [56]. We determine that the percentage point difference between p_{astro} using our method and the SNR-only approach to be $\approx 3\%$, although this includes events with tiny absolute shifts due to them being very close to either 0% or 100% in the first place.

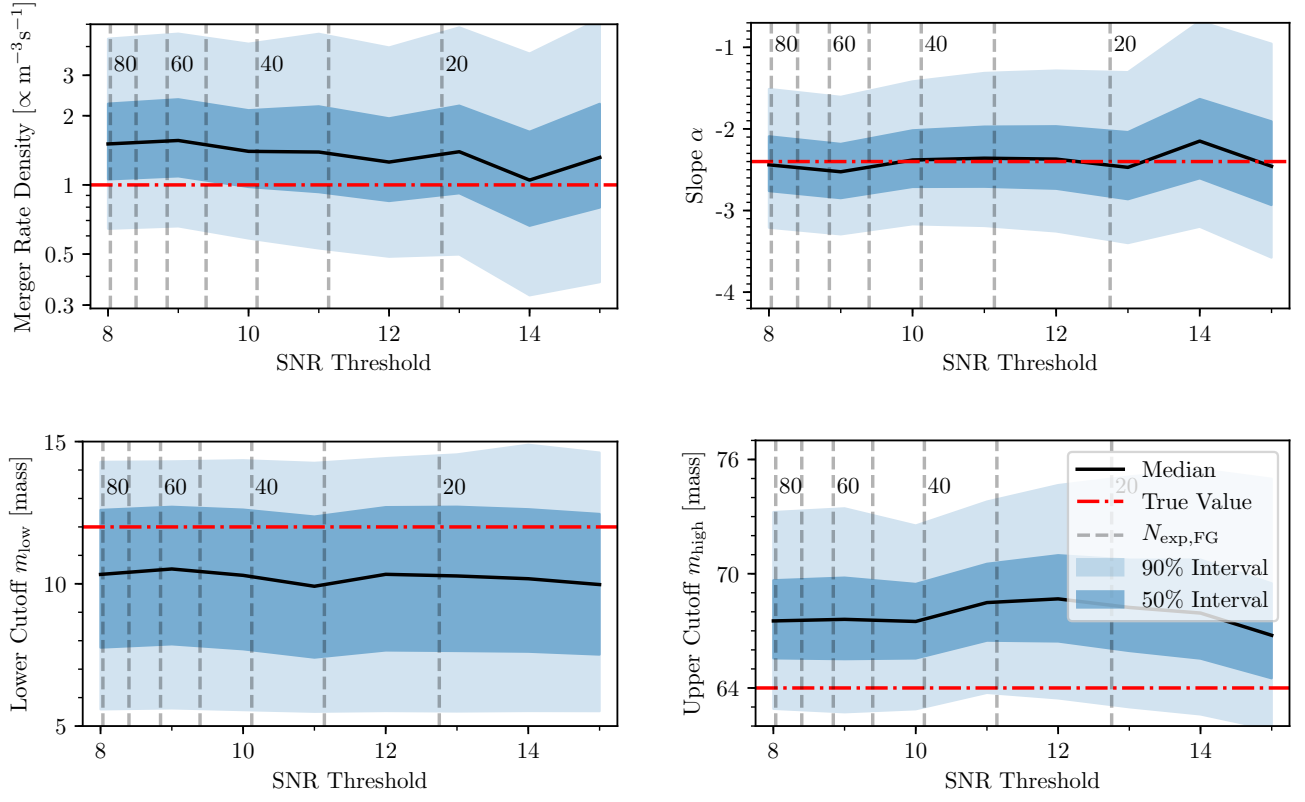


FIG. 2. Confidence intervals for individual parameters of one realisation of the truncated power-law model, as a function of SNR threshold. The parameters shown are the inferred astrophysical merger rate (upper left) and power law slope (upper right), as well as the low (lower left) and high (lower right) mass cut-offs. The red dash-dotted line indicates the true value for the underlying population. Dashed grey lines indicate the expected number of foreground events at the given SNR threshold.




FIG. 3. The inferred mass distribution of the foreground population. The bands indicate the given percentiles in the probability density at any given mass across all posterior samples.




FIG. 4. The probability that an event is caused by detector noise rather than being of astrophysical origin, $1 - p_{\text{astro}}$, versus SNR. The foreground population model is a truncated power-law distribution. Blue and red dots represent foreground and background events respectively. The dash-dotted line shows the probability that would be inferred by a SNR based estimate [18], assuming the relative number of expected foreground and background events is known perfectly. The inset focuses on the region with background events and emphasises events which are unlikely to be astrophysical by using a linear scale.

B. Gaussian distribution

The second core population considered is a simple Gaussian with a very small width. This population was chosen to test the inference on hard-to-infer parameters and to see the effect a very distinctive distribution has on the discriminating power of our method. As such, we chose the width to be very narrow with a standard deviation of 1.6 around a mean mass of 27. This means that the width of the population is smaller than the typical width of the mass posterior of $\approx 2 - 3$ for each individual event, which we expect to be hard to estimate with a relatively small number of events. On the other hand, the discriminating power of using information from the mass estimates should be much greater than for a wide distribution like the truncated power law used in the previous section.

The parameter estimates for a single realisation are shown in figure 5. We observe that true width of the population σ is contained comfortably within the inferred posterior, though the uncertainty is rather large. It is generally overestimated slightly. Similarly the mean of the population is found well with an uncertainty comparable to the population width. The rate is constrained much better than in case of the truncated power law as this model lack the degeneracy between the rate and a poorly constrained population parameter. The lack of a strong correlation between a population parameter and the merger rate density also highlights its linear relation to the estimated number of foreground events contained within the sample.

When lowering the SNR threshold from 15 down to 8, the sizes of the 90% confidence intervals of the population parameters and merger rate density decrease by median factors of $\approx 1.3 - 1.7$. The true mass distribution is well within the confidence interval shown in figure 6, though the true distribution is somewhat more narrow than inferred as seen previously in figure 5.

The comparison of the estimated p_{astro} as shown in figure 7 show how important the inclusion of masses is for this population. We can clearly identify the band of foreground events for which $1 - p_{\text{astro}}$ is smaller by factors of a few up to 10 compared to the SNR based estimate. Only three out of ≈ 74 foreground events lost any p_{astro} . There are some background events which, too, have their p_{astro} increased, though most are demoted and often down to effectively 0.

C. Incorrect models - No background component

Previous analyses of gravitational wave populations (such as the power law model used in [24]) use a high threshold to ensure a high probability that the events used are of astrophysical origin, in effect neglecting the possibility of




FIG. 5. Parameter estimates for a single realisation. The foreground population model is a Gaussian with two free parameters, the mean μ and width σ . The true expected numbers of events above a SNR threshold 8 is 1620, 5% of which are expected to be foreground events. The black lines show the kernel density estimate of the posterior, the thin blue lines the corresponding histogram. The red dash-dotted lines indicates the true values for the underlying population.

background. Here we investigate the behaviour of our toy model with the background component disabled, corresponding to such a scenario. This shows the results one would obtain if simply fitting the foreground model to a contaminated dataset. The underlying population is a truncated power law identical to the one used in the first set of results presented in section IV A.

The results are shown in Fig. 8, where we observe the inferred distribution to be very different from the true one when the lowest SNR threshold of 8 is used and the dataset is 95% polluted (left panel). The mass cut-offs are extended to the edges of the prior ranges to incorporate noise events at those values. The confidence interval includes the true value as long as the SNR threshold is sufficiently high since the number of background events is negligible, but trends towards -3 as the threshold is lowered. This is expected since the background dominates the low SNR region and has a power law slope of 0 in mass. This slope corresponds to an actual slope of -3 when selection effects are considered. In the right panel we see the effect on the estimation of the power law slope as the threshold is varied. At a pollution fraction of 50% (SNR 13.7) the statistical uncertainty of the slope becomes large enough that the systematic bias is not visible.




FIG. 6. The inferred mass distribution of the foreground population. The bands indicate the given percentiles in the probability density at any given mass across all posterior samples.




FIG. 7. The probability of an event being caused by detector noise rather than being of astrophysical origin, $1 - p_{\text{astro}}$, versus SNR. The foreground population model is a narrow Gaussian distribution. Blue and red dots represent foreground and background events respectively. The dash-dotted line shows the probability that would be inferred by a SNR based estimate [18], assuming the relative number of expected foreground and background events is known perfectly. The inset focuses on the region with background events and emphasises events which are unlikely to be astrophysical by using a linear scale.

D. Incorrect models - Neglected selection effects

The second kind of error we considered was the neglect to properly account for the mass dependence of selection effects. In case of a power law distribution this is trivial, as it simply adds +3 to the inferred value of the slope. Therefore we chose a Gaussian as the population, and we increased the width to 9 to highlight the impact of selection effects on the inferred population. Figure 9 shows that the selection effects effectively shift the distribution towards higher masses. This is a general feature as the m^3 term strongly favours high mass events in the observed set of events. Depending on the population the may also affect the width of the population, which happened to be a very




FIG. 8. Results for the power law model when we neglect to account for contamination due to noise. Left: Inferred mass distribution at the lowest SNR threshold of 8, the PDF percentiles are calculated across the population posterior at any given mass. Right: Inferred power law slope as a function of SNR threshold. Vertical dashed lines indicate the expected number of true events above a particular SNR.




FIG. 9. The inferred mass distribution for the Gaussian model without compensating for the mass dependence of selection effects, using the lowest SNR threshold of 8. The PDF percentiles are calculated across the population posterior at any given mass.

minor effect in this case.

Together with the previous section IV C this illustrates that population inference can be made impossible even when the model matches the underlying distribution. Accounting for the presence of noise and selection effects is essential for correct inference and to avoid bias when attempting to lower the SNR threshold.

V. ADVANCED LIGO ENGINEERING RUN

After the successful tests using the simple toy model we also applied this method to more realistic data. For this we used simulated data from the fourth advanced LIGO engineering run (ER4). LIGO Hanford noise data were simulated as Gaussian noise using the LIGO design sensitivity noise spectrum [4], while LIGO Livingston data were derived from an instrumental channel monitoring the power-stabilized laser, recoloured to the same average target spectrum. Simulated gravitational wave signals were then added (“injected”) to the recorded noise data streams. The injected population of binary black hole mergers was chosen to be uniform in both component masses between limits of 5 and 20 M_{\odot} , and uniform in volume with no cosmological effects included. The EOBNRv2HM approximant tuned to




FIG. 10. Cumulative number of (simulated) events detected in LIGO ER4 engineering run data vs. threshold search detection statistic ρ_c . Black steps indicate the search results, with $\pm\sqrt{N}$ bands indicating expected counting fluctuations. Dark red and light blue lines indicate power-law models of signal and noise distributions, respectively. Dotted light blue lines indicate empirical estimates of the noise distribution from each of 3 disjoint analysis periods, showing that the background model $p(\rho_c|\eta_i=B) \propto \rho_c^{-54.8}$ is sufficiently accurate in the range of interest. Dark green dashed lines show the total expected number of events (signal+noise) as a 90% credible band.

numerical relativity [61] was used to simulate binary black hole mergers including non-dominant GW emission modes, for non-spinning binary components. The intended astrophysical rate corresponding to the injected merger signals was $5 \text{ Gpc}^{-3}\text{yr}^{-1}$.⁵

We searched the data for signals using the `pycbc` [5, 6, 62, 63] pipeline to cover binary mergers of non-spinning components with masses between 3 and 50 M_\odot ; this range also defined the prior for parameter estimation performed on each event using `LALInference` [64]. The search detection statistic for candidate events, ρ_c , is the quadrature sum of χ^2 -reweighted SNRs $\hat{\rho}_{H,L}$ over single-detector events having consistent component masses and times of arrival between the two detectors [27, 65]. The number of events we chose to analyse is limited by computational cost; we impose a threshold $\rho_c > 8$ leaving us with 100 events in ≈ 37 days of LHO-LLO coincident observing time; 51 of these events correspond to known injected signals, with the remainder due to noise fluctuations.⁶

We first determine the rates of signal and noise events and the relative probabilities of signal vs. noise origin for each event [2, 3, 18], given only the ρ_c value of each event, models of the signal and noise event distributions over ρ_c , and an estimate of the total rate of noise events derived from time-shifted analyses [63, 65]. The result of this estimate is summarized in Fig. 10. We find 53 events with a signal probability p_{astro} above 50%, of which 47 have $p_{\text{astro}} > 90\%$. This analysis is comparable to those used to estimate the rate and p_{astro} for binary black hole mergers in the first Advanced LIGO Observing period [2], and does not use information about the mass distributions of signal or noise events, besides the assumption that the signal population is contained within the analysis mass limits.

We now turn to our analysis, which estimates the rate and population model parameters simultaneously. The population model used here is a power law in each component mass, of the form

$$p(m_1, m_2|\theta_F, \eta_i = F) \propto \begin{cases} m_1^\alpha m_2^\beta & \text{if } m_{\text{low}} < m_2 \leq m_1 < m_{\text{high}} \\ 0 & \text{else} \end{cases} \quad (26)$$

where α and β are the two power law slopes, both with true values equal to 0. The mass cut-offs m_{low} and m_{high} are shared between both power laws, resulting in four free parameters in our population model. The selection effects were simulated numerically using the `LALsimulation` [66] implementation of the `IMRPhenomPv2` waveform [67, 68] to implement the method described by Finn & Chernoff [69].

⁵ Due to a software error the amplitude of injected signals was a factor 2 higher than intended, effectively simulating a true merger rate of $40 \text{ Gpc}^{-3}\text{yr}^{-1}$; however in the results presented here, we rescale our rate estimates to compensate for this error.

⁶ In reality we will not have access to an independent record listing all true signals!

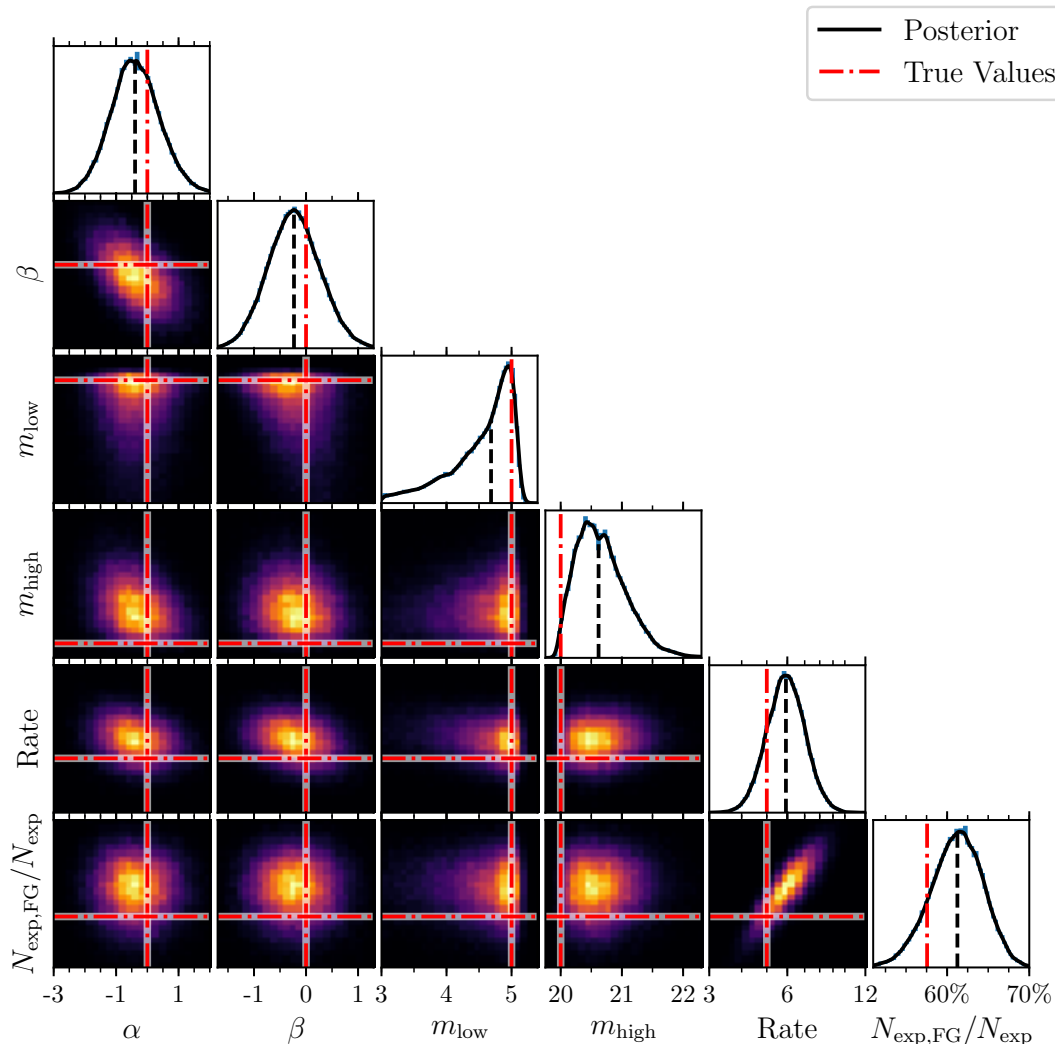


FIG. 11. Parameter estimates for the ER4 dataset. The foreground population model is power law in both component masses with separate slopes α and β and shared cut-offs m_{low} and m_{high} . The cut-off masses and merger rate density are given in units of M_{\odot} and $\text{Gpc}^{-3}\text{yr}^{-1}$ respectively. The black lines show the kernel density estimate of the posterior, the thin blue lines the corresponding histogram. The red dash-dotted line indicates the true value.

As the background model, we used a power law fit to the distribution of ρ_c values from the `pycbc` time-shifted analysis, giving a slope of ≈ -54.8 for $\rho_c > 8$. In a separate step we constructed two dimensional fit to the distribution of component masses in the search results. We find empirically that the background distribution can be approximated as the product of a power-law in chirp mass and a exponential distribution in mass ratio. Using the masses as determined by the search is not strictly the correct approach, which would be to run the time-shifted data through the same LALInference analysis as used for the zero-lag events: this was, though, computationally infeasible. Therefore the search results serve as a proxy for the optimal analysis. We have found empirically that small changes to the background mass distribution do not have a strong effect on the result, which is expected given the dissimilarity of foreground and background distributions. We do not include any additional uncertainty on the background model rate or mass distribution.

The results of estimating the population parameters as displayed in figure 11 show that we successfully recover the true population parameters. The slopes are underestimated slightly which causes the inferred merger rate density to be elevated, although it still encompasses the true value. The mass cut-offs are found well with some tails to lower or higher masses for m_{low} and m_{high} respectively, since the uniform distribution of injections covers all regions of the mass range without major gaps and the cut-offs are shared between both component masses.

Since the slope of the background SNR distribution is much larger than the slope of the astrophysical foreground,

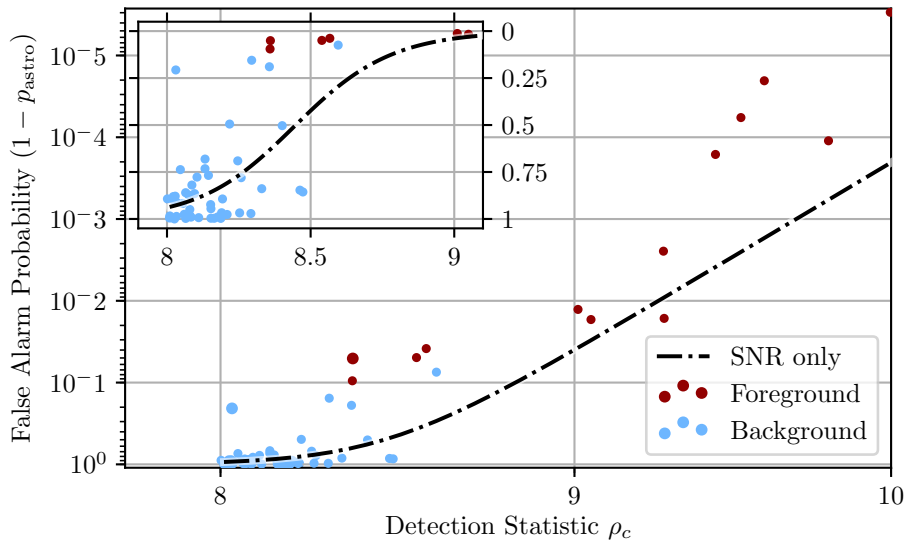


FIG. 12. The probability of an event recovered from the ER4 data being caused by noise instead of corresponding to an astrophysical event, which is $1 - p_{\text{astro}}$. Blue and red dots represent foreground and background events respectively. The dash-dotted line shows the probability that would be inferred by a SNR based estimate [18]. The inset focuses on the region with background events and emphasises events which are unlikely to be astrophysical by using a linear scale.

the transition region where events may belong to either source category with comparable probabilities is quite small. This means there are few events which fall in between the region of certain background and certain foreground, limiting the gains that can be made by our method in this case. Nevertheless, the fact that the foreground mass distribution is quite distinct from the background causes a significant increase in p_{astro} relative to the ρ_c -based approach, as can be seen in figure 12.

As with the previous ranking statistic based analysis we can count the events found with a p_{astro} value above some given threshold. We find 56 events above a threshold of 50%, which is 3 events more than before. The number of included events which we later identified as noise triggers rises from 4 to 5. When the threshold is set to $p_{\text{astro}} > 90\%$, the number of events increases by 5 to 52, though it now includes one noise trigger. Thus, we find that applying our new method to realistic data not only reproduces the results obtained using established methods, but identifies additional foreground events. We also successfully recover the injected population parameters.

VI. OUTLOOK

In this work we have derived a new technique for simultaneous estimation of parameters defining the shape of two or more sub-populations and their expected contribution to the overall number of events. This technique allows us to extract information from formerly sub-threshold events without biasing the result due to uncertainties in classifying their origin. The method is agnostic to the specific choice of threshold, and lowering the threshold will allow more information to be included from events of an uncertain nature, improving the result until the noise completely dominates the signals.

The greatest gains over existing methods are found when there is a large number of events for which the source classification gives comparable probabilities for at least two categories, while the distributions in secondary parameters are very distinct. This behaviour near the transition between populations is likely to be especially useful in the characterisation of weak event populations, such as unresolvable binary mergers at cosmological distances. This is of particular interest when determining whether the source population evolves with redshift. However, this also implies that gains are expected to be small when the primary source classification is very potent and population models are uncertain; in this case our method converges to that with a single population.

While such thresholded analyses that ignore the possible presence of background cannot be guaranteed free of systematic bias, the expected size of bias can be bounded by considering the rate of background events above threshold, as well as the degree of divergence between foreground and background distributions over the parameters of interest. Controlling the bias of a thresholded analysis thus still requires accurate background estimation. In particular, for the

small number of events thus far detected by LIGO-Virgo, possible biases on population inference due to neglecting background contamination are expected to be well below statistical errors.

We have successfully tested this new model on different binary merger mass distributions in an artificial universe, as well as to synthetic LIGO data from an engineering run. This demonstrates the feasibility of applying this method to existing and future LIGO-Virgo observing runs, which should allow a better joint determination of source event rates and distributions. The method itself is, however, not limited to the realm of gravitational wave astronomy, and can be useful whenever a set of data points contains multiple populations.

ACKNOWLEDGEMENTS

We are grateful to the LSC for allowing the use of LIGO software engineering run data hosted on the CIT cluster, and access to LIGO Data Grid computing resources. We are grateful to Chris Pankow for useful discussions and assistance in the ER4 analysis. We would also like to thank Jonathan Gair and Richard O’Shaughnessy for their helpful comments and feedback. This work was supported by STFC grants ST/K005014/2 and ST/M004090/2.

-
- [1] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev. Lett.* **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
- [2] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev. X* **6**, 041015 (2016), arXiv:1606.04856 [gr-qc].
- [3] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Astrophys. J.* **833**, L1 (2016), arXiv:1602.03842 [astro-ph.HE].
- [4] B. P. Abbott *et al.*, *Living Reviews in Relativity* **21**, 3 (2018).
- [5] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev.* **D93**, 122003 (2016), arXiv:1602.03839 [gr-qc].
- [6] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, *et al.*, *Class. Quant. Grav.* **33**, 215004 (2016), arXiv:1508.02357 [gr-qc].
- [7] K. Cannon, R. Cariou, A. Chapman, M. Crispin-Ortuzar, *et al.*, *Astrophys. J.* **748**, 136 (2012), arXiv:1107.2665 [astro-ph.IM].
- [8] K. Cannon, C. Hanna, and J. Peoples, ArXiv e-prints (2015), arXiv:1504.04632 [astro-ph.IM].
- [9] C. Messick, K. Blackburn, P. Brady, P. Brockill, *et al.*, *Phys. Rev.* **D95**, 042001 (2017), arXiv:1604.04324 [astro-ph.IM].
- [10] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, *Astrophys. J.* **849**, 118 (2017), arXiv:1705.01513 [gr-qc].
- [11] B. Allen, *Phys. Rev. D* **71**, 062001 (2005), arXiv:gr-qc/0405045 [gr-qc].
- [12] A. H. Nitz, *Class. Quant. Grav.* **35**, 035016 (2018), arXiv:1709.08974 [gr-qc].
- [13] C. Capano, T. Dent, C. Hanna, M. Hendry, *et al.*, *Phys. Rev. D* **96**, 082002 (2017).
- [14] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Class. Quant. Grav.* **33**, 134001 (2016), arXiv:1602.03844 [gr-qc].
- [15] S. Fairhurst and P. Brady, *Class. Quant. Grav.* **25**, 105002 (2008), arXiv:0707.2410 [gr-qc].
- [16] V. Tiwari, (2017), arXiv:1712.00482 [astro-ph.HE].
- [17] K. K. Y. Ng, S. Vitale, A. Zimmerman, K. Chatziioannou, D. Gerosa, and C.-J. Haster, (2018), arXiv:1805.03046 [gr-qc].
- [18] W. M. Farr *et al.*, *Phys. Rev. D* **91**, 023005 (2015), arXiv:1302.5341 [astro-ph.IM].
- [19] R. Smith and E. Thrane, *Phys. Rev.* **X8**, 021019 (2018), arXiv:1712.00688 [gr-qc].
- [20] T. G. F. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, *et al.*, *Phys. Rev.* **D85**, 082003 (2012), arXiv:1110.0530 [gr-qc].
- [21] W. Del Pozzo, T. G. F. Li, M. Agathos, C. Van Den Broeck, and S. Vitale, *Phys. Rev. Lett.* **111**, 071101 (2013), arXiv:1307.8338 [gr-qc].
- [22] C. Messenger and J. Veitch, *New J. Phys.* **15**, 053027 (2013), arXiv:1206.3461 [astro-ph.IM].
- [23] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Astrophys. J. Suppl.* **227**, 14 (2016), arXiv:1606.03939 [astro-ph.HE].
- [24] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **119**, 141101 (2017).
- [25] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Astrophys. J.* **832**, L21 (2016), arXiv:1607.07456 [astro-ph.HE].
- [26] R. Biswas, P. R. Brady, J. D. E. Creighton, and S. Fairhurst, *Class. Quant. Grav.* **26**, 175009 (2009), [Erratum: *Class. Quant. Grav.*30,079502(2013)], arXiv:0710.0465 [gr-qc].
- [27] Abadie, J. *et al.* (Virgo, LIGO Scientific), *Phys. Rev. D*, 082002 (2011).
- [28] R. O’Shaughnessy, B. Vaishnav, J. Healy, and D. Shoemaker, *Phys. Rev. D* **82**, 104006 (2010).
- [29] M. Dominik, E. Berti, R. O’Shaughnessy, I. Mandel, *et al.*, *Astrophys. J.* **806**, 263 (2015), arXiv:1405.7016 [astro-ph.HE].
- [30] V. Kalogera, K. Belczynski, C. Kim, R. O’Shaughnessy, and B. Willms, *Physics Reports* **442**, 75 (2007), the Hans Bethe Centennial Volume 1906-2006.
- [31] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *The Astrophysical Journal Letters* **818**, L22 (2016).
- [32] I. Mandel and A. Farmer (2018) arXiv:1806.05820 [astro-ph.HE].
- [33] K. Belczynski *et al.*, *Astron. Astrophys.* **594**, A97 (2016), arXiv:1607.03116 [astro-ph.HE].
- [34] K. Belczynski, J. Klencki, G. Meynet, C. L. Fryer, *et al.*, ArXiv e-prints (2017), arXiv:1706.07053 [astro-ph.HE].
- [35] K. Belczynski, D. E. Holz, T. Bulik, and R. O’Shaughnessy, *Nature (London)* **534**, 512 (2016), arXiv:1602.04531 [astro-ph.HE].
- [36] M. Spera, M. Mapelli, and A. Bressan, *Mon. Not. Roy. Astron. Soc.* **451**, 4086 (2015), arXiv:1505.05201 [astro-ph.SR].

- [37] J. W. Barrett, S. M. Gaebel, C. J. Neijssel, A. Vigna-Gómez, *et al.*, (2017), 10.1093/mnras/sty908, arXiv:1711.06287 [astro-ph.HE].
- [38] M. Zevin, C. Pankow, C. L. Rodriguez, L. Sampson, *et al.*, *Astrophys. J.* **846**, 82 (2017), arXiv:1704.07379 [astro-ph.HE].
- [39] D. Wysocki, D. Gerosa, R. O’Shaughnessy, K. Belczynski, *et al.*, *Phys. Rev.* **D97**, 043014 (2018), arXiv:1709.01943 [astro-ph.HE].
- [40] D. Wysocki, J. Lange, and R. O’Shaughnessy, (2018), arXiv:1805.06442 [gr-qc].
- [41] C. Talbot and E. Thrane, *Astrophys. J.* **856**, 173 (2018), arXiv:1801.02699 [astro-ph.HE].
- [42] S. R. Taylor and D. Gerosa, *ArXiv e-prints* (2018), arXiv:1806.08365 [astro-ph.HE].
- [43] I. Mandel, W. M. Farr, A. Colonna, S. Stevenson, *et al.*, *Mon. Not. Roy. Astron. Soc.* **465**, 3254 (2017), arXiv:1608.08223 [astro-ph.HE].
- [44] R. O’Shaughnessy, D. Gerosa, and D. Wysocki, *Phys. Rev. Lett.* **119**, 011101 (2017), arXiv:1704.03879 [astro-ph.HE].
- [45] I. Mandel and R. O’Shaughnessy, *Classical and Quantum Gravity* **27**, 114007 (2010), arXiv:0912.1074 [astro-ph.HE].
- [46] D. Gerosa, M. Kesden, E. Berti, R. O’Shaughnessy, and U. Sperhake, *Phys. Rev. D* **87**, 104028 (2013), arXiv:1302.4442 [gr-qc].
- [47] S. Stevenson, F. Ohme, and S. Fairhurst, *Astrophys. J.* **810**, 58 (2015), arXiv:1504.07802 [astro-ph.HE].
- [48] W. M. Farr, S. Stevenson, M. Coleman Miller, I. Mandel, *et al.*, *Nature* **548**, 426 (2017), arXiv:1706.01385 [astro-ph.HE].
- [49] S. Stevenson, C. P. L. Berry, and I. Mandel, *Mon. Not. Roy. Astron. Soc.* **471**, 2801 (2017), arXiv:1703.06873 [astro-ph.HE].
- [50] D. Gerosa and E. Berti, *Phys. Rev.* **D95**, 124046 (2017), arXiv:1703.06223 [gr-qc].
- [51] M. Fishbach, D. E. Holz, and B. Farr, *Astrophys. J.* **840**, L24 (2017), arXiv:1703.06869 [astro-ph.HE].
- [52] C. Talbot and E. Thrane, *Phys. Rev.* **D96**, 023012 (2017), arXiv:1704.08370 [astro-ph.HE].
- [53] I. Mandel, *Phys. Rev.* **D81**, 084029 (2010), arXiv:0912.5531 [astro-ph.HE].
- [54] C. Cutler and E. E. Flanagan, *Phys. Rev.* **D49**, 2658 (1994), arXiv:gr-qc/9402014 [gr-qc].
- [55] K. Cannon, C. Hanna, and D. Keppel, *Phys. Rev. D* **88**, 024025 (2013), arXiv:1209.0718 [gr-qc].
- [56] T. Dent and J. Veitch, *Phys. Rev.* **D89**, 062002 (2014), arXiv:1311.7174 [gr-qc].
- [57] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *PASP* **125**, 306 (2013), 1202.3665.
- [58] J. Goodman and J. Weare, *Commun. Appl. Math. Comput. Sci.* **65** (2010).
- [59] W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, *et al.*, *Astrophys. J.* **741**, 103 (2011), arXiv:1011.1459 [astro-ph.GA].
- [60] Z. Barkat, G. Rakavy, and N. Sack, *Physical Review Letters* **18**, 379 (1967).
- [61] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, *et al.*, *Phys. Rev.* **D84**, 124052 (2011), arXiv:1106.1021 [gr-qc].
- [62] T. Dal Canton *et al.*, *Phys. Rev.* **D90**, 082004 (2014), arXiv:1405.6731 [gr-qc].
- [63] A. Nitz, I. Harry, D. Brown, C. M. Biwer, *et al.*, “gwastro/pycbc: Post-o2 release 12c,” (2018).
- [64] J. Veitch, V. Raymond, B. Farr, W. Farr, *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
- [65] S. Babak, R. Biswas, P. R. Brady, D. A. Brown, *et al.*, *Phys. Rev. D* **87**, 024033 (2013).
- [66] LIGO Scientific Collaboration, “LSC Algorithm Library,” <https://wiki.ligo.org/DASWG/LALSuite>.
- [67] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, *et al.*, *Phys. Rev. Lett.* **113**, 151101 (2014), arXiv:1308.3271 [gr-qc].
- [68] S. Khan, S. Husa, M. Hannam, F. Ohme, *et al.*, *Phys. Rev. D* **93**, 044007 (2016), arXiv:1508.07253 [gr-qc].
- [69] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993).