

Calibrating signal-to-noise ratio detection thresholds using gravitational-wave catalogs

Matthew Mould^{1,2,*} Christopher J. Moore³ and Davide Gerosa^{4,5,3}

¹*Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA*

²*LIGO Laboratory, Massachusetts Institute of Technology, 185 Albany Street, Cambridge, Massachusetts 02139, USA*

³*School of Physics and Astronomy and Institute for Gravitational Wave Astronomy, University of Birmingham, Birmingham B15 2TT, United Kingdom*

⁴*Dipartimento di Fisica “G. Occhialini,” Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, 20126 Milano, Italy*

⁵*INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, 20126 Milan, Italy*



(Received 20 November 2023; accepted 23 February 2024; published 12 March 2024)

Searching for gravitational-wave signals is a challenging and computationally intensive endeavor undertaken by multiple independent analysis pipelines. While detection depends only on observed noisy data, it is sometimes inconsistently defined in terms of source parameters that in reality are unknown, e.g., by placing a threshold on the optimal signal-to-noise ratio (SNR). We present a method to calibrate unphysical thresholds to search results by performing Bayesian inference on real observations using a model that simultaneously parametrizes the intrinsic network optimal SNR distribution and the effect of search sensitivity on it. We find consistency with a fourth-order power law and detection thresholds of $10.5_{-2.4}^{+2.1}$, $11.2_{-1.4}^{+1.2}$, and $9.1_{-0.5}^{+0.5}$ (medians and 90% credible intervals) for events with false-alarm rates less than 1 yr^{-1} in the first, second, and third LIGO-Virgo-KAGRA observing runs, respectively. Though event selection can only be self-consistently reproduced by physical searches, employing our inferred thresholds allows approximate observation-calibrated selection criteria to be applied when efficiency is required and injection campaigns are infeasible.

DOI: [10.1103/PhysRevD.109.063013](https://doi.org/10.1103/PhysRevD.109.063013)

I. INTRODUCTION

Assembling the LIGO–Virgo–KAGRA (LVK) [1–3] gravitational-wave (GW) catalogs [4–7] requires searching for signals. Classification of a signal as being of astrophysical origin typically involves information from several detectors, including the matched-filter signal-to-noise ratio (SNR) optimized over large template banks, coincidence analyses, data-quality checks, and estimates of the expected noise and astrophysical population. Searches with the multiple independent pipelines are therefore computationally expensive [8–11].

Instead of imposing a selection threshold based on data-dependent quantities such as the observed signal-to-noise ratio (SNR), false-alarm rate (FAR), or probability of astrophysical origin (p_{astro}), it is common to take ranking statistics that are far more simple to compute but that additionally depend on the true binary source parameters, such as the optimal SNR [12,13]. This approach is inconsistent with the data generation process because real searches never have access to the true source properties due

to parameter degeneracies and detector noise [14]. Nevertheless, it has led to several approximations for rapidly computing selection effects [15–19], and has been employed for population inference (e.g., [20–23]) and to process population-synthesis simulations (e.g., [24–28]). Doing so requires selecting a threshold value for the chosen ranking statistic that should well represent the full searches that produce observational GW catalogs.

One approach is to generate a set of fake GW signals, inject them into the search pipelines [29], filter the injections with the physical data-dependent threshold, and optimize the approximate threshold such that the resulting distribution of sources most closely matches the filtered injections. This necessitates running the search pipelines on a large set of sources which carries a significant computational cost, which the goal of approximating selection effects was to avoid (though injection sets generated for other uses can be repurposed [30]). Another approach is to analytically approximate the noise distribution and its effects on physical ranking statistics such as the observed SNR [30], which can be used as an efficient threshold in place of full search pipelines [31,32].

We present a different approach that instead makes use of observational GW catalogs to which the real searches have

*mmould@mit.edu

already been applied. Given a set of astrophysical sources that pass the pipeline threshold and their associated parameter uncertainties, we use the likelihood that the GW data were detected by the real searches but subsequently parametrize detection in terms of a ranking statistic such as the network optimal SNR. We simultaneously infer the detectable and intrinsic properties of the population of this ranking statistic, thus calibrating the unphysical detection threshold against full searches on real signals while accounting for the systematic uncertainties between the two and the statistical uncertainty in each event. Modelers who wish to apply selection effects with an unphysical model such as a threshold in the network optimal SNR can therefore use our inference, which is fully consistent with the aforementioned injection-calibration method [30], to ensure the application of selection effects is representative of full searches.

In Sec. II we describe our Bayesian approach to simultaneous inference of the detectable and intrinsic properties of the ranking-statistic distribution, in this case the network optimal SNR. The results are presented in Sec. III and a concluding discussion is made in Sec. IV.

II. INFERENCE

In the following, we use π to denote prior distributions, \mathcal{L} to denote likelihoods (i.e., data distributions), P to denote discrete distributions (e.g., for detectability which has a binary outcome), \mathcal{P} for posteriors, and p for mixed distributions of observed and latent variables.

A. Hierarchical posterior

We consider the joint hierarchical distribution over detectable data $\{\text{det}_n, d_n\}_{n=1}^N$ and latent variables $\{\rho_n\}_{n=1}^N$, where the data for each GW event is labeled with a flag “det” denoting that it has been included in the catalog according to some criterion. Each ρ_n is drawn from a common underlying population $\pi(\rho|\lambda)$ characterized by global parameters λ . While a physical model defines detectability in terms of the data alone to produce the catalog, we subsequently model detection as depending on ρ and additional global parameters κ . For example, if detection is taken as a step function in ρ , the parameter κ could indicate the location of the threshold.

Assuming the data from each event are independent, the GW population likelihood [33,34] is

$$p(\{\text{det}_n, d_n, \rho_n\}|\kappa, \lambda) \propto P(\text{det}|\kappa, \lambda)^{-N} \times \prod_{n=1}^N P(\text{det}_n|\rho_n, \kappa) \mathcal{L}(d_n|\rho_n) \pi(\rho_n|\lambda), \quad (1)$$

where $\{x_n\}$ is shorthand for $\{x_n\}_{n=1}^N$. Note that unlike a physically-motivated model for detection, i.e., $\mathcal{L}(\text{det}, d|\rho) = P(\text{det}|d) \mathcal{L}(d|\rho)$, detection-dependent terms

$P(\text{det}|\rho, \kappa)$ are retained above due to the assumed dependence on the parameters ρ_n and κ . This is equivalent to a modification of the prior $\pi(\rho|\lambda)$, whereby regions of support are explicitly downweighted or truncated *a priori* [14], according to Bayes’ theorem:

$$\pi(\rho|\text{det}, \kappa, \lambda) = \frac{P(\text{det}|\rho, \kappa)}{P(\text{det}|\kappa, \lambda)} \pi(\rho|\lambda). \quad (2)$$

Consequently, the global detection probability

$$P(\text{det}|\kappa, \lambda) = \int P(\text{det}|\rho, \kappa) \pi(\rho|\lambda) d\rho \quad (3)$$

does not depend on the noise model through the marginalization of the data distribution but does depend on our assumed parametrization of detection through κ . Since we will take simple functional forms for $P(\text{det}|\rho, \kappa)$ and $\pi(\rho|\lambda)$, the integral in Eq. (3) can be computed in closed form. In other words, we are using the observed data distribution, given by $\{\text{det}_n, d_n\}_{n=1}^N$, to infer the observable distribution of ρ , given by $\pi(\rho|\text{det}, \kappa, \lambda)$.

Before performing a population analysis with the joint likelihood above, separate parameter-estimation (PE) analyses are used to infer ρ_n for each event individually under priors $\tilde{\pi}_n(\rho_n)$ that may differ event-to-event and also differ from $\pi(\rho|\lambda)$. However, using the fact that the single-event likelihood is unchanged, the individual posteriors are $\tilde{\mathcal{P}}(\rho_n|d_n) \propto \mathcal{L}(d_n|\rho_n) \tilde{\pi}_n(\rho_n)$. We can therefore write the population likelihood as

$$p(\{\text{det}_n, d_n, \rho_n\}|\kappa, \lambda) \propto P(\text{det}|\kappa, \lambda)^{-N} \times \prod_{n=1}^N P(\text{det}_n|\rho_n, \kappa) \frac{\tilde{\mathcal{P}}(\rho_n|d_n)}{\tilde{\pi}_n(\rho_n)} \pi(\rho_n|\lambda). \quad (4)$$

We do not directly infer each ρ_n from PE runs, which instead infer the posteriors $\tilde{\mathcal{P}}(\theta_n|d_n)$ of the source parameters θ_n under default priors $\tilde{\pi}_n(\theta_n)$. While each of the original priors for θ_n may have the same shape, they could impose different cuts in the parameter space. Each of these distributions directly implies corresponding distributions for ρ_n . In practice, we usually have access to the posteriors $\tilde{\mathcal{P}}(\theta_n|d_n)$ through discrete samples $\{\theta_{ni}\}_{i=1}^{N_n}$ rather than as continuous distributions, where N_n is the number of posterior samples for event n . However, if we wish to evaluate the density in Eq. (4), we must be able to evaluate the individual posterior densities.

Using Bayes’ theorem, the joint posterior is

$$\mathcal{P}(\{\rho_n\}, \kappa, \lambda|\{\text{det}_n, d_n\}) \propto p(\{\text{det}_n, d_n, \rho_n\}|\kappa, \lambda) \pi(\kappa, \lambda), \quad (5)$$

where $\pi(\kappa, \lambda) = \pi(\kappa) \pi(\lambda)$ is the prior over the population-level parameters that we take as independent. We sample the posteriors using NumPyro [35]. In Sec. III we present

results inferred using the joint distributions in Eq. (4), but in Appendix A we show that they are fully consistent with the posteriors inferred using a numerically marginalized likelihood.

B. Optimal signal-to-noise ratio

We will take as our ranking statistic the network optimal SNR ρ . The optimal SNRs for each interferometer individually are given by

$$\rho_i^2(\theta) = 4 \int \frac{|h_i(f, \theta)|^2}{S_i(f)} df, \quad (6)$$

where h_i is the GW signal model at the detector as a function of the source parameters and frequency f , $S_i(f)$ is the one-sided power spectral density (PSD) of the noise, and i indexes the interferometers. For a set of I detectors the network optimal SNR is

$$\rho(\theta) = \sqrt{\sum_{i=1}^I \rho_i^2(\theta)}. \quad (7)$$

We take IMRPhenomXPHM [36] as the waveform model and use PyCBC [37] and Bilby [38] to compute SNRs.

Note that the optimal SNR should also depend on the observed data via estimation of the PSD. However, we make the approximation that it is given by fixed representative noise curves for the first (O1), second (O2), and third (O3) observing runs of the LIGO Hanford,¹ LIGO Livingston,² and Virgo³ detectors. For each event we compute the network optimal SNR with the detectors that were observing. This is simply fixing our choice of ranking statistic; one could simplify further by instead taking a single hypothetical design PSD, computing the single-detector optimal SNR, and calibrating this to the observed data, because the ranking statistic can be anything that reasonably encodes detectability.

For the joint hierarchical likelihood, we construct continuous single-event likelihood approximations. We fit each single-event posterior $\tilde{\mathcal{P}}(\rho_n | d_n)$ for the network optimal SNR as a univariate truncated normal distribution (such that $\rho_n \geq 0$). Their means μ_n and variances σ_n^2 are taken as the sample means and variances of the posterior samples $\{\rho_{ni}\}_{i=1}^N$. For the single-event priors for each ρ_n we take log-normal distributions and use SciPy [39] to fit for their shape and scale parameters by minimizing the cross entropy between the true distributions implied by prior samples and the fitted approximations. This likelihood

approximation is verified by comparison to a numerically marginalized likelihood in Appendix A.

C. Priors

While the individual event SNR priors $\tilde{\pi}_n(\rho_n)$ are determined by the original uninformative PE priors imposed over the source parameters, the population prior $\pi(\rho | \lambda)$ can be reasonably well-motivated astrophysically [40,41]. In the local universe, luminosity distances D approximately follow a geometric prior $\pi(D) \propto D^2$. Since the SNR scales as $\rho \sim 1/D$ the corresponding SNR distribution scales as

$$\pi(\rho) = \pi(D) \left| \frac{dD}{d\rho} \right| \sim \rho^{-4}. \quad (8)$$

Our default model is therefore a power-law distribution in the network optimal SNR. As we are primarily interested in inferring a single threshold value in our chosen ranking statistic, our default detection model will be a Heaviside step function in ρ . We describe in detail the various models we consider below and summarize them in Table I, where the model label indicates the number of its parameters (apart from the flow-based model).

1. Model 1

Our first model has a single parameter—the detection threshold τ —while the intrinsic distribution is a power law with slope $\sigma = 4$, as above. Written explicitly, this is

$$\pi(\rho | \sigma = 4) \propto \rho^{-\sigma}, \quad (9)$$

$$P(\text{det} | \rho, \tau) = \begin{cases} 1 & \text{if } \rho > \tau, \\ 0 & \text{if } \rho \leq \tau. \end{cases} \quad (10)$$

We take a broad uniform prior $\pi(\tau) = \mathcal{U}(0, 20)$.

TABLE I. Summary of the models considered for the intrinsic and detectable populations of the network optimal SNR ρ . Listed are the parameters of the models and corresponding priors, where \mathcal{U} and \mathcal{N} denote uniform and normal distributions, respectively. Apart from Model F, the model labels indicate the number of parameters.

Model	Intrinsic	Detection	Parameters	Priors
1	Power law	Heaviside	Slope σ Threshold τ	= 4 $\mathcal{U}(0, 20)$
2	Power law	Heaviside	Slope σ Threshold τ	$\mathcal{U}(1, 10)$ $\mathcal{U}(0, 20)$
3	Power law	Sigmoid	Slope σ Threshold τ Width ω	$\mathcal{U}(1, 10)$ $\mathcal{U}(0, 20)$ $\mathcal{U}(0, \tau)$
4	Power law	Heaviside	Slope σ Thresholds $\tau_{1,2,3}$	$\mathcal{U}(1, 10)$ $\mathcal{U}(0, 20)$
F	Normalizing flow		Weights $w_{1,\dots,42}$	$\mathcal{N}(0, 1)$

¹<https://dcc.ligo.org/LIGO-G1500622/public>.

²<https://dcc.ligo.org/LIGO-G1401390/public>.

³<https://dcc.ligo.org/P1800374/public> for O2 and <https://dcc.ligo.org/LIGO-T2000012/public> for O3.

2. Model 2

Our two-parameter model is the same as above except the slope σ is also free to vary, for which we take a uniform prior $\pi(\sigma) = \mathcal{U}(1, 10)$.

3. Model 3

Model 3 is used to test the assumption of a strict detection threshold by replacing it with a smooth transition function. Above the threshold we take the same power law as in Model 2, while below it the detection probability,

$$P(\text{det} | \rho, \tau, \omega) = \begin{cases} S(\rho, \tau, \omega) & \text{if } \tau - \omega \leq \rho \leq \tau, \\ 0 & \text{if } \rho < \tau - \omega, \end{cases} \quad (11)$$

is parametrized by a sigmoid function,

$$S(\rho, \tau, \omega) = \frac{1}{2} \left[1 + \sin \left(\frac{\pi}{2} + \frac{\rho - \tau}{\omega} \pi \right) \right], \quad (12)$$

that symmetrically ramps up to unity at $\rho = \tau$ over an interval of width ω . Along with the same priors as in Model 2, we take $\pi(\omega|\tau) = \mathcal{U}(0, \tau)$.

4. Model 4

Given that the detection catalog contains observations from three distinct observing runs between which the detector sensitivities improved [42], we also consider an extension of Model 2 in which there are separate thresholds— τ_1 , τ_2 , and τ_3 —for events in O1, O2, and O3, respectively. This model contains four parameters; one slope and three Heaviside thresholds. We take the same priors as in Model 2.

5. Model F

Finally, we consider a more flexible model that jointly parametrizes the detection probability and intrinsic distribution, as in Eq. (2), with a normalizing flow. In particular, we use a block neural autoregressive flow [43] with a single hidden layer of ten units, composed with an exponential function to enforce positivity, as implemented in Refs. [35,44]. This model is deliberately overparametrized with 42 parameters (the weights of the neural network, for which we take standard normal priors) to test our assumptions on the shape of the detectable distribution for ρ . In particular, we can test whether a power law is a representative model for larger SNRs, the location of the peak of the distribution, and its width at lower SNRs.

D. Event selection

For consistency with the SNR computation in Sec. II B, we use the public parameter-estimation samples from the IMRPhenomXPHM analyses [45,46] in the combined GWTC-2.1 [6] and GWTC-3 [7] catalogs. Of these, three were observed in O1, seven in O2, and 78 in O3, while two

events excluded by this requirement are GW170817 and GW190425. GW170817 in particular is the event with the highest SNR, implying it is informative for the upper tail of the SNR distribution and could therefore impact inferred power-law slopes, but we do not expect its exclusion to impact detection thresholds which are informed by events with low SNRs. To the best of our knowledge, the released posteriors for GW170817 do not include all parameters required to compute SNRs, meaning it could not be included in our analysis even taking a different waveform model. Events are included in these catalogs if they are determined to be of astrophysical origin with probability $p_{\text{astro}} > 0.5$ in at least one of the search pipelines. Following Ref. [22], we instead consider the events with $\text{FAR} < 1 \text{ yr}^{-1}$ in at least one of the pipelines, which excludes some events from O3 and results in a catalog of 72 events. We compare results using the two catalogs in Appendix B.

III. RESULTS

In the following we report medians and 90% symmetric credible intervals for inferred posteriors.

A. Detection threshold

In Fig. 1 we display the posteriors inferred from Model 1. We find a detection threshold of $\tau = 9.2_{-0.4}^{+0.4}$. Using the joint likelihood in Eq. (4) we simultaneously infer the

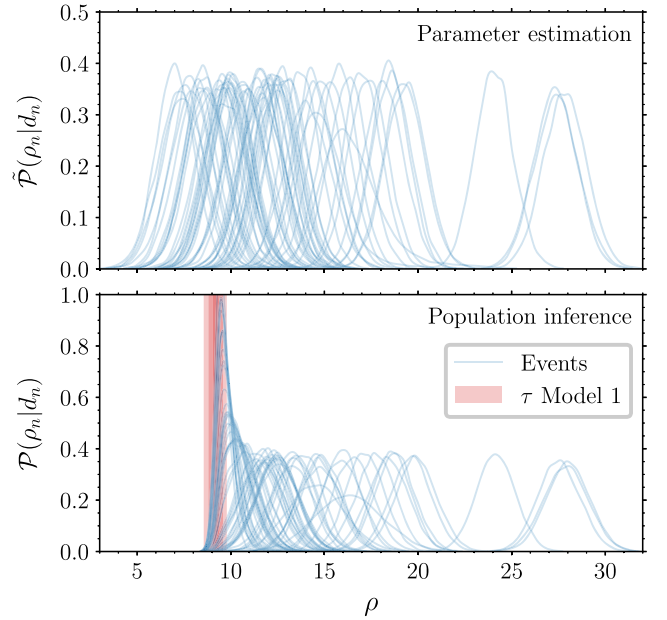


FIG. 1. The detection threshold τ inferred using Model 1. The shaded red bands show from darker to lighter the symmetric 50%, 90%, and 99% credible intervals. The blue distributions show the network optimal SNRs ρ_n as inferred by the original parameter estimation runs in the top row and by the joint population inference in the bottom row.

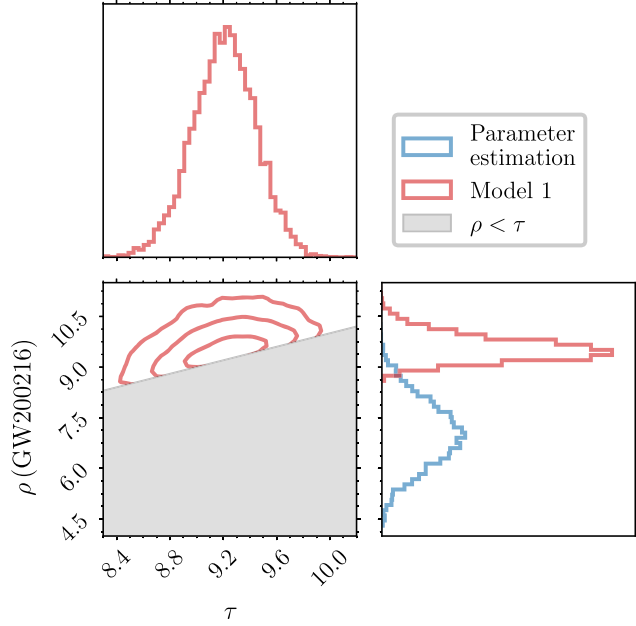


FIG. 2. The joint posterior distribution of the network optimal SNR ρ for the event with the lowest median SNR, GW200216_220804, and the detection threshold τ for Model 1. Red distributions show the joint population inference result, blue shows the original parameter estimation, and the gray region is excluded by the model. The panels along the diagonal show one-dimensional marginal posteriors while the lower left panel shows the two-dimensional marginal, with contours at the 50%, 90%, and 99% credible levels.

network optimal SNRs ρ_n for each event (equivalent to population-reweighted posteriors in marginal analyses as in, e.g., Ref. [47]). The population constraint for the strict detection threshold forces the tail of lower SNR events at $\rho \approx 4$ to be above threshold with $\rho \gtrsim 8$.

We further demonstrate this effect in Fig. 2, where we show the joint posterior for the threshold τ and the SNR for the event with the lowest median, GW200216_220804 (shortened to GW200216 in the plot for brevity). While originally $\rho = 7.0^{+1.7}_{-1.6}$ as inferred from parameter estimation samples which is largely in the region $\rho < \tau$ excluded by the assumed population model, the population inference places a constraint of $\rho = 9.5^{+0.7}_{-0.5}$, the two being consistent only at the 97% one-sided credible level.

B. Astrophysical distribution

We verify the power-law slope assumed in Model 1 by letting this also be a free parameter in Model 2. Our posteriors are displayed in Fig. 3. We find a measurement of the detection threshold only slightly broader than for Model 1 ($\tau = 9.3^{+0.4}_{-0.5}$) and a slope $\sigma = 4.2^{+0.8}_{-0.7}$ that is fully consistent with the astrophysical expectation of Eq. (8). In particular, we find that the simpler Model 1 is preferred over Model 2 with a Bayes factor of 7.7 as computed using the Savage-Dickey density ratio [48].

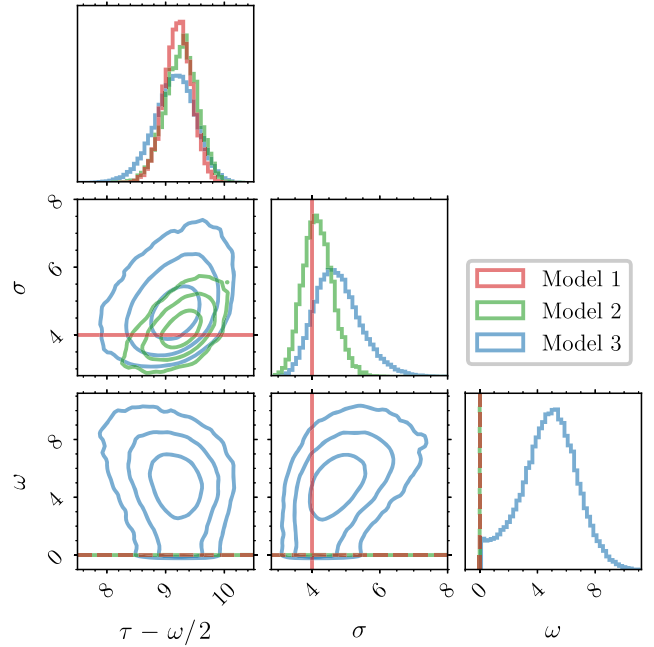


FIG. 3. Posterior distributions for the power-law slope σ of the intrinsic network optimal SNR population, the centroid $\tau - \omega/2$ of the detection threshold τ , and the detection probability width ω that parametrize Models 1 (red), 2 (green), and 3 (blue). Models 1 and 2 assume $\omega = 0$, and Model 1 assumes $\sigma = 4$; these constraints are displayed with vertical and horizontal lines for comparison. The diagonal panels display one-dimensional marginal posteriors and the lower left panels display two-dimensional marginals with 50%, 90%, and 99% credible levels.

C. Detectable distribution

We now test the detection threshold assumed by Models 1 and 2. When instead assuming that the detection probability is not a step function, we infer a width $\omega = 4.9^{+3.0}_{-3.7}$, as parametrized by Model 3 in Eq. (12). This suggests a preference for nonzero widths, though we also find nonzero support for $\omega = 0$ in Fig. 3. This broadening shifts the peak of the distribution to higher ρ , with $\tau = 11.6^{+1.4}_{-1.8}$ for Model 3, since the nonzero density in the detectable population can be accounted for not just by the power law component but also by the sigmoid detection probability, unlike for Models 1 and 2. The centroid of the threshold remains consistent between all three models, however, with $\tau - \omega/2 = 9.2^{+0.5}_{-0.6}$ for Model 3. This also results in a steeper but still consistent power-law slope, $\sigma = 4.7^{+1.3}_{-0.9}$, since it is additionally constrained at the other end of the distribution by the absence of events with $\rho \gtrsim 30$. The optimal SNR reflects only the (approximate) center of the expected distribution of observed SNRs. But beyond this, the catalog is constructed using a different ranking statistic, in this case the FAR, which does not directly translate into a strict threshold in other quantities, such as the network optimal SNR. The $\omega = 0$ slice of the




FIG. 4. Detectable distributions of the network optimal SNR ρ and associated posterior uncertainty inferred using Models 2 (green), 3 (blue), and F (orange). Solid lines indicate the posterior median and shaded regions the symmetric 90% credible intervals.

Model 3 posterior remains consistent with Models 1 and 2, as seen in Fig. 3.

The resulting posterior population distribution, computed by evaluating the prior models $\pi(\rho|\text{det}, \kappa, \lambda)$ at draws from the posterior $\mathcal{P}(\kappa, \lambda|\{d\})$, is displayed in Fig. 4. Compared to Model 2, the distribution inferred using Model 3 has larger uncertainties, suggesting the parametrization of the former may be overly restrictive. They are consistent at larger SNRs where they assume the same parametrization, but below the detection threshold inferred by Model 2 become inconsistent beyond the 90% uncertainty.

We verify the assumptions of Model 3 by placing a more flexible prior over the detectable population of ρ with Model F, whose inferred distribution is also displayed in Fig. 4 for comparison. The results for Models 3 and F are broadly consistent, peaking at the same location ($11.3^{+0.8}_{-0.8}$ for Model F, cf. above) and staying consistent below and above it within the 90% posterior uncertainty, though the latter prefers a broader distribution overall. With Model F we find a slower rise with a width $5.8^{+1.5}_{-1.3}$ below the peak of the distribution. The overall uncertainty in Model F is lower than that of Model 3 since, due to its over parametrization and resulting sparsity [49], many weight posteriors are prior driven. However, we use this model parametrized by a neural network only to broadly check that the functional forms assumed for the previous models are reasonable. Choosing a smaller network does not change the resulting posterior population distribution, but changing the shape of the priors to, e.g., uniform, result in slightly larger uncertainties around the peak of the distribution.




FIG. 5. Posteriors of the parameters for Models 2 (green) and 4 (blue). Both are characterized by a power law with slope σ and a strict detection threshold τ , shared between all three observing runs for Model 2 but separated into O1, O2, and O3 for Model 4 (τ_1 , τ_2 , and τ_3 , respectively). Diagonal panels show each one-dimensional posterior while lower left panels show two-dimensional marginals with 50%, 90%, and 99% credible regions.

D. Differences between observing runs

Finally, we consider another extension to Model 2 that still assumes a strict detection threshold, but takes different values for the three LVK observing runs (Model 4). The posterior for Models 2 and 4 are compared in Fig. 5. Assuming a shared power-law slope σ leaves its measurement unchanged, with $\sigma = 4.3^{+0.8}_{-0.7}$. The detection threshold for O3, $\tau_3 = 9.1^{+0.5}_{-0.5}$, is fully consistent with the shared thresholds in Model 1 and 2, implying the inferences for latter are primarily driven by the O3 events that make up the majority of the catalog. The thresholds for O1 and O2 are more uncertain, however, due to the lower number of events (3 and 7, respectively), with $\tau_1 = 10.5^{+2.1}_{-2.4}$ and $\tau_2 = 11.2^{+1.2}_{-1.4}$. These thresholds are consistent with each other and, while their posterior constraints do overlap with that for O3, they are typically larger. This means a signal passing the FAR cut in O3 could be included in the catalog with a lower equivalent network optimal SNR than it would in O1 or O2, given the assumed representative PSDs.

IV. DISCUSSION

In both inference and modeling settings, GW detectability is very often unphysically thresholded using the (network) optimal SNR, a quantity that depends on the data only through estimation of the detector noise properties and moreover on the true binary source parameters. This is in

contrast to real searches that depend only on the observed GW data. By simultaneously parametrizing the detection probability and intrinsic distribution of a selected ranking statistic such as the network optimal SNR, we showed that corresponding detection thresholds can be calibrated against observational catalogs that were assembled using real search pipelines. In other words, we inferred the threshold of an unphysical model that best reproduces the physical model.

However, we emphasize that the only way to fully reproduce search results is to process simulated sources with software injections and there is no guarantee that, e.g., an optimal SNR threshold will filter sources across the parameter space in the same way. Such thresholds make strong assumptions on the dependence of detectability on the source parameters and reference PSD through the SNR calculation that searches do not. In reality, detection also depends on other parameters that this method does not account for, such as the background rate which is also inhomogeneous across the binary parameter space (e.g., more massive binaries produce shorter signals which are more easily mistaken for noise transients, and vice versa). Because we do not know the true source properties of search triggers, the only way to benchmark unphysical selection criteria is to construct the distribution of physical ranking statistics given the unphysical one, which would anyway require performing software injections in many noise realizations. The primary purpose of this work, then, is not to replace injection campaigns where accurate reproduction of search pipelines is necessary (e.g., population inference), but to infer appropriate thresholds for unphysical models when efficiency is paramount (e.g., processing large population synthesis simulations into mock catalogs). Our parametrization also provides a convenient functional form for the detection probability tailored to current observations.

Other approaches that model detection with some approximation for the noise properties allow selection effects to be computed while depending only on observables as opposed to the true source properties [14,31,32]. If one were to construct observational catalogs in this way, by thresholding on, e.g., the observed SNR, it is trivial to perform self-consistent downstream analyses, such as population inference. However, the assumptions that produce (semi) analytic detection estimates, though reasonable, are not exact. Moreover, catalogs are more commonly assembled using statistics that take account of information beyond just the SNR, such as the FAR [22] or p_{astro} [6,7]. When applied to such catalogs, the above approach therefore still requires the choice of an approximating threshold, because the ranking statistics used to assemble the catalog and then to estimate detection biases for, e.g., population inference or simulated reconstruction of observed events, are different. The method we present is general and can be used to infer catalog-calibrated detection thresholds no matter the chosen approximating model.

We find a marked agreement between our inferred O3-only network optimal SNR threshold with that found by optimization against injections [30]— $9.1_{-0.5}^{+0.5}$ and 9.2, respectively (see their Fig. 6)—despite systematic differences between the two. Firstly, our catalog is composed of real signals that are compared to a waveform model, whereas injections are both generated and recovered using waveform models. Secondly, while both methods took fixed PSDs to calculate SNRs, Ref. [30] also uses those to draw noise realizations for injections, whereas there is some variability in real observations. Thirdly, the observational catalog is contaminated at some level by noise transients [22], which is not true of injection campaigns. Our posterior measurements are therefore conditioned on this systematic waveform and noise mis-modeling (though we expect the systematic uncertainty to be subdominant with respect to the statistical uncertainty), meaning downstream use of the same ranking statistic and thresholds are self consistent.

We encoded the full vector of binary source properties into the population distribution of a single representative variable, here taken to be the network optimal SNR. Even if the intrinsic population is uncorrelated between source parameters, this is not true of the observed population. A full joint inference of all binary source parameters along with global parameters that characterize the detectable and intrinsic populations will be an interesting approach for future work.

ACKNOWLEDGMENTS

We thank Salvatore Vitale, Jack Heinzl, and Maya Fishbach for helpful comments. M.M. is supported by LIGO Laboratory through the National Science Foundation Award PHY-1764464. C.J.M. acknowledges the support of the UK Space Agency Grant No. ST/V002813/1. D.G. is supported by ERC Starting Grant No. 945155–GWmining, Cariplo Foundation Grant No. 2021-0555, MUR PRIN Grant No. 2022-Z9X4XS, MSCA Fellowship No. 101064542–StochRewind, MSCA Fellowship No. 101149270–ProtoBH, Leverhulme Trust Grant No. RPG-2019-350, and the ICSC National Research Centre funded by NextGenerationEU. Computational work was performed at CINECA with allocations through INFN and Bicocca. This research has made use of data or software obtained from the Gravitational Wave Open Science Center ([50]), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation, as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO

was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan.

APPENDIX A: MARGINAL POSTERIOR

Instead of simultaneously inferring κ , λ , and $\{\rho_n\}$, if we are uninterested in the event-level parameters we can marginalize the likelihood over each to obtain

$$\mathcal{L}(\{\det_n, d_n\}|\kappa, \lambda) \propto P(\det|\kappa, \lambda)^{-N} \times \prod_{n=1}^N \int P(\det_n|\rho_n, \kappa) \mathcal{L}(d_n|\rho_n) \pi(\rho_n|\lambda) d\rho_n. \quad (\text{A1})$$

Rewriting the individual-event likelihoods in terms of the original parameter-estimation posteriors, we have

$$\mathcal{L}(\{\det_n, d_n\}|\kappa, \lambda) = P(\det|\kappa, \lambda)^{-N} \times \prod_{n=1}^N \frac{1}{N_n} \sum_{i=1}^{N_n} P(\det_n|\rho_{ni}, \kappa) \frac{\pi(\rho_{ni}|\lambda)}{\tilde{\pi}_n(\rho_{ni})}, \quad (\text{A2})$$

where each integral is approximated as a Monte Carlo average over the posterior samples $\{\rho_{ni}\}_{i=1}^{N_n} \sim \tilde{\mathcal{P}}(\rho_n|d_n)$. The advantage of this approach is that we do not need to use fitted posteriors for the individual events as above, instead relying directly on the posterior samples. We do still require the original prior densities, but since they do not depend on κ or λ the prior densities evaluated at the posterior samples, $\{\{\tilde{\pi}_n(\rho_{ni})\}_{i=1}^{N_n}\}_{n=1}^N$, can be computed once and stored ahead of time.

As an example, we show the Model 3 posteriors resulting from the full hierarchical distribution of Eq. (4) and the marginal likelihood of Eq. (A2) in Fig. 6. The two are almost identical, implying our assumed single-event truncated-normal posteriors are valid. The only slight discrepancy between them is that the posterior for the detection-threshold width ω inferred with the joint likelihood has nonzero support at $\omega = 0$, whereas that for the marginal likelihood does not. This is likely because the truncated normal posteriors have support for ρ_n across the positive real line whereas the parameter-estimation results are subject to finite-sampling effects.

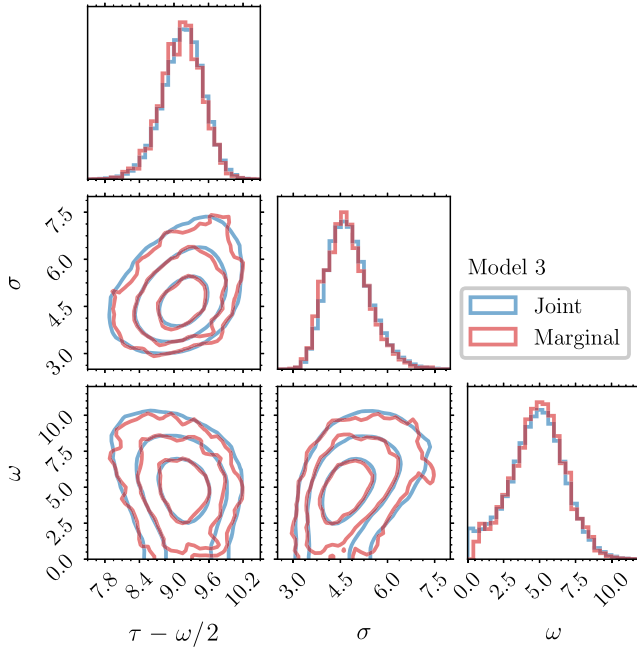


FIG. 6. Posterior distributions for the parameters of Model 3 as inferred in a fully hierarchical analysis (blue) and a numerically marginalized one (red). Diagonal panels contain histograms for each parameter individually while the others contain two-dimensional posteriors with contours at the 50%, 90%, and 99% credible levels.




FIG. 7. Posterior distributions for the parameters of Model 4 inferred using catalogs where events are selected based on a threshold of $\text{FAR} < 1 \text{ yr}^{-1}$ (blue) or $p_{\text{astro}} > 0.5$ (red). One- and two-dimensional marginal posteriors are given in the diagonal and lower-left panels, respectively, with the latter containing credible regions at the 50%, 90%, and 99% levels.

APPENDIX B: CATALOG DEFINITION

The results in Sec. III are all produced using a catalog where events are selected if they satisfy $\text{FAR} < 1 \text{ yr}^{-1}$, as in Ref. [22]. In other words, we are inferring the parameters κ and λ of the detectable distribution of network optimal SNRs ρ that best reproduces this FAR threshold. Another common catalog threshold is for events to be more likely of astrophysical origin than otherwise, i.e., $p_{\text{astro}} > 0.5$ [7]. We assess the impact of these two choices in Fig. 7 by running the inference with Model 4 on each catalog. In the catalog defined by $p_{\text{astro}} > 0.5$ there are several events whose network optimal SNR posteriors clearly do not satisfy the assumed truncated-normal approximation employed in Sec. II B. We therefore use the marginal likelihood from Eq. (A2) in which the finite set of

parameter-estimation samples are used directly, thus avoiding approximating the posterior distributions.

The only events that do not make the $\text{FAR} < 1 \text{ yr}^{-1}$ threshold compared to the larger $p_{\text{astro}} > 0.5$ catalog are from O3. Therefore, the parameters τ_1 and τ_2 that define the detection thresholds for O1 and O2, respectively, are unchanged with respect to this choice of catalog. Due to the inclusion of more low-SNR events from O3, the corresponding threshold τ_3 is lowered from $\tau_3 = 9.1^{+0.5}_{-0.5}$ to $\tau_3 = 8.2^{+0.4}_{-0.4}$. The power-law slope is also mostly influenced by the O3 events in the catalog due to the larger number of them, and becomes $\sigma = 4.0^{+0.7}_{-0.6}$ for the $p_{\text{astro}} > 0.5$ catalog, compared to $\sigma = 4.2^{+0.9}_{-0.7}$ for $\text{FAR} < 1 \text{ yr}^{-1}$. In both cases the uncertainties are slightly reduced due to the addition of 18 events.

-
- [1] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
- [2] F. Acernese *et al.* (VIRGO Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
- [3] T. Akutsu *et al.* (KAGRA Collaboration), *Nat. Astron.* **3**, 35 (2019).
- [4] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 031040 (2019).
- [5] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **11**, 021053 (2021).
- [6] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **109**, 022001 (2024).
- [7] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), *Phys. Rev. X* **13**, 041039 (2023).
- [8] G. S. Davies, T. Dent, M. Tápai, I. Harry, C. McIsaac, and A. H. Nitz, *Phys. Rev. D* **102**, 022004 (2020).
- [9] K. Cannon *et al.*, *SoftwareX* **14**, 100680 (2021).
- [10] F. Aubin *et al.*, *Classical Quantum Gravity* **38**, 095004 (2021).
- [11] S. Klimenko *et al.*, *Phys. Rev. D* **93**, 042004 (2016).
- [12] L. S. Finn and D. F. Chernoff, *Phys. Rev. D* **47**, 2198 (1993).
- [13] D. F. Chernoff and L. S. Finn, *Astrophys. J. Lett.* **411**, L5 (1993).
- [14] R. Essick and M. Fishbach, *Astrophys. J.* **962**, 169 (2024).
- [15] C. Talbot and E. Thrane, *Astrophys. J.* **927**, 76 (2022).
- [16] D. Gerosa, G. Pratten, and A. Vecchio, *Phys. Rev. D* **102**, 103020 (2020).
- [17] K. W. K. Wong, K. K. Y. Ng, and E. Berti, [arXiv:2007.10350](https://arxiv.org/abs/2007.10350).
- [18] M. Mould, D. Gerosa, and S. R. Taylor, *Phys. Rev. D* **106**, 103013 (2022).
- [19] C. E. A. Chapman-Bird, C. P. L. Berry, and G. Woan, *Mon. Not. R. Astron. Soc.* **522**, 6043 (2023).
- [20] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **882**, L24 (2019).
- [21] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **913**, L7 (2021).
- [22] R. Abbott *et al.* (KAGRA, VIRGO, and LIGO Scientific Collaborations), *Phys. Rev. X* **13**, 011048 (2023).
- [23] M. Mould and D. Gerosa, *Phys. Rev. D* **105**, 024076 (2022).
- [24] J. W. Barrett, S. M. Gaebel, C. J. Neijssel, A. Vigna-Gómez, S. Stevenson, C. P. L. Berry, W. M. Farr, and I. Mandel, *Mon. Not. R. Astron. Soc.* **477**, 4685 (2018).
- [25] J. Riley *et al.* (COMPAS Team), *Astrophys. J. Suppl. Ser.* **258**, 34 (2022).
- [26] F. S. Broekgaarden *et al.*, *Mon. Not. R. Astron. Soc.* **516**, 5737 (2022).
- [27] M. Dominik, E. Berti, R. O’Shaughnessy, I. Mandel, K. Belczynski, C. Fryer, D. E. Holz, T. Bulik, and F. Pannarale, *Astrophys. J.* **806**, 263 (2015).
- [28] V. Baibhav, E. Berti, D. Gerosa, M. Mapelli, N. Giacobbo, Y. Bouffanais, and U. N. Di Carlo, *Phys. Rev. D* **100**, 064060 (2019).
- [29] LIGO Scientific, Virgo, and KAGRA Collaborations, [10.5281/zenodo.7890437](https://zenodo.org/record/7890437) (2023).
- [30] R. Essick, *Phys. Rev. D* **108**, 043011 (2023).
- [31] M. Fishbach, W. M. Farr, and D. E. Holz, *Astrophys. J. Lett.* **891**, L31 (2020).
- [32] A. M. Farah, B. Edelman, M. Zevin, M. Fishbach, J. M. Ezquiaga, B. Farr, and D. E. Holz, *Astrophys. J.* **955**, 107 (2023).
- [33] I. Mandel, W. M. Farr, and J. R. Gair, *Mon. Not. R. Astron. Soc.* **486**, 1086 (2019).
- [34] S. Vitale, D. Gerosa, W. M. Farr, and S. R. Taylor, in *Handbook of Gravitational Wave Astronomy* (Springer, New York, 2020).
- [35] D. Phan, N. Pradhan, and M. Jankowiak, [arXiv:1912.11554](https://arxiv.org/abs/1912.11554).
- [36] G. Pratten *et al.*, *Phys. Rev. D* **103**, 104056 (2021).
- [37] S. A. Usman *et al.*, *Classical Quantum Gravity* **33**, 215004 (2016).
- [38] G. Ashton *et al.*, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [39] P. Virtanen *et al.*, *Nat. Methods* **17**, 261 (2020).

- [40] B. F. Schutz, *Classical Quantum Gravity* **28**, 125023 (2011).
- [41] H.-Y. Chen and D. E. Holz, [arXiv:1409.0522](https://arxiv.org/abs/1409.0522).
- [42] B. P. Abbott *et al.* (KAGRA, LIGO Scientific, Virgo, and VIRGO Collaborations), *Living Rev. Relativity* **21**, 3 (2018).
- [43] N. De Cao, I. Titov, and W. Aziz, in *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence* (2019), [arXiv:1904.04676](https://arxiv.org/abs/1904.04676).
- [44] M. Mould, github.com/mdmould/shallow, Zenodo, [10.5281/zenodo.7873398](https://zenodo.org/record/7873398) (2023).
- [45] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *SoftwareX* **13**, 100658 (2021).
- [46] R. Abbott *et al.* (KAGRA, VIRGO, and LIGO Scientific Collaborations), *Astrophys. J. Suppl. Ser.* **267**, 29 (2023).
- [47] C. J. Moore and D. Gerosa, *Phys. Rev. D* **104**, 083008 (2021).
- [48] J. M. Dickey, *Ann. Math. Stat.* **42**, 204 (1971).
- [49] T. Hoefler, D. Alistarh, T. Ben-Nun, N. Dryden, and A. Peste, *J. Mach. Learn. Res.* **22** (2021).
- [50] <https://gwosc.org/>.