

One to many: comparing single gravitational-wave events to astrophysical populations

Matthew Mould ^{1*}, Davide Gerosa ^{2,3,1}, Marco Dall’Amico ^{4,5}, Michela Mapelli ^{4,5,6}

¹*School of Physics and Astronomy & Institute for Gravitational Wave Astronomy, University of Birmingham, Birmingham, B15 2TT, UK*

²*Dipartimento di Fisica ‘G. Occhialini’, Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, 20126 Milano, Italy*

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

⁴*Physics and Astronomy Department Galileo Galilei, University of Padova, Vicolo dell’Osservatorio 3, I-35122 Padova, Italy*

⁵*INFN-Padova, Via Marzolo 8, I-35131 Padova, Italy*

⁶*Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik, Albert-Ueberle-Str. 2, D-69120 Heidelberg, Germany*

8 September 2023

ABSTRACT

Gravitational-wave observations have revealed sources whose unusual properties challenge our understanding of compact-binary formation. Inferring the formation processes that are best able to reproduce such events may therefore yield key astrophysical insights. A common approach is to count the fraction of synthetic events from a simulated population that are consistent with some real event. Though appealing owing to its simplicity, this approach is flawed because it neglects the full posterior information, depends on an ad-hoc region that defines consistency, and fails for high signal-to-noise detections. We point out that a statistically consistent solution is to compute the posterior odds between two simulated populations, which crucially is a relative measure, and show how to include the effect of observational biases by conditioning on source detectability. Applying the approach to several gravitational-wave events and simulated populations, we assess the degree to which we can conclude model preference not just between distinct formation pathways but also between subpopulations within a given pathway.

Key words: gravitational waves – black hole mergers – methods: statistical

1 INTRODUCTION

In the first three observing runs (Abbott et al. 2019, 2021c,a,b) of the LIGO (Aasi et al. 2015), Virgo (Acernese et al. 2015), and KAGRA (Kagra Collaboration et al. 2019) gravitational-wave (GW) interferometers there have been nearly 100 detection candidates of transient signals originating from stellar-mass compact-object binaries. Despite the growing catalogue of observations, there are still new events whose unusual or unique properties make them especially interesting. For example: GW200129 is a binary black hole (BH) merger whose remnant is constrained to have a large gravitational kick (Varma et al. 2022) and which may (or may not, Payne et al. 2022) exhibit evidence for relativistic spin precession (Hannam et al. 2022); GW190521 consists of two BHs with particularly large component masses (Abbott et al. 2020c) and features possible signatures of orbital eccentricity (Romero-Shaw et al. 2020; Gayathri et al. 2022); GW190412 confidently contains two stellar-mass BHs but with clearly unequal masses (Abbott et al. 2020b); and for GW190814 (Abbott et al. 2020d) the lighter component is either an unusually heavy neutron star

or light BH, making its formation with very unequal masses difficult to explain.

The atypical properties of these events have important implications for their progenitors, as astrophysical formation environments must facilitate the production of such sources. In the case of GW190521, for example, several possibilities have been proposed to explain its formation (Abbott et al. 2020e; Romero-Shaw et al. 2020; Fishbach & Holz 2020; Nitz & Capano 2021; Bustillo et al. 2021; Olsen et al. 2021; Estellés et al. 2022; Gayathri et al. 2022; Gamba et al. 2023). This event is particularly problematic because, while it confidently has a large total mass, the short signal duration leads to ambiguity in the nature of the source (Biscoveanu et al. 2021; Romero-Shaw et al. 2023).

Given the vast array of possible – if not plausible – formation scenarios, a suitable metric is needed to assess the relative degree of consistency between GW measurements and different models of formation. A common method of building such models is to simulate a population of mergers. One may therefore be interested in investigating whether or not such simulations are capable of reproducing populations that are consistent with a given GW event, and conversely whether a given event is consistent with a synthetic population (see, e.g., Gerosa et al. 2020b; Mapelli et al. 2021; Dall’Amico et al.

* mmould@star.sr.bham.ac.uk

2021; Tagawa et al. 2021; Zevin & Bavera 2022; Broekgaarden et al. 2022; Zhang et al. 2023).

But how does one consistently compare a *distribution* of simulated sources with a *single* observed event?

A typical approach proceeds by (i) constructing a synthetic population of mergers (e.g., through population synthesis simulations), (ii) defining a region in the parameter posterior of a GW event of interest [e.g., 90% credible intervals (CIs) in masses and spins], and (iii) counting the number of simulated mergers inside this region. A high fraction of simulated samples lying inside the chosen region is taken as an indication that the detected GW event is consistent with the simulated formation scenario. Conversely, one could define the confidence region in terms of the simulated mergers and count the number of posterior samples from the event inside that region.

However, this method is at best approximate because: for simplicity, the confidence region is typically assembled heuristically as the product of marginalized one-dimensional (1D) CIs, and the result depends on the choice of quantile (e.g., 90%); density gradients inside this region are neglected; posterior support outside the region is neglected; the influence of the original parameter estimation (PE) prior is not addressed; and there is an ambiguity in which of the two comparisons to make, as they do not agree.

These issues become especially clear in the limit of high signal-to-noise ratios (SNRs). For a high-SNR detection the posterior distribution is narrow and, therefore, the fraction of simulated mergers inside any region of posterior support will be arbitrarily close to zero, even if the posterior lies entirely within the extremal simulated sources. Conversely, since the posterior support is bounded well within the range of simulated sources, the fraction of posterior samples inside any given region of the simulated population will be arbitrarily close to unity.

In this work we point out a statistically consistent alternative provided by working entirely within the framework of Bayesian PE and model comparison. In short, the posterior odds can be used when comparing the posterior of a GW event with a model of formation, the latter acting as a Bayesian prior. An important point is that one should not ask to what extent a GW event matches a particular model of formation in absolute terms, and should instead quantify how much more or less likely one model is over another. The statistics required to perform Bayesian PE and model comparison (and hierarchical inference) have already been extensively reviewed elsewhere (see, e.g., Thrane & Talbot 2019; Mandel et al. 2019; Vitale et al. 2022). We summarize the salient details and point out that the analysis required is simple and, in the context of GW astronomy, readily performed with publicly available data products combined with a predictive astrophysical model.

1.1 Practical summary

One first identifies a GW event of interest, perhaps because it is an outlier with respect to other observations (Fishbach et al. 2020; Essick et al. 2022; Romero-Shaw et al. 2022). Then, irrespective of the remaining catalogue, we wish to compare models of formation for this individual event.

Performing a comparison between astrophysical population models – say, A and B – and deciding which is more consistent with the observed GW event with parameters θ (e.g., masses,

spins, redshift, etc.) requires computing the likelihoods that the GW data are drawn from those populations. The required steps are:

- (i) Simulate an astrophysical set of mergers from model A ; $\{\theta_{Aj}\}_{j=1}^{N_A}$.
- (ii) Construct a probability density estimate of the simulated merger source properties; $\pi(\theta|A)$.
- (iii) Evaluate whether each simulated source is detectable or not; $\{P(\text{det}|\theta_{Aj})\}_{j=1}^{N_A}$.
- (iv) Repeat steps (i)–(iii) for model B .
- (v) Download the PE samples for the GW event of interest; $\{\theta_i\}_{i=1}^N$. If not already known, construct a density estimate for the uninformative PE prior; $\pi(\theta|U)$.
- (vi) Compute the detectability-weighted Bayes factor

$$\mathcal{D}_{A/B} = \frac{P(\text{det}|B)}{P(\text{det}|A)} \mathcal{B}_{A/B}, \quad (1)$$

where, by summing over simulated sources, the average detection probability for population A is

$$P(\text{det}|A) = \frac{1}{N_A} \sum_{j=1}^{N_A} P(\text{det}|\theta_{Aj}) \quad (2)$$

and, by summing over posterior samples for the real event, the Bayes factor is

$$\mathcal{B}_{A/B} = \frac{\sum_{i=1}^N \pi(\theta_i|A)/\pi(\theta_i|U)}{\sum_{i=1}^N \pi(\theta_i|B)/\pi(\theta_i|U)}. \quad (3)$$

If $\mathcal{D}_{A/B} > 1$ (< 1) then, assuming equal prior odds, the detected GW event is a posteriori more consistent with population A (B) – the key statement of interest. Since this posterior odds (or Bayes factor) is inherently a relative measure, considering a single model is not possible; the absence of a second astrophysical model implies a comparison to the original PE prior.

In Sec. 2, we formally define the heuristic fraction used previously and demonstrate why it is unsuitable. We summarize Bayesian model comparison in the present context of GW astronomy and derive the above quantities step by step. In Sec. 3, we apply this formalism to several stellar-mass mergers detected by LIGO/Virgo and assess to what degree we can conclude that they are consistent with various models of formation. We end on some final thoughts in Sec. 4. Code to reproduce these results is publicly available (github.com/mdmould/popodds; Mould 2023).

2 BAYESIAN MODEL COMPARISON

Using Bayes' theorem, the measured properties of a GW event are given by the posterior distribution

$$p(\theta|\mathbf{d}, U) = \frac{\mathcal{L}(\mathbf{d}|\theta)\pi(\theta|U)}{\mathcal{Z}(\mathbf{d}|U)}, \quad (4)$$

where $\mathcal{L}(\mathbf{d}|\theta)$ is the likelihood that the vector of waveform parameters θ (e.g., masses, spins, distance, sky position, etc.) produce the GW data \mathbf{d} , $\pi(\theta|U)$ represents our typically uninformative prior belief on what the waveform parameters could be, and

$$\mathcal{Z}(\mathbf{d}|U) = \int \mathcal{L}(\mathbf{d}|\theta')\pi(\theta'|U)d\theta' \quad (5)$$

is the marginal likelihood – or evidence – that normalizes the posterior by ensuring $\int p(\boldsymbol{\theta}|\mathbf{d}, U) d\boldsymbol{\theta} \equiv 1$ and represents the probability that the chosen PE prior produces the observed data.

We only explicitly condition prior-dependent terms on the uninformative model U to represent their dependence on the shape of the chosen prior distribution. The result also depends on, e.g., the waveform model that in turn determines those parameters $\boldsymbol{\theta}$ under consideration. We omit these extra model conditionals that we consider fixed to avoid unnecessarily lengthy expressions.

Given such inference on the binary properties $\boldsymbol{\theta}$, the question is whether the measurement $p(\boldsymbol{\theta}|\mathbf{d}, U)$ is consistent with a particular model of binary formation. We describe such approaches in the following sections.

2.1 Heuristic approximation

We begin with the heuristic fraction that is often used to assess the consistency of GW events with populations models. It is computed by counting the number of sources simulated from some astrophysical model that lie inside a chosen region R of the measured properties $p(\boldsymbol{\theta}|\mathbf{d}, U)$. Such simulations can be phrased as a finite set of draws from an unknown underlying astrophysical distribution $\pi(\boldsymbol{\theta}|A)$ that describes the result of the simulation. The fraction inside R is

$$f_A(R) = \frac{N_A(R)}{N_A} \approx \int \pi(\boldsymbol{\theta}|A) d\boldsymbol{\theta} = \int \mathcal{I}(\boldsymbol{\theta}, R) \pi(\boldsymbol{\theta}|A) d\boldsymbol{\theta}, \quad (6)$$

where N_A is the total number of simulated sources, $N_A(R)$ is the number of those sources that fall within the targeted region R of the posterior, and $\mathcal{I}(\boldsymbol{\theta}, R)$ is an indicator function equal to unity if $\boldsymbol{\theta} \in R$ and zero otherwise. Large values of $f_A(R)$ can be taken as an indication that the event posterior is consistent with the simulated model, though the value at which one may define significance is arbitrary. A similar conclusion holds conversely for the fraction of GW posterior samples lying inside a parameter region defined in terms of the simulated sources; this fraction is not usually used as population models are typically broader than GW event posteriors, meaning the fraction will always be close to unity.

The region R is often constructed heuristically by joining the 90% CIs of each 1D marginalization of the posterior. However, this choice is complicated by the addition of higher dimensions and is not necessarily a good approximation of the true full-dimensional 90% credible region. Additionally, as one includes more parameters in the comparison while keeping the simulation size fixed, the curse of dimensionality means that fewer simulated sources will fall within the chosen region, leading to the naive conclusion that the model is harder to reconcile with the data.

We illustrate the setup of this heuristic fraction in Fig. 1. We take different two-dimensional (2D) normal distributions in (x, y) as both the posterior and simulated population model. In red we highlight the heuristic region R assembled as the product of the 1D 90% posterior CIs; note that this region does not match the joint-dimensional 90% credible region of the posterior. Sources simulated from the fiducial model are plotted as blue markers, with those within the heuristic region highlighted in bold. The heuristic fraction is found as the fraction of the simulated points that lie inside the red square. This computation neglects any posterior support outside the

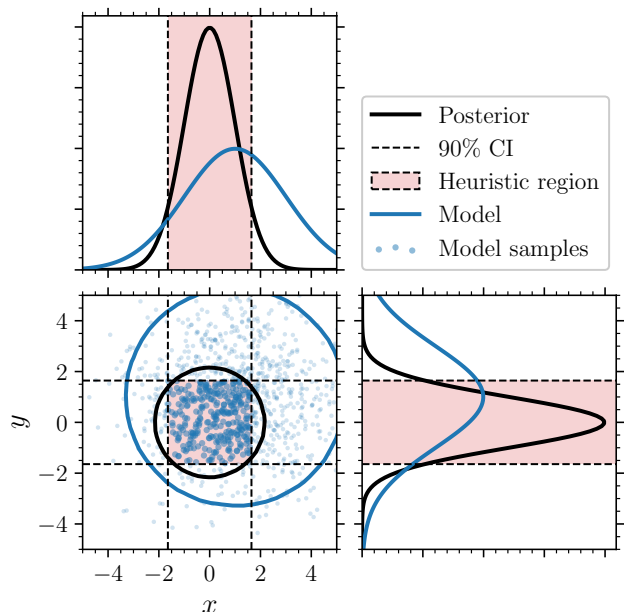


Figure 1. Illustration of the heuristic fraction. The posterior from a single observation (solid black) and simulated fiducial model (solid blue) are 2D normal distributions. In the lower left panel are their 90% credible regions while the diagonal panels display their 1D marginals. The dashed lines are located at the boundaries of the symmetric 90% CIs of the marginalized posteriors, while the red region highlights the heuristic region assembled as the product of the individual intervals. Blue points represent sources simulated from the model, with those inside the red region being counted for the heuristic fraction.

red region and also neglects variation of the posterior density inside it. One can see that if the posterior were to become more precise (due to, e.g., a higher SNR), the red region would shrink and fewer simulated points lie inside it, even though the posterior is fully enclosed within the simulated model. Additionally, the heuristic fraction is an absolute measure for a single model in that the influence of the original prior on the posterior is not accounted for. As such, it is not clear when a value of $f_A(R)$ should be considered significant.

2.2 Bayes factors

We now seek a statistically consistent alternative to the heuristic fraction. Namely, one can perform Bayesian model selection to find which models are preferred relative to others.

Starting from the GW event posterior in Eq. (4), we can reconstruct the posterior that would have been inferred had we imposed our informative astrophysical model A as a prior by re-weighting the posterior that we do have. Since the likelihood is unchanged, the new posterior may be written as

$$p(\boldsymbol{\theta}|\mathbf{d}, A) = p(\boldsymbol{\theta}|\mathbf{d}, U) \frac{\pi(\boldsymbol{\theta}|A) \mathcal{Z}(\mathbf{d}|U)}{\pi(\boldsymbol{\theta}|U) \mathcal{Z}(\mathbf{d}|A)}. \quad (7)$$

Given samples from the original posterior $p(\boldsymbol{\theta}|\mathbf{d}, U)$, the new posterior is constructed using those samples with weights proportional to the prior ratio $\pi(\boldsymbol{\theta}|A)/\pi(\boldsymbol{\theta}|U)$.

Since the posterior is normalized, integrating over $\boldsymbol{\theta}$ and rearranging yields the ratio of evidences – also called the

Bayes factor – that quantifies the relative likelihood between the two prior models:

$$\mathcal{B}_{A/U} = \frac{\mathcal{Z}(\mathbf{d}|A)}{\mathcal{Z}(\mathbf{d}|U)} = \int p(\boldsymbol{\theta}|\mathbf{d}, U) \frac{\pi(\boldsymbol{\theta}|A)}{\pi(\boldsymbol{\theta}|U)} d\boldsymbol{\theta}. \quad (8)$$

A value of $\mathcal{B}_{A/U} > 1$ (< 1) implies that the astrophysical model is more (less) likely to produced the GW data \mathbf{d} than the original PE prior, with values further from unity indicating a stronger preference; a typical scale indicating a decisive conclusion is $|\ln \mathcal{B}_{A/U}| \gtrsim 5$ (Jeffreys 1939). One may notice that this Bayes factor is proportional to the per-event term entering the usual GW population likelihood (Thrane & Talbot 2019; Mandel et al. 2019; Vitale et al. 2022), though here we consider population models to be fixed. By reusing the posterior inferred under the default uninformative prior in this construction, we do not need to re-evaluate the computationally expensive likelihood $\mathcal{L}(\mathbf{d}|\boldsymbol{\theta})$ or perform new PE runs.

We describe how to compute the Bayes factor in practice in Appendix A and how to handle changes of variables in Appendix B.

Unlike the heuristic fraction, the Bayes factor is an explicit comparative measure between two models – the original uninformative prior, U , and the new informative one, A . By comparison with Eq. (8), we see that the heuristic fraction $f_A(R)$ in Eq. (6) has the form of a Bayes factor with $p(\boldsymbol{\theta}|\mathbf{d}, U)/\pi(\boldsymbol{\theta}|U)$ replaced with the indicator function $\mathcal{I}(\boldsymbol{\theta}, R)$. To be valid, this would require the posterior $p(\boldsymbol{\theta}|\mathbf{d}, U)$ to have no support outside the region R and be equal to the original prior $\pi(\boldsymbol{\theta}|U)$ inside R . One case this holds is when the posterior only recovers the prior and the region R is selected as the entire prior volume – but this is clearly not true in general. Equation (6) also has the form of an evidence as in Eq. (5) with the likelihood $\mathcal{L}(\mathbf{d}|\boldsymbol{\theta})$ replaced with the indicator function (also an invalid approximation). One could therefore consider the relative fraction

$$\mathcal{F}_{A/U}(R) = \frac{f_A(R)}{f_U(R)} = \frac{\int \mathcal{I}(\boldsymbol{\theta}, R) \pi(\boldsymbol{\theta}|A) d\boldsymbol{\theta}}{\int \mathcal{I}(\boldsymbol{\theta}, R) \pi(\boldsymbol{\theta}|U) d\boldsymbol{\theta}}, \quad (9)$$

as this is a comparative measure that has a similar interpretation to the Bayes factor with a natural scale: $\mathcal{F}_{A/U}(R) > 1$ (< 1) implies A is favoured (disfavoured) with respect to U over the heuristically chosen region R . However, this relative heuristic fraction still suffers from the other issues enumerated in Sec. 2.1. We consider a simple example and compare the heuristic fraction, the relative fraction, and the Bayes factor in Appendix C.

2.3 Posterior odds

A comparison can be made between two different informative models – say, A and B – from neither of which we have inferred a posterior. The Bayes factor between A and B is just the ratio of the two Bayes factors with respect to the original uninformative prior U :

$$\mathcal{B}_{A/B} = \frac{\mathcal{Z}(\mathbf{d}|A)}{\mathcal{Z}(\mathbf{d}|B)} = \frac{\mathcal{B}_{A/U}}{\mathcal{B}_{B/U}}. \quad (10)$$

Note, then, that while the Bayes factor between A and B does not depend on the original prior, its computation requires us to explicitly divide it out to compute $\mathcal{B}_{A/U}$ and $\mathcal{B}_{B/U}$ using Eq. (8) because it is the only one for which we have

inferred a posterior. Equation (10) has the advantage of being independent of the normalization of the original PE prior, however, which only enters through $\mathcal{Z}(\mathbf{d}|U)$. This fact is useful when only the shape and not the normalization of $\pi(\boldsymbol{\theta}|U)$ is known. For example, in the context of GW astronomy one may choose uniform priors in detector-frame component masses but place limits (which may vary between events and waveform models) on the chirp mass and mass ratio (Abbott et al. 2023), resulting in a prior whose shape within the cut remains flat but whose overall normalization is altered. This is irrelevant for inference but crucial for comparisons.

Finally, we arrive at the metric we sought: when comparing Bayesian models one may employ the posterior odds,

$$\mathcal{O}_{A/B} = \frac{p(A|\mathbf{d})}{p(B|\mathbf{d})} = \frac{\pi(A)}{\pi(B)} \mathcal{B}_{A/B}, \quad (11)$$

where $\pi(A)$ represents our prior on the hypothesis that A is the correct model generating GW events and the ratio $\pi(A)/\pi(B)$ is the prior odds, i.e., how much weight we a priori assign to model A over B . In many situations we are unable to reasonably estimate such priors and instead default to equal prior odds, in which case $\mathcal{O}_{A/B} = \mathcal{B}_{A/B}$ and the posterior odds and Bayes factor coincide. In this case, comparing two astrophysical models is typically more useful as the PE prior may often be preferred owing to its broad shape, even though it is not a realistic model of formation; this is where the prior odds would downweight preference for the PE prior if it could be reasonably chosen.

However, we stress that posterior measurements – in particular, the posterior odds – are the Bayesian metrics that should be used in these circumstances. By defaulting to the Bayes factor, we must be mindful of the implicit assumption of equal prior odds and that an alternative conclusion may be drawn if we impose different model priors.

2.4 Astrophysical populations and selection effects

So far we have not considered selection effects – the fact that detected GW events are subject to observational biases. This means that certain signals are more easily detected than others, e.g., from sources which are closer to the detectors. When considering an observed GW event, we know that this source must belong to the observable population – not just the intrinsic astrophysical population – and this knowledge can be folded into the model comparison.

We seek a generalization to Eq. (11) for the posterior odds $\mathcal{O}_{A/B}$ that takes into account selection effects. As in Eqs. (8) and (10), the central quantity we require is the evidence for an astrophysical model A , but now conditioned on detectability (\det). Using Bayes' theorem, we may write this as

$$\mathcal{Z}(\mathbf{d}|A, \det) = \frac{P(\det|\mathbf{d}) \mathcal{Z}(\mathbf{d}|A)}{P(\det|A)}; \quad (12)$$

cf. Eq. (10) of Loredò (2004). The probability of successfully detecting a signal in the data is $P(\det|\mathbf{d}) = 1$ for the detected GW event under consideration, where we are making the assumption of a deterministic detection threshold as a function of the data. Given only the set of parameters $\boldsymbol{\theta}$, the source detection probability is found by marginalizing over the data distribution, i.e.,

$$P(\det|\boldsymbol{\theta}) = \int P(\det|\mathbf{d}') \mathcal{L}(\mathbf{d}'|\boldsymbol{\theta}) d\mathbf{d}'. \quad (13)$$

The prior-averaged detection probability

$$P(\text{det}|A) = \int P(\text{det}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}|A)d\boldsymbol{\theta} \quad (14)$$

is the fraction of sources from the population A that is detectable.

In Eq. (12), the numerator contains information about the *detected* source, being that its terms depend on the data, while the denominator accounts for the distribution of *detectable* sources; all detections are detectable by definition, but not all detectable sources end up being detected. Indeed, this detectability-conditioned evidence accounts for the fact that extending the model into unobservable regions of the parameter space should not alter its consistency with an observed event.

Putting the above pieces together yields

$$\frac{\mathcal{Z}(\mathbf{d}|A, \text{det})}{\mathcal{Z}(\mathbf{d}|U)} = \frac{\mathcal{B}_{A/U}}{P(\text{det}|A)}, \quad (15)$$

which again matches the population-level likelihood but now including selection effects (Thrane & Talbot 2019; Mandel et al. 2019; Vitale et al. 2022). Note here that as we have used $P(\text{det}|\mathbf{d}) = 1$, \mathbf{d} represents a particular realization of the data rather than a variable. As such, the evidences in Eq. (15) are evaluations of the marginal data likelihoods and are not normalized distributions.

We again refer the reader to Appendix A which describes how to compute the quantities in Eq. (15) in practice.

Similar to Sec. 2.3, we can compare two astrophysically-informed models A and B , neither of which we have posteriors for, with

$$\mathcal{D}_{A/B} = \frac{\mathcal{Z}(\mathbf{d}|A, \text{det})}{\mathcal{Z}(\mathbf{d}|B, \text{det})} = \frac{P(\text{det}|B)}{P(\text{det}|A)} \mathcal{B}_{A/B}. \quad (16)$$

Under equal prior odds, this represents our a posteriori preference for model A over B in producing the data with a detectable GW signal. Under unequal prior odds, $\mathcal{D}_{A/B}$ should be multiplied by the prior ratio, as in Eq. (11).

3 EXAMPLES IN GRAVITATIONAL-WAVE ASTRONOMY

We now apply our model-comparison formalism to a few example populations and GW events. We use public GW data from O3 (Abbott et al. 2023), in particular the reanalysis of Abbott et al. (2021a). We use posterior samples generated with the IMRPHEMOMXPHM waveform model (Pratten et al. 2021) that have been reweighted to a redshift prior corresponding to a uniform source-frame volumetric merger rate.

3.1 Three-body interactions in young star clusters

Dall’Amico et al. (2021) explore the possibility of forming stellar-mass binary BH mergers through three-body dynamics and assess whether this formation channel can produce mergers whose source properties are consistent with those of GW190521. They performed 2×10^5 direct N -body simulations of binary–single interactions including post-Newtonian terms up to the 2.5 order (Memmesheimer et al. 2004; Mikkola & Merritt 2008) and a prescription for relativistic kicks (Healy & Lousto 2018). The BH population is generated from the

simulations of young star clusters (YSCs) by Di Carlo et al. (2019), including BHs in the pair-instability mass gap between 60–120 M_{\odot} (Woosley 2017) formed via stellar collisions.

The mergers resulting from these dynamical interactions can occur through two main processes: flyby events or exchanges. In flybys, a tertiary perturber trades energy and momentum with the original binary, leaving the system configuration unchanged. In exchanges, one component of the original binary is expelled and replaced by the perturber during the interaction. We label the three possible scenarios as primary–secondary (m_1 – m_2), primary–tertiary (m_1 – m_3), and secondary–tertiary (m_2 – m_3) mergers; i.e., the label indicates the BH components of the final merger.

A small fraction ($\sim 0.1\%$) of the YSC mergers contain a BH remnant from a previous merger. We exclude these sources since their low number would give unreliable statistical estimates. This results in a total population of 7177 binaries. We neglect the inclusion of BH spins of the YSC binaries. Dall’Amico et al. (2021) investigated the effect of spins in post-processing by assuming various Maxwellian distributions for the spin magnitudes, while keeping their directions isotropic. We reflect the uncertainty – due to our (lack of) understanding of, e.g., core–envelope coupling, tides, and accretion (Qin et al. 2018; Fuller & Ma 2019; Steinle & Kesden 2021; Olejak & Belczynski 2021; Shao & Li 2022; Zevin & Bavera 2022; Ma & Fuller 2023) – in the astrophysical origin of spin (Callister et al. 2022; Mould et al. 2022b; Tong et al. 2022; Périgois et al. 2023) or lack thereof (Roulet et al. 2021; Galaudage et al. 2021) by taking the astrophysical spin prior to match the original PE prior (uniform up to maximum magnitudes and isotropic in direction); see Appendix B.

In Fig. 2, we show the distributions of source-frame component masses of mergers from the YSC simulations, classified according to the type of dynamical interaction. Overplotted is the GW190521 posterior. Simply by visual inspection it is clear that mergers in which the lighter component of the original binary is exchanged in the dynamical encounter (m_1 – m_3) reach higher masses. This YSC subpopulation is therefore most likely to accommodate the high masses of an event like GW190521 – $m_1 = 98_{-22}^{+34} M_{\odot}$ and $m_2 = 57_{-30}^{+27} M_{\odot}$ (medians and symmetric 90% CIs) – compared to m_1 – m_2 and m_2 – m_3 binaries, as pointed out in by Dall’Amico et al. (2021). Flyby encounters (m_1 – m_2) produce binaries with the lowest total masses and thus the lowest overlap with the mass posterior of GW190521 because the tertiary masses in this model are likely higher than the secondary masses (Dall’Amico et al. 2021).

We formalize these conclusions with the Bayes factors listed in Table 1. Remembering that these are relative measures, we quantify the likelihood that each dynamical encounter is preferred as a model for GW190521 over the m_1 – m_2 flyby channel using Eqs. (10, 16); see also Eqs. (A5, A6). Indeed, in support of the qualitative discussion above, we find that dynamical exchanges in which the original primary forms a merging binary with the interacting tertiary are favoured over tertiary flybys with $\ln \mathcal{B} = 2.8$ – a moderately strong but not decisive conclusion. Secondary–tertiary mergers are also preferred over primary–secondary mergers but only with mild evidence ($\ln \mathcal{B} = 1.2$). These results are in broad agreement with Dall’Amico et al. (2021) (see their Table 2), but are now consistent with the Bayesian PE framework.

We compute the detection-conditioned Bayes factors by




Figure 2. Distributions of the source-frame primary and secondary component masses of YSC mergers, classified by the type of dynamical interaction: tertiary flyby (m_1-m_2 , blue), original secondary exchanged (m_1-m_3 , orange), and original primary exchanged (m_2-m_3 , green). Overplotted is the posterior for GW190521 (purple). The panels on the diagonal present the 1D marginal distributions; for the YSC simulations the contours contain 99% of sources, while for GW190521 the shading represents darker to lighter the 50%, 90%, and 99% credible regions.

Merger	$\ln \mathcal{B}$	$\ln \mathcal{D}$
m_1-m_2	0	0
m_1-m_3	2.8	2.1
m_2-m_3	1.2	1.1

Table 1. Astrophysical (\mathcal{B}) and detection-conditioned (\mathcal{D}) Bayes factors comparing subpopulations of mergers within YSCs for GW190521. Dynamical flybys are labelled m_1-m_2 and exchanges in which the original binary BH secondary (primary) is replaced with the tertiary BH is labelled m_1-m_3 (m_2-m_3). Values are quoted relative to the m_1-m_2 sub-channel.

averaging over 10^5 draws from the priors (equal to the original PE priors) assumed over the unmodelled parameters (spins, redshifts, and other extrinsic parameters) for each simulated merger, as in Appendix A2. The inclusion of selection effects via Eq. (16) reduces the strength of model preference but does not change the qualitative results (see Table 1); the distribution of detectable m_1-m_3 (m_1-m_2) mergers is most (least) favoured, but only with mild evidence since $|\ln \mathcal{D}| \lesssim 2$ in all cases.

3.2 Hierarchical mergers

Hierarchical BH mergers (Gerosa & Fishbach 2021) – in which at least one binary component is the remnant of a previous

merger (or multiple previous mergers) – may contribute to the binary BH merger rate observed by LIGO/Virgo (Kimball et al. 2021; Mould et al. 2022a; Wang et al. 2022; Li et al. 2023). This scenario crucially depends on the properties of the environments in which so-called higher-generation binaries may form (Mapelli et al. 2021) and in particular on their escape speeds, as post-merger remnants receive GW recoils (Fitchett 1983) that can eject them their hosts and prevent repeated mergers. If retained, hierarchical mergers result in BHs with distinctively large masses and spins (Gerosa & Fishbach 2021). To a lesser degree, asymmetric masses form through inter-generational pairings, though even cluster dynamics tend to preferentially lead to nearly equal-mass systems (O’Leary et al. 2016; Di Carlo et al. 2019, 2020; Antonini et al. 2023) due to, e.g., mass segregation. The hierarchical merger scenario has therefore been invoked as a possible origin for the formations of GW190412 (Abbott et al. 2020b; Gerosa et al. 2020b; Rodriguez et al. 2020; Tagawa et al. 2021; Liu & Lai 2021) and GW190521 (Abbott et al. 2020e; Fragione et al. 2020; Dall’Amico et al. 2021; Kimball et al. 2021; Arca-Sedda et al. 2021; Tagawa et al. 2021; Liu & Lai 2021).

Labelling BHs that formed as the endpoint of stellar evolution as “first-generation” (1G) BHs and the merger product of two 1G BHs as a “second-generation” (2G) BH, we test which merger generations are most consistent with the GW events GW190412 and GW190521 using a simple model of hierarchical BH formation (Gerosa & Mould 2023; Mould & Gerosa 2023). In this model we define a “cluster” as a collection of BHs that can each pair to form a merging binary hosted in an environment with a common escape speed v_{esc} . We generate 10^3 distinct clusters whose escape speeds are chosen uniform randomly between 0 km/s and 10^3 km/s, each with an initial distribution of 10^4 1G BHs whose masses are drawn from a powerlaw between $5M_\odot$ and $50M_\odot$ (reflecting cutoffs due to pair-instability supernovae) with index -2 (roughly tracing the stellar initial mass function) and whose spins have uniform magnitudes between 0 and 0.2 (assuming low natal spins). One by one, BHs are drawn and paired with a less massive binary partner, both with uniform probabilities. Their spin directions are sampled isotropically. The properties of the merger remnant – its mass, spin magnitude, and kick magnitude – are then estimated (Gerosa & Kesden 2016). If the kick exceeds the assigned escape speed of the host cluster the remnant is ejected, but otherwise it is retained. The steps above are then iterated to exhaustion of the available BHs and this is repeated for all the clusters. Finally, the mergers from all 10^3 clusters are concatenated to form one population containing $\sim 7 \times 10^6$ mergers.

This simple procedure is able to rapidly generate large populations of hierarchical mergers. However, we note that the various assumptions in our parametrization are not necessarily realistic in an astrophysical setting. However, it is useful in its computational efficiency and produces population distributions reflecting the key features of hierarchical BH formation, serving as an illustrative example for this work.

Figure 3 displays the output of these simple simulations. We split the full population into four distinct subsets: those in which the merging binary contained two 1G BHs (1G+1G), one 1G BH and one 2G BH (1G+2G), two 2G remnants (2G+2G), and at least one component which has already undergone more than one previous merger ($>2G$). We plot the population distributions of the binary chirp mass M_c , mass

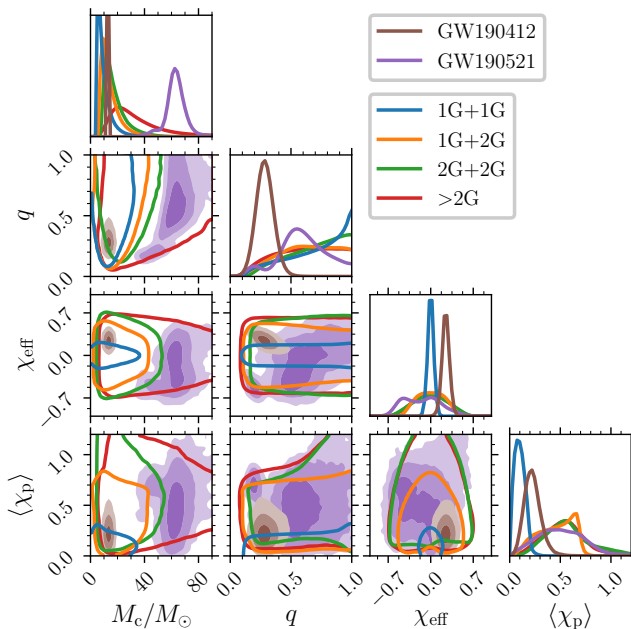


Figure 3. Distributions of chirp mass M_c , mass ratio q , effective aligned spin χ_{eff} , and averaged effective precessing spin $\langle\chi_p\rangle$. Diagonal and lower panels present one- and 2D marginals, respectively. The 2D shaded regions for GW190412 (brown) and GW190521 (purple) represent the 50%, 90%, and 99% credible regions, from darker to lighter. The simulated hierarchical mergers are split into subpopulations of binaries containing two 1G BHs (1G+1G, blue), a 1G BH and a 1G+1G remnant (1G+2G, orange), two 2G BHs (2G+2G, green), or at least one BH having undergone more than one previous merger (>2G, red). Their 2D distributions are given by contours containing 99% of sources.

ratio q , effective aligned spin χ_{eff} (Racine 2008), and averaged effective precessing spin $\langle\chi_p\rangle$ (Gerosa et al. 2021); the first three are conserved quantities while $\langle\chi_p\rangle$ varies over the inspiral (but not over a spin-precession cycle) and is bounded in $[0, 2]$. The effective spins in particular have been shown to present features identifying hierarchical mergers (Mapelli et al. 2021; Baibhav et al. 2021). Over plotted are the posteriors of these parameters for GW190412 and GW190521. The higher-generation distributions unsurprisingly reach higher masses and spins which can accommodate the heavy BHs and possibility of tilted spins in GW190521 and non-zero χ_{eff} of GW190412. The 1G+2G and >2G subpopulations overlap most with the marginal mass-ratio distributions of GW190412 as 1G+1G and 2G+2G formations are more likely to pair equal-mass BHs.

We again formalize these notions with the Bayes factors from Eqs. (10, 16). In this case, since we have a large number ($\sim 10^7$) of total mergers, when computing the detection fractions we assign a single value drawn from the PE prior for the set of unmodelled parameters (redshift, inclination, right ascension, declination, and polarization) to each merger, rather than averaging over a set of Monte Carlo realizations. We compute astrophysical and detection-conditioned Bayes factors with respect to the 1G+1G subpopulation and report the results in Table 2.

For GW190412, 1G+2G formation is favoured with re-

Generation	GW190412		GW190521	
	$\ln \mathcal{B}$	$\ln \mathcal{D}$	$\ln \mathcal{B}$	$\ln \mathcal{D}$
1G+1G	0	0	0	0
1G+2G	3.2	2.3	94	93
2G+2G	1.2	-0.1	98	97
>2G	1.7	-0.7	102	100

Table 2. Intrinsic (\mathcal{B}) and detectability-conditioned (\mathcal{D}) Bayes factors for higher-generation over 1G origins for GW190412 and GW190521.

spect to other hierarchical mergers, with moderate evidence ($\ln \mathcal{B} = 3.2$) over first-generation binaries and milder evidence with respect to higher generations. Owing to the confidently unequal masses ($q = 0.28_{-0.07}^{+0.09}$) of GW190412, any population that pairs unequal-mass BHs while allowing for its positively aligned spin ($\chi_{\text{eff}} = 0.25_{-0.10}^{+0.10}$) and precise chirp mass ($M_c = 13.3_{-0.4}^{+0.5} M_\odot$) will be preferred. In particular, the mass-ratio distributions of 1G+2G and >2G binaries are flatter compared to equal-generation pairings, while 1G+2G binaries are more densely located in the region of GW190412’s chirp mass (see Fig. 3). The detectable populations of 2G+2G and >2G binaries become disfavoured with respect to 1G+1G binaries in describing GW190412 with $|\ln \mathcal{D}| \lesssim 1$, while the 1G+2G subpopulation remains favoured, though more weakly.

In the case of GW190521, we note that 1G+1G binaries often do not overlap with the GW posterior due to the maximum 1G mass $m_{\text{max}} = 50 M_\odot$. This means that, in computing the Bayes factors in Eqs. (A5, A6), the result becomes sensitive to the number of posterior samples used due to low effective counts; increasing the posterior sample size likely increases the denominator of Eq. (A1) without increasing the numerator. We use all the samples available in the public data ($\sim 2 \times 10^5$ for GW190412 and $\sim 2 \times 10^4$ for GW190521), which implies that the Bayes factors for GW190521 in Table 2 represent lower limits. The result, however, is clear: in our simplified model of hierarchical mergers with our selected parametrization, GW190521 is confidently favoured to contain a higher-generation BH over our assumed 1G population, with $\ln \mathcal{B} \approx 100$. In fact, the subpopulation of binaries in which at least one BH has undergone more than one merger (>2G) is most favoured, confidently ($\ln \mathcal{B} > 5$) over 1G+2G binaries but less so ($\ln \mathcal{B} < 5$) over 2G+2G binaries. This conclusion, or its strength, would change if, e.g., we increased the maximum allowed mass of 1G BHs in our parametrization. Including selection effects, we find these results are almost unchanged for the detectable hierarchical-merger distributions (see Table 2).

Furthermore, we compute the Bayes factors for the populations of both intrinsic and detectable mergers *within* each cluster – i.e., as a function of the host escape speed v_{esc} – over the binary merger generations. The results are shown in Fig. 4. We first note that the noise is due to the fact that each cluster produces a number of mergers that is at most one less than the initial number (10^4) of 1G BHs and reflects the statistical uncertainty in the Monte Carlo approximation. At low escape speeds there are not many – if any – higher-generation mergers, implying the approximated Bayes factors (see Appendix A) cannot be trusted there.




Figure 4. Intrinsic (\mathcal{B} , top row) and detectability-conditioned (\mathcal{D} , bottom row) Bayes factors for higher-generation origins of GW190412 (left column) and GW190521 (right column) as a function of the host cluster escape speed v_{esc} . The Bayes factors for binaries with one 1G and one 2G BH (1G+2G, orange), two 2G BHs (2G+2G, green), and at least one BH of even higher generation (>2G, red) are given with respect to 1G+1G binaries (blue). Shaded lines represent the Monte Carlo estimation while solid lines give a moving average. The horizontal dotted black lines mark the boundaries of equal preference.

For GW190412, >2G binaries become disfavoured relative to other higher-generation populations at larger escape speeds because such clusters more efficiently retain larger masses that are inconsistent with those of the observed event; for this same reason, in the case of GW190521 the astrophysical and detectable >2G Bayes factors increase on average with larger cluster escape speeds. The astrophysical (detectable) >2G subpopulation remains favoured (becomes disfavoured) over the 1G+1G subpopulation for increasing v_{esc} , while the Bayes factors for 1G+2G and 2G+2G across all clusters are consistent with the overall population in Table 2. On the other hand, for GW190521 the per-cluster preference for higher-generation binaries is reduced compared to the overall population, but there is still a decisive conclusion against first-generation binaries. We reiterate that this conclusion depends on our choice of parametrization.

4 CONCLUSIONS

Even as the number of detected GW events grows into the triple digits, there are individual detections which are interesting on their own merit. They may have unusual source properties that make them outliers with respect to the rest of the inferred population. For these detections we may wish to individually assess the degree to which they are reproducible by a given formation scenario. Some previous analyses have employed the heuristic fraction of synthetic mergers from a simulated population that lie within a chosen volume of the

measured parameters of a real GW event. This metric is unsuitable as it does not account for the full Bayesian posterior and prior information in the original PE, it depends on a heuristically constructed confidence region, and it fails in the limit of precise measurements (Sec. 2.1).

We point out that a statistically consistent replacement is given by the posterior odds (Secs. 2.2 and 2.3), with a tractable computation in practice (Appendix A). We also show how the influence of selection biases can be accommodated in this statistical framework (Sec. 2.4), consistent with the usual hierarchical GW population likelihood. In fact, rather than starting from a single GW event and working up, one can similarly derive Eqs. (15, 16) starting from the full population likelihood and working down by considering the limiting case of a single observation and fixed population models.

Crucially, the posterior odds is a relative or comparative measure; one can only test whether a particular astrophysical model is more likely to produce mergers consistent with a given GW event compared to another, but not whether the model is good in absolute terms – a natural outcome of working within the Bayesian framework. We used this fact to assess the degree to which selected GW observations are better explained by various simulated formation pathways relative to one another (Sec. 3). In particular, we demonstrated comparisons between subpopulations (e.g., first-generation against second-generation mergers) within a particular formation environment (e.g., dynamical clusters).

As mentioned in Sec. 1.1, the calculation we presented is actually a two-step procedure, the second of which we derived in this work: first, one selects an event of interest, then re-analyses it in the context of competing population models to deduce which better describes the data. We neglected a quantitative inclusion of the first step as we were interested in the information provided by a single event irrespective of other events in the combined observational catalogue. One can account for the information provided by the remaining catalogue, assuming that the single-event selection is equally probable among all events, with population-informed priors (Moore & Gerosa 2021). Otherwise, the selection of that event must be accounted for to avoid biases, akin to those due to search sensitivity which results in the additional detection-probability factors in Eq. (12). To take this approach, one can follow Essick et al. (2022) who derive the population-level likelihood for a catalogue of detections but marginalise over a single event that may be an outlier while self-consistently accounting for the very selection of that event. The problem tackled here is the converse: one analyses a single interesting event while marginalising over the remaining catalogue, accounting for the same selection effect just mentioned; cf. their Eqs. (7) and (10). Defining an event as interesting may depend on the remaining catalogue – e.g., it is clearly separated in parameter space from all others – or on astrophysical prior knowledge – e.g., it lies in a predefined data-independent mass gap.

This framework can be used not just when selecting a single event from a larger catalogue of observations, but also when there is only a single observation available in the first place. This may be the case for LISA (Amaro-Seoane et al. 2017), where the most pessimistic forecasts predict just $\sim \mathcal{O}(1)$ massive BH merger and extreme mass-ratio inspiral observations annually (Amaro-Seoane et al. 2023). The degree to which we can conclude model preference in the presence of

observational biases in this scenario of limited available data will therefore be an interesting topic for future work.

ACKNOWLEDGEMENTS

We thank Christopher Moore, Riccardo Buscicchio, Floor Broekgaarden, and Christopher Berry for helpful discussions. M.Mo. and D.G. are supported by European Union's H2020 ERC Starting Grant No. 945155–GWmining, Cariplo Foundation Grant No. 2021-0555, and the ICSC National Research Centre funded by NextGenerationEU. D.G. is supported by Leverhulme Trust Grant No. RPG-2019-350. M.D. is supported by Cariparo Foundation Grant. No. 55440. M.Ma. is supported by European Union's H2020 ERC Consolidator Grant No. 770017–DEMOBLACK and by the German Excellence Strategy via the Heidelberg Cluster of Excellence (EXC 2181 - 390900948) STRUCTURES.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author. We make available a reproducible tool to perform the computation along with examples at github.com/mdmould/popodds; Mould 2023.

REFERENCES

- Aasi J., et al., 2015, *Classical Quant. Grav.*, 32, 074001 (arXiv:1411.4547)
- Abbott B. P., et al., 2019, *Phys. Rev. X*, 9, 031040 (arXiv:1811.12907)
- Abbott R., et al., 2020a, *Living Rev. Relativ.*, 23, 3
- Abbott R., et al., 2020b, *Phys. Rev. D*, 102, 043015 (arXiv:2004.08342)
- Abbott R., et al., 2020c, *Phys. Rev. Lett.*, 125, 101102 (arXiv:2009.01075)
- Abbott R., et al., 2020d, *Astrophys. J. Lett.*, 896, L44 (arXiv:2006.12611)
- Abbott R., et al., 2020e, *Astrophys. J. Lett.*, 900, L13 (arXiv:2009.01190)
- Abbott R., et al., 2021a, (arXiv:2108.01045)
- Abbott R., et al., 2021b, (arXiv:2111.03606)
- Abbott R., et al., 2021c, *Phys. Rev. X*, 11, 021053 (arXiv:2010.14527)
- Abbott R., et al., 2023, *Astrophys. J. Supp.*, 267, 29 (arXiv:2302.03676)
- Acernese F., et al., 2015, *Classical Quant. Grav.*, 32, 024001 (arXiv:1408.3978)
- Amaro-Seoane P., et al., 2017, (arXiv:1702.00786)
- Amaro-Seoane P., et al., 2023, *Living Rev. Relativ.*, 26, 2 (arXiv:2203.06016)
- Antonini F., Gieles M., Dosopoulou F., Chattopadhyay D., 2023, *Mon. Not. R. Astron. Soc.*, 522, 466 (arXiv:2208.01081)
- Arca-Sedda M., Rizzuto F. P., Naab T., Ostriker J., Giersz M., Spurzem R., 2021, *Astrophys. J.*, 920, 128 (arXiv:2105.07003)
- Baibhav V., Berti E., Gerosa D., Mould M., Wong K. W. K., 2021, *Phys. Rev. D*, 104, 084002 (arXiv:2105.12140)
- Biscoveanu S., Isi M., Varma V., Vitale S., 2021, *Phys. Rev. D*, 104, 103018 (arXiv:2106.06492)
- Broekgaarden F. S., Stevenson S., Thrane E., 2022, *Astrophys. J.*, 938, 45 (arXiv:2205.01693)
- Bustillo J. C., et al., 2021, *Phys. Rev. Lett.*, 126, 081101 (arXiv:2009.05376)
- Callister T. A., Miller S. J., Chatziioannou K., Farr W. M., 2022, *Astrophys. J. Lett.*, 937, L13 (arXiv:2205.08574)
- Chapman-Bird C. E. A., Berry C. P. L., Woan G., 2023, *Mon. Not. R. Astron. Soc.*, 522, 6043 (arXiv:2212.06166)
- Dall'Amico M., Mapelli M., Di Carlo U. N., Bouffanais Y., Rastello S., Santoliquido F., Ballone A., Arca Sedda M., 2021, *Mon. Not. R. Astron. Soc.*, 508, 3045 (arXiv:2105.12757)
- Dax M., Green S. R., Gair J., Macke J. H., Buonanno A., Schölkopf B., 2021, *Phys. Rev. Lett.*, 127, 241103 (arXiv:2106.12594)
- Di Carlo U. N., Giacobbo N., Mapelli M., Pasquato M., Spera M., Wang L., Haardt F., 2019, *Mon. Not. R. Astron. Soc.*, 487, 2947 (arXiv:1901.00863)
- Di Carlo U. N., et al., 2020, *Mon. Not. R. Astron. Soc.*, 498, 495 (arXiv:2004.09525)
- Essick R., Farah A., Galadage S., Talbot C., Fishbach M., Thrane E., Holz D. E., 2022, *Astrophys. J.*, 926, 34 (arXiv:2109.00418)
- Estellés H., et al., 2022, *Astrophys. J.*, 924, 79 (arXiv:2105.06360)
- Fishbach M., Holz D. E., 2020, *Astrophys. J. Lett.*, 904, L26 (arXiv:2009.05472)
- Fishbach M., Farr W. M., Holz D. E., 2020, *Astrophys. J. Lett.*, 891, L31 (arXiv:1911.05882)
- Fitchett M. J., 1983, *Mon. Not. R. Astron. Soc.*, 203, 1049
- Fragione G., Loeb A., Rasio F. A., 2020, *Astrophys. J. Lett.*, 902, L26 (arXiv:2009.05065)
- Fuller J., Ma L., 2019, *Astrophys. J. Lett.*, 881, L1 (arXiv:1907.03714)
- Galadage S., Talbot C., Nagar T., Jain D., Thrane E., Mandel I., 2021, *Astrophys. J. Lett.*, 921, L15 (arXiv:2109.02424)
- Gamba R., Breschi M., Carullo G., Albanesi S., Rettugno P., Bernuzzi S., Nagar A., 2023, *Nat. Astron.*, 7, 11 (arXiv:2106.05575)
- Gayathri V., et al., 2022, *Nat. Astron.*, 6, 344
- Gerosa D., Fishbach M., 2021, *Nat. Astron.*, 5, 749 (arXiv:2105.03439)
- Gerosa D., Kesden M., 2016, *Phys. Rev. D*, 93, 124066 (arXiv:1605.01067)
- Gerosa D., Mould M., 2023, (arXiv:2305.04987)
- Gerosa D., Pratten G., Vecchio A., 2020a, *Phys. Rev. D*, 102, 103020 (arXiv:2007.06585)
- Gerosa D., Vitale S., Berti E., 2020b, *Phys. Rev. Lett.*, 125, 101103 (arXiv:2005.04243)
- Gerosa D., Mould M., Gangardt D., Schmidt P., Pratten G., Thomas L. M., 2021, *Phys. Rev. D*, 103, 064067 (arXiv:2011.11948)
- Green S. R., Gair J., 2021, *Machine Learning: Science and Technology*, 2, 03LT01
- Hannam M., et al., 2022, *Nature*, 610, 652 (arXiv:2112.11300)
- Healy J., Lousto C. O., 2018, *Phys. Rev. D*, 97, 084002 (arXiv:1801.08162)
- Jeffreys H., 1939, *Theory of Probability*. Oxford University Press
- Kagra Collaboration et al., 2019, *Nature Astronomy*, 3, 35 (arXiv:1811.08079)
- Kimball C., et al., 2021, *Astrophys. J. Lett.*, 915, L35 (arXiv:2011.05332)
- Kobyzev I., Prince S. J. D., Brubaker M. A., 2019, (arXiv:1908.09257)
- Li Y.-J., Wang Y.-Z., Tang S.-P., Fan Y.-Z., 2023, (arXiv:2303.02973)
- Liu B., Lai D., 2021, *Mon. Not. R. Astron. Soc.*, 502, 2049 (arXiv:2009.10068)
- Loredo T. J., 2004, *AIP Conf. Proc.*, 735, 195 (arXiv:astro-ph/0409387)
- Ma L., Fuller J., 2023, *Astrophys. J.*, 952, 53 (arXiv:2305.08356)
- Mandel I., Farr W. M., Gair J. R., 2019, *Mon. Not. R. Astron. Soc.*, 486, 1086 (arXiv:1809.02063)
- Mapelli M., et al., 2021, *Mon. Not. R. Astron. Soc.*, 505, 339 (arXiv:2103.05016)
- Memmesheimer R.-M., Gopakumar A., Schäfer G., 2004,

Phys. Rev. D, 70, 104011 (arXiv:gr-qc/0407049)

Mikkola S., Merritt D., 2008, *Astron. J.*, 135, 2398 (arXiv:0709.3367)

Moore C. J., Gerosa D., 2021, *Phys. Rev. D*, 104, 083008 (arXiv:2108.02462)

Mould M., 2023, *mdmould/popodds*, doi:10.5281/zenodo.7647704

Mould M., Gerosa D., 2023, *mdmould/qluster*, doi:10.5281/zenodo.7807210

Mould M., Gerosa D., Taylor S. R., 2022a, *Phys. Rev. D*, 106, 103013 (arXiv:2203.03651)

Mould M., Gerosa D., Broekgaarden F. S., Steinle N., 2022b, *Mon. Not. R. Astron. Soc.*, 517, 2738 (arXiv:2205.12329)

Nitz A. H., Capano C. D., 2021, *Astrophys. J. Lett.*, 907, L9 (arXiv:2010.12558)

O’Leary R. M., Meiron Y., Kocsis B., 2016, *Astrophys. J. Lett.*, 824, L12 (arXiv:1602.02809)

Olejak A., Belczynski K., 2021, *Astrophys. J. Lett.*, 921, L2 (arXiv:2109.06872)

Olsen S., Roulet J., Chia H. S., Dai L., Venumadhav T., Zackay B., Zaldarriaga M., 2021, *Phys. Rev. D*, 104, 083036 (arXiv:2106.13821)

Papamakarios G., Nalisnick E., Jimenez Rezende D., Mohamed S., Lakshminarayanan B., 2019, (arXiv:1912.02762)

Payne E., Thrane E., 2023, *Phys. Rev. Res.*, 5, 023013 (arXiv:2210.11641)

Payne E., Hourihane S., Golomb J., Udall R., Davis D., Chatzioanou K., 2022, *Phys. Rev. D*, 106, 104017 (arXiv:2206.11932)

Pérgois C., Mapelli M., Santoliquido F., Bouffanais Y., Rufolo R., 2023, (arXiv:2301.01312)

Pratten G., et al., 2021, *Phys. Rev. D*, 103, 104056 (arXiv:2004.06503)

Qin Y., Fragos T., Meynet G., Andrews J., Sørensen M., Song H. F., 2018, *Astron. Astrophys.*, 616, A28 (arXiv:1802.05738)

Racine É., 2008, *Phys. Rev. D*, 78, 044021 (arXiv:0803.1820)

Rodriguez C. L., et al., 2020, *Astrophys. J. Lett.*, 896, L10 (arXiv:2005.04239)

Romero-Shaw I., Lasky P. D., Thrane E., Calderón Bustillo J., 2020, *Astrophys. J. Lett.*, 903, L5 (arXiv:2009.04771)

Romero-Shaw I. M., Thrane E., Lasky P. D., 2022, *Publ. Astron. Soc. Aust.*, 39, e025 (arXiv:2202.05479)

Romero-Shaw I. M., Gerosa D., Loutrel N., 2023, *Mon. Not. R. Astron. Soc.*, 519, 5352 (arXiv:2211.07528)

Roulet J., Chia H. S., Olsen S., Dai L., Venumadhav T., Zackay B., Zaldarriaga M., 2021, *Phys. Rev. D*, 104, 083010 (arXiv:2105.10580)

Scott D. W., 2015, *Multivariate Density Estimation: Theory, Practice, and Visualization*. John Wiley and Sons

Shao Y., Li X.-D., 2022, *Astrophys. J.*, 930, 26 (arXiv:2203.14529)

Steinle N., Kesden M., 2021, *Phys. Rev. D*, 103, 063032 (arXiv:2010.00078)

Tagawa H., Kocsis B., Haiman Z., Bartos I., Omukai K., Samsing J., 2021, *Astrophys. J.*, 908, 194 (arXiv:2012.00011)

Talbot C., Golomb J., 2023, (arXiv:2304.06138)

Talbot C., Thrane E., 2022, *Astrophys. J.*, 927, 76

Thrane E., Talbot C., 2019, *Publ. Astron. Soc. Aust.*, 36, e010 (arXiv:1809.02293)

Tong H., Galadage S., Thrane E., 2022, *Phys. Rev. D*, 106, 103019

Varma V., et al., 2022, *Phys. Rev. Lett.*, 128, 191102 (arXiv:2201.01302)

Virtanen P., et al., 2020, *Nat. Methods*, 17, 261

Vitale S., Gerosa D., Farr W. M., Taylor S. R., 2022, in *Handbook of Gravitational Wave Astronomy*. p. 45, doi:10.1007/978-981-15-4702-7_45-1

Wang Y.-Z., Li Y.-J., Vink J. S., Fan Y.-Z., Tang S.-P., Qin Y., Wei D.-M., 2022, *Astrophys. J. Lett.*, 941, L39 (arXiv:2208.11871)

Wong K. W. K., Ng K. K. Y., Berti E., 2020, (arXiv:2007.10350)

Woolsey S. E., 2017, *Astrophys. J.*, 836, 244 (arXiv:1608.08939)

Zevin M., Bavera S. S., 2022, *Astrophys. J.*, 933, 86

(arXiv:2203.02515)

Zhang R. C., Fragione G., Kimball C., Kalogera V., 2023, *Astrophys. J.*, 954, 23 (arXiv:2302.07284)

APPENDIX A: PRACTICAL COMPUTATION

A1 Bayes factor

The integral in the Bayes factor of Eq. (8) can be computed as a Monte Carlo summation in two ways that are equivalent (modulo numerical issues – see below), unlike for the heuristic fraction:

$$\mathcal{B}_{A/U} \approx \frac{1}{N} \sum_{i=1}^N \frac{\pi(\boldsymbol{\theta}_i|A)}{\pi(\boldsymbol{\theta}_i|U)} \approx \frac{1}{N_A} \sum_{j=1}^{N_A} \frac{p(\boldsymbol{\theta}_{A_j}|\mathbf{d}, U)}{\pi(\boldsymbol{\theta}_{A_j}|U)}. \quad (\text{A1})$$

The first expression can be computed if one has posterior samples $\{\boldsymbol{\theta}_i\}_{i=1}^N \sim p(\boldsymbol{\theta}|\mathbf{d}, U)$ for a given GW event and can evaluate the population density $\pi(\boldsymbol{\theta}|A)$. The second expression can be computed if one can evaluate the posterior density $p(\boldsymbol{\theta}|\mathbf{d}, U)$ of the GW event and has synthetic sources $\{\boldsymbol{\theta}_{A_j}\}_{j=1}^{N_A} \sim \pi(\boldsymbol{\theta}|A)$ simulated from an astrophysical population. In either case we must be able to evaluate the prior density $\pi(\boldsymbol{\theta}|U)$ from the original PE.

Besides this, the better expression to compute is the one in which the additional uncertainty introduced by the integral approximation is lower. This depends on the variances of the GW posterior and simulated population, and on the number of samples over which the summation is taken (Talbot & Golomb 2023). If one already has an explicit form for $\pi(\boldsymbol{\theta}|A)$ given by, e.g., a parametrized model, the first expression of Eq. (A1) should be used. Otherwise, one should use the first expression of Eq. (A1) if the error on the GW event posterior $p(\boldsymbol{\theta}|\mathbf{d}, U)$ is smaller than the range of the astrophysical population $\pi(\boldsymbol{\theta}|A)$, and vice versa. The former is most often true in the case of GW data and as such is our default choice (e.g., a single binary mass measurement is more narrow than the entire range of astrophysical masses, though this is not necessarily true for parameters that are typically not well measured such as BH spins).

It may be the case that we do not have access to a closed-form expression for the uninformative PE prior $\pi(\boldsymbol{\theta}|U)$ but do have samples drawn from it – in this case a density estimator is required. If building an informative prior from astrophysical simulations, we also need to construct an estimator for $\pi(\boldsymbol{\theta}|A)$ from the finite set of simulated sources. For the examples following in Sec. 3 we use Gaussian kernel density estimates (KDEs) (Scott 2015; Virtanen et al. 2020); for more complex or higher-dimensional distributions other density estimation techniques may be more accurate, e.g., normalizing flows (Kobyzev et al. 2019; Papamakarios et al. 2019). On the other hand, if a continuous representation of the posterior density $p(\boldsymbol{\theta}|\mathbf{d}, U)$ is already available (e.g., through neural posterior estimation as in Green & Gair 2021; Dax et al. 2021), the discrete simulated sources can be used directly.

It should be noted, however, that the expression on the right of Eq. (A1) has a numerical divergence if $\pi(\boldsymbol{\theta}|A)$ has support beyond the support of $\pi(\boldsymbol{\theta}|U)$ – i.e., there may be simulated sources $\boldsymbol{\theta}_{A_j}$ for which $\pi(\boldsymbol{\theta}_{A_j}|U) = 0$, whereas the equivalent integrand $\mathcal{L}(\mathbf{d}|\boldsymbol{\theta}_{A_j})/\mathcal{Z}(\mathbf{d}|U)$ is perfectly well defined. This does not occur in the left-hand expression, for which we are guaranteed that $\pi(\boldsymbol{\theta}_i|U) > 0$ by definition.

A2 Detectable fractions

We also need to compute the detectable fractions in Eq. (16). In principle these quantities must be estimated by injecting signals with source properties θ into the full detection pipeline, with $p(\text{det}|\theta) = 1$ (0) for signals which are (not) recovered. However, this presents prohibitive computational expense when required for large populations. Additionally, the detection probability depends on all the source properties – both intrinsic and extrinsic – but we might only be interested in modelling the population of the former, or even just a subset of them. Each source must be assigned single values for the unmodelled parameters – which increases the uncertainty in the population-averaged detectability – or these parameters must be manually marginalized over – a procedure which must be done numerically and increases the computational burden.

In practice, a source is often approximated as detectable if its SNR is above a chosen threshold value. And recent approaches have utilized machine learning to enable efficient evaluations of source detectability (Gerosa et al. 2020a; Wong et al. 2020; Talbot & Thrane 2022; Chapman-Bird et al. 2023). The process of generating a waveform, projecting the signal into the detector frame, computing the detection statistic (e.g., SNR), and evaluating the detectability can be reduced to the evaluation of a neural network trained on sources already assessed in the full detection pipeline. In particular, such approximants can process large batches of inputs very efficiently. They are therefore well suited for averaging detection probabilities over large populations of simulated sources and, furthermore, for handling unmodelled parameters that require an additional averaging step, as mentioned above. We therefore choose to use the neural network implementation of Gerosa et al. (2020a), assuming a three-detector LIGO/Virgo network with sensitivities corresponding to the third observing run (O3) and a network SNR threshold of 12 (Abbott et al. 2020a). For each binary, the network takes as input (θ) the detector-frame component masses, spin vectors, redshift, orbital inclination, right ascension, declination, and GW polarization, and outputs the classification $p(\text{det}|\theta) = 1$ (0) for sources which are (not) detectable.

When the detection probabilities of each simulated binary are estimated directly, whether in this way or with a different pipeline, the detection fractions required in Eq. (16) can then be computed as expectation values over the population. For our set of simulated mergers $\{\theta_{Aj}\}_{j=1}^{N_A} \sim \pi(\theta|A)$, Eq. (14) can be approximated with

$$P(\text{det}|A) \approx \frac{1}{N_A} \sum_{j=1}^{N_A} P(\text{det}|\theta_{Aj}). \quad (\text{A2})$$

A3 Numerical implementation

Putting everything together to compare the GW event posterior $\{\theta_i\}_{i=1}^N \sim p(\theta|\mathbf{d}, U)$ to two astrophysical simulations $\{\theta_{Aj}\}_{j=1}^{N_A} \sim \pi(\theta|A)$ and $\{\theta_{Bk}\}_{k=1}^{N_B} \sim \pi(\theta|B)$, we get for the Bayes factor between the intrinsic populations

$$\mathcal{B}_{A/B} \approx \frac{\sum_{i=1}^N \pi(\theta_i|A)/\pi(\theta_i|U)}{\sum_{n=1}^N \pi(\theta_n|B)/\pi(\theta_n|U)}, \quad (\text{A3})$$

and between the detectable populations

$$\mathcal{D}_{A/B} \approx \frac{N_A \sum_{k=1}^{N_B} P(\text{det}|\theta_{Bk})}{N_B \sum_{j=1}^{N_A} P(\text{det}|\theta_{Aj})} \mathcal{B}_{A/B}. \quad (\text{A4})$$

It may be more numerically stable to evaluate the logarithms of the probability densities in Eq. (A3) and then compute the logarithms of the integrals using a log-sum-exp implementation (our default). The Bayes factors between the two populations are then given by

$$\begin{aligned} \ln \mathcal{B}_{A/B} &\approx \ln \sum_{i=1}^N \exp[\ln \pi(\theta_i|A) - \ln \pi(\theta_i|U)] \\ &\quad - \ln \sum_{i=1}^N \exp[\ln \pi(\theta_i|B) - \ln \pi(\theta_i|U)], \end{aligned} \quad (\text{A5})$$

$$\ln \mathcal{D}_{A/B} = \ln \mathcal{B}_{A/B} - \ln P(\text{det}|A) + \ln P(\text{det}|B). \quad (\text{A6})$$

APPENDIX B: MISSING AND TRANSFORMED PARAMETERS

In many situations, the new prior may model a subset or transformed versions of the original parameters, or both. An example from GW astronomy is that, while θ contains the detector-frame binary masses, component spins, luminosity distance, and sky location, etc., a model of compact-binary formation (such as a population-synthesis simulation) may only model source-frame masses and redshift, without modelling the other extrinsic parameters or perhaps the spins. These differences enter the calculation of the Bayes factor.

B1 Missing parameters

Suppose that the astrophysical prior only models a subset ϕ (e.g., intrinsic binary parameters) of the parameters $\theta = (\phi, \psi)$, where ψ are the unmodeled parameters (e.g., extrinsic binary parameters). One cannot simply replace θ with ϕ in the integral of Eq. (8). Since ϕ and ψ are independent of each other in our astrophysical prior and if we take equal priors $\pi(\psi|A) = \pi(\psi|U)$ (as is usual in population inference; Vitale et al. 2022), then

$$\mathcal{B}_{A/U} = \int \pi(\phi|A) \int \frac{p(\phi, \psi|\mathbf{d}, U)}{\pi(\phi|\psi, U)} d\psi d\phi. \quad (\text{B1})$$

Only if the original PE prior over ϕ is independent of ψ , i.e., $\pi(\phi|\psi, U) = \pi(\phi|U)$, may we compute the Bayes factor by replacing θ with ϕ ; this may be true in some cases but is not in general. For example, consider a model of source-frame binary masses when the original prior was placed over detector-frame masses; here, one must account for the redshift dependence.

B2 Transformed parameters

When dealing with a such transformation $\tilde{\theta} = \tilde{\theta}(\theta)$ (e.g., source-frame masses and redshift) of the original parameters θ (e.g., detector-frame masses and luminosity distance), one can either convert the original posterior $p(\theta|\mathbf{d}, U)$ and prior $\pi(\theta|U)$ to distributions over the new prior parameters $\tilde{\theta}$, or convert the new prior $\pi(\tilde{\theta}|A)$ in terms of the original parameters θ . The corresponding transformation of the distributions is $p(\tilde{\theta}) = p(\theta)|d\theta/d\tilde{\theta}|$, where the scaling is given by the Jacobian determinant of the inverse transformation. In fact, either parametrization is equivalent because

$$\int p(\theta|\mathbf{d}, U) \frac{\pi(\theta|A)}{\pi(\theta|U)} d\theta = \int p(\tilde{\theta}|\mathbf{d}, U) \frac{\pi(\tilde{\theta}|A)}{\pi(\tilde{\theta}|U)} d\tilde{\theta}. \quad (\text{B2})$$

Given the original posterior samples $\theta \sim p(\theta|\mathbf{d}, U)$, one can readily construct the transformed posterior samples as

$\tilde{\theta}(\boldsymbol{\theta}) \sim p(\tilde{\theta}|\mathbf{d}, U)$. The Bayes factor from Eq. (8) may be written as

$$\mathcal{B}_{A/U} = \int p(\boldsymbol{\theta}|\mathbf{d}, U) \frac{\pi(\tilde{\theta}(\boldsymbol{\theta})|A)}{\pi(\boldsymbol{\theta}|U)} \left| \frac{d\tilde{\theta}}{d\boldsymbol{\theta}} \right| d\boldsymbol{\theta} \quad (\text{B3})$$

$$= \int p(\tilde{\theta}|\mathbf{d}, U) \frac{\pi(\tilde{\theta}|A)}{\pi(\boldsymbol{\theta}(\tilde{\theta})|U)} \left| \frac{d\tilde{\theta}}{d\boldsymbol{\theta}} \right| d\tilde{\theta}, \quad (\text{B4})$$

where we have written the first expression in terms of the original parameters $\boldsymbol{\theta}$ but kept the new prior model A in terms of its parameters $\tilde{\theta}$, and we have written the second expression in terms of the transformed parameters $\tilde{\theta}$ but kept the original prior U in terms of its parameters $\boldsymbol{\theta}$.

As a concrete (and rather common – see, e.g., [Abbott et al. 2019, 2021c,a,b](#)) example, consider setting a prior over the redshifted component masses $m_{1z} = (1+z)m_1$ and $m_{2z} = (1+z)m_2$ of a GW event as observed in the detector frame and then modelling the source-frame masses $m_1 \geq m_2$ in the population without modelling the redshift z . As usual, we assume priors over the remaining binary parameters that are independent of masses and redshift, such that the corresponding terms in the Bayes factor cancel, as in Eq. (B1). The Jacobian for the transformation $(m_{1z}, m_{2z}, z) \mapsto (m_1, m_2, z)$ is $(1+z)^2$. In this example, we would therefore have that

$$\begin{aligned} \mathcal{B}_{A/U} &= \int dm_{1z} dm_{2z} dz (1+z)^2 p(m_{1z}, m_{2z}, z|\mathbf{d}, U) \\ &\quad \times \frac{\pi(m_{1z}/(1+z), m_{2z}/(1+z)|A)}{\pi(m_{1z}, m_{2z}|U)} \\ &= \int dm_1 dm_2 dz (1+z)^2 p(m_1, m_2, z|\mathbf{d}, U) \\ &\quad \times \frac{\pi(m_1, m_2|A)}{\pi(m_1(1+z), m_2(1+z)|U)}, \end{aligned} \quad (\text{B5})$$

where one should note the assumed cancellation $\pi(z|A)/\pi(z|U) = 1$. To reiterate, this is only valid because we have accounted for the transformation of variables.

APPENDIX C: SIMPLE EXAMPLE

C1 Bayes factor

We now consider an illustrative univariate example in which a uniform prior (U) on a parameter x between $a = 0$ and $b = 100$ results in a Gaussian posterior at $\hat{x} = 50$ with standard deviation $\hat{\sigma} = 5$. We take normal distributions with means μ and standard deviations σ as alternative priors (A) and compute the Bayes factor over the uniform prior. In this simplified case, the Bayes factor has the closed form

$$\mathcal{B}_{A/U} = \frac{b-a}{\sqrt{2\pi(\hat{\sigma}^2 + \sigma^2)}} \exp\left(-\frac{1}{2} \frac{(\hat{x} - \mu)^2}{\hat{\sigma}^2 + \sigma^2}\right). \quad (\text{C1})$$

In Fig. C1 we show the prior, posterior, and a few examples of the fiducial model in the top panel. The Bayes factor is plotted as a function of μ and σ in the bottom panel. Priors which assign more support further from the bulk of the posterior are disfavoured with respect to the original flat prior, meaning $\mathcal{B}_{A/U} < 1$. In the limit of broad models as $\sigma \rightarrow \infty$, the Gaussian priors become disfavoured since their densities are always lower than that of the flat prior. Models become favoured if well localized and centred on the posterior. The changeover occurs at $\mathcal{B}_{A/U} = 1$; this boundary (see below)

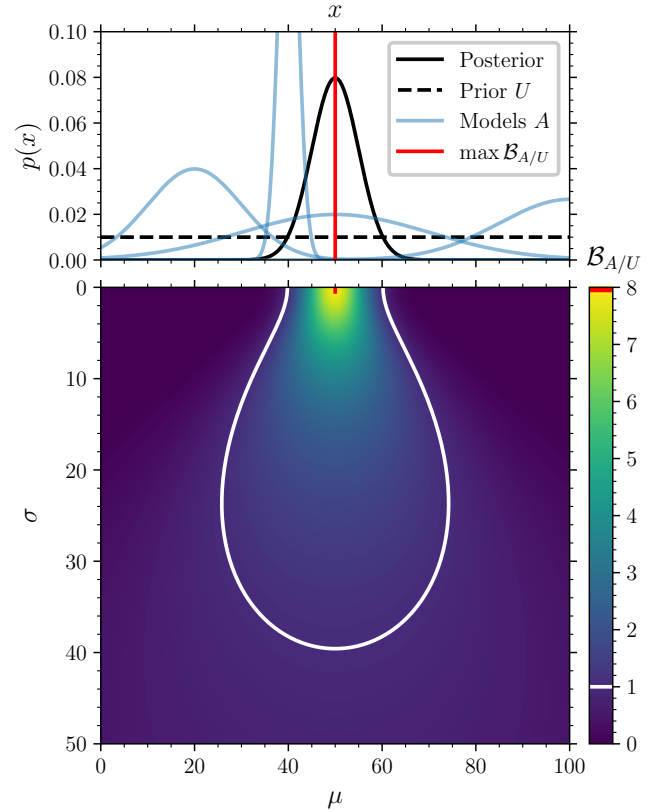


Figure C1. Bayes factors $\mathcal{B}_{A/U}$ in a simple example with a univariate parameter x (Sec. C). In the top panel we display the fiducial posterior for x (solid black curve), the assumed uniform prior from which this posterior was inferred (dashed black line), and several example Gaussians which act as alternative priors (blue curves). In red is the delta distribution corresponding to the model with the maximum Bayes factor. In the bottom panel we display the Bayes factor as a function of the mean and standard deviation of the Gaussian prior, with white marks denoting the equal preference $\mathcal{B}_{A/U} = 1$ boundary inside which the alternative prior is favoured.

is marked in white in Fig. C1, inside (outside) which the Gaussian prior is favoured (disfavoured) over the flat prior. The maximum Bayes factor occurs for a delta distribution placed at the peak of the posterior ([Payne & Thrane 2023](#)), i.e., for $\mu = \hat{\mu}$ and $\sigma = 0$, as indicated in red.

According to Eq. (C1), the $\mathcal{B}_{A/U} = 1$ boundary is given by

$$\mu = \hat{x} \pm \sqrt{2(\hat{\sigma}^2 + \sigma^2) \ln \frac{b-a}{\sqrt{2\pi(\hat{\sigma}^2 + \sigma^2)}}} \quad (\text{C2})$$

and is marked in white in Fig. C1. The largest value of σ with $\mathcal{B}_{A/U} = 1$ occurs when $\sigma = \sqrt{(b-a)^2/(2\pi) - \hat{\sigma}^2} \approx 40$ and $\mu = \hat{x} = 50$, while the largest and smallest values of μ with $\mathcal{B}_{A/U} = 1$ occur when $\sigma = \sqrt{(b-a)^2/(2\pi e) - \hat{\sigma}^2} \approx 24$ and $\mu = \hat{x} \pm (b-a)/\sqrt{2\pi e} \approx 50 \pm 24$. The $\mathcal{B}_{A/U} = 1$ turning point at $\sigma = 0$ gives $\mu \approx 40, 60$. The maximum Bayes factor model, indicated in red, is a delta distribution at $x = \hat{x}$ ([Payne & Thrane 2023](#)) with $\max \mathcal{B}_{A/U} = (b-a)/(\sqrt{2\pi}\hat{\sigma}) \approx 8$.

Note that, though our assumed posterior has support outside the prior range (which in reality is impossible), the prior boundaries are ten standard deviations away from




Figure C2. Relative heuristic fractions $\mathcal{F}_{A/U}(R)$ in a simple example for a univariate parameter x , to be compared against the Bayes factor $\mathcal{B}_{A/U}$ in Fig. C1. We assume a Gaussian posterior with mean $\hat{x} = 50$ and standard deviation $\hat{\sigma} = 5$, and uniform prior between $a = 0$ and $b = 100$. We consider alternative Gaussian priors with means μ and standard deviations σ and mark in white the equal preference $\mathcal{F}_{A/U}(R) = 1$ boundary inside which the Gaussian prior is favoured. The region R is chosen to be the $q = 79\%$ symmetric CI of the posterior.

the posterior median such that the effect is negligible. If instead we truncate the posterior on $[a, b]$, the expression for the Bayes factor that generalizes Eq. (C1) has an extra factor $[\text{erf}(\frac{b-\nu}{\sqrt{2\tau}}) - \text{erf}(\frac{a-\nu}{\sqrt{2\tau}})] / [\text{erf}(\frac{b-\hat{x}}{\sqrt{2\hat{\sigma}}}) - \text{erf}(\frac{a-\hat{x}}{\sqrt{2\hat{\sigma}}})]$, where $\nu = (\mu\hat{\sigma}^2 + \hat{x}\sigma^2) / (\hat{\sigma}^2 + \sigma^2)$, and $\tau = \hat{\sigma}\sigma / \sqrt{\hat{\sigma}^2 + \sigma^2}$. If the Gaussian priors were also truncated on $[a, b]$ then we would have $\mathcal{B}_{A/U} \rightarrow 1$ as $\sigma \rightarrow \infty$.

C2 Comparison to heuristic fraction

In Fig. C2 we show the relative heuristic fraction $\mathcal{F}_{A/U}(R)$ from Eq. (9) for the same univariate example. We choose the region R to be the symmetric 79% ($q = 0.79$) credible interval of the posterior; this value was chosen as it results in $\max \mathcal{F}_{A/U}(R) \approx \max \mathcal{B}_{A/U}$ for this specific example. The heuristic fraction can be expressed entirely in terms of error functions; from Eq. (6) we have that the fraction of prior samples from $\pi(x|A)$ inside the posterior region R is

$$f_A(R) = \frac{1}{2} \left[\text{erf} \left(\frac{\hat{x} - \mu + \sqrt{2}\hat{\sigma} \text{erf}^{-1}(q)}{\sqrt{2}\sigma} \right) + \text{erf} \left(\frac{\mu - \hat{x} + \sqrt{2}\hat{\sigma} \text{erf}^{-1}(q)}{\sqrt{2}\sigma} \right) \right]. \quad (\text{C3})$$

A similar calculation for the flat prior $\pi(x|U)$ therefore yields for the relative fraction in Eq. (9)

$$\mathcal{F}_{A/U}(R) = \frac{(b-a)f_A(R)}{2\sqrt{2}\hat{\sigma} \text{erf}^{-1}(q)}. \quad (\text{C4})$$

The relative fraction saturates much earlier for narrow models, i.e., $\sigma \approx 0$, compared to the Bayes factor, for which there is a unique maximum. There is a sharp transition from




Figure C3. Heuristic fractions in the simple univariate Gaussian example, to be compared against the Bayes factor $\mathcal{B}_{A/U}$ in Fig. C1. The left (right) column shows the fraction $f_A(R)$ ($f_{d,U}(R)$) of model (posterior) samples inside the q symmetric CI of the posterior (model). The quantile $q = 50\%, 90\%, 99\%$ is increased down the rows. White lines indicate the arbitrarily chosen $f_A(R), f_{d,U}(R) = 1/2$ boundaries.

$\mathcal{F}_{A/U}(R) = 0$ to $\mathcal{F}_{A/U}(R) = \max \mathcal{F}_{A/U}(R)$ at $\sigma = 0$ because the gradient of the posterior density is neglected in the heuristic computation; the model lies entirely within the posterior region R once μ is within the symmetric interval defined by q . The contours of constant $\mathcal{F}_{A/U}(R)$ also do not capture the full behaviour of the $\mathcal{B}_{A/U}$ contours, in particular for small σ . Otherwise, $\mathcal{F}_{A/U}(R)$ represents a reasonable approximation of $\mathcal{B}_{A/U}$ in this simple example but, most importantly, the quantile q was hand picked to result in relative fractions similar to the Bayes factor. If one were to select a different quantile, e.g., the arbitrary choice of 90%, the resulting distribution of relative heuristic fractions in Fig. C2 would be altered and, in particular, less similar to the Bayes factors in Fig. C1.

In Fig. C3 we show the values of the heuristic confidence fractions $f_A(R)$ – which counts the fraction of model A samples inside a region R of the posterior – and $f_{d,U}(R)$ – which counts the fraction of posterior samples inside a chosen region of the model prior. We again take the regions R as symmetric CIs. We display the fractions as functions of both μ and R by selecting the $q = 50\%, 90\%, 99\%$ symmetric intervals. The expression for $f_{d,U}(R)$ is found by interchange of \hat{x} and

$\hat{\sigma}$ with μ , and σ , respectively, in Eq. (C3). Figure C3 shows that these two fractions are indeed different. While $f_A(R)$ is maximized for narrow models ($\sigma \rightarrow 0$), $f_{a,U}(R)$ is maximized for broad models ($\sigma \rightarrow \infty$). In the former case, narrow alternative priors $\pi(x|A)$ place more (less) density inside R if the mean μ is inside (outside) this region. On the other hand, in this case broadening $\pi(x|A)$ places more density outside R and thus $f_A(R) \rightarrow 0$ as $\sigma \rightarrow \infty$. In the latter case, broadening the prior encapsulates more of the posterior support and thus $f_{a,U}(R) \rightarrow 1$ as $\sigma \rightarrow \infty$. The two disparate fractions agree only when there is no overlap between the alternative prior and the posterior, as seen in the top left and right corners of the panels in Fig. C3. The heuristic fraction $f_A(R)$ represents the Bayes factor much worse than the relative fraction $\mathcal{F}_{A/U}(R)$, in particular in the contours of constant $f_A(R)$. This also highlights the dependence on the arbitrary choice of R through q , which changes the variation of the heuristic fractions over the parameters μ and σ . As $q \rightarrow 1$, we find $f_A(R), f_{a,U}(R) \rightarrow 1$ since in the present example the posterior and priors have unbounded supports.