

Bridging the Gap: Categorizing Gravitational-Wave Events at the Transition Between Neutron Stars and Black Holes

AMANDA FARAH,¹ MAYA FISHBACH,^{2,*} REED ESSICK,³ DANIEL E. HOLZ,^{1,4} AND SHANIKA GALAUDAGE^{5,6}

¹*Department of Physics, University of Chicago, Chicago, IL 60637, USA*

²*Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) and Department of Physics and Astronomy, Northwestern University, 1800 Sherman Ave, Evanston, IL 60201, USA*

³*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario, Canada, N2L 2Y5*

⁴*Department of Astronomy and Astrophysics, Enrico Fermi Institute, and Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA*

⁵*School of Physics and Astronomy, Monash University, Clayton VIC 3800, Australia*

⁶*OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Clayton VIC 3800, Australia*

(Dated: November 4, 2021)

ABSTRACT

We search for features in the mass distribution of detected compact binary coalescences which signify the transition between neutron stars and black holes. We analyze all gravitational wave detections by LIGO-Virgo made through the end of the first half of the third observing run, and find clear evidence for two different populations of compact objects based solely on gravitational wave data. We confidently (99.3%) find a deviation from a single power law describing both neutron stars and low-mass black holes at $2.4_{-0.5}^{+0.5} M_{\odot}$, which is consistent with many predictions for the maximum neutron star mass. We find suggestions of the purported lower mass gap between the most massive neutron stars and the least massive black holes, but are unable to conclusively resolve it with current data. If it exists, we find the lower mass gap’s edges to lie at $2.2_{-0.5}^{+0.7} M_{\odot}$ and $6.0_{-1.4}^{+2.4} M_{\odot}$. We then analyze events that have been deemed “exceptional” by the LIGO-Virgo-KAGRA collaborations in the context of these features. We analyze GW190814 self-consistently in the context of the full population of compact binaries, finding support for its secondary mass to be either a NS or a lower mass gap object, consistent with previous claims. Our models are the first to be able to accommodate this event, which is an outlier with respect to the binary black hole population. We find that the NSBH events GW200105_162426 and GW200115_042309 probe the edges of, and may have components within, the lower mass gap. As future data improves the global population models, the classification of these events will also become more precise.

1. INTRODUCTION

The LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration (LVK) continue to expand the catalog of confidently detected Gravitational-Wave (GW) transients. To date, there have been 46 unambiguous detections of binary black hole (BBH) mergers (Abbott et al. 2019a, 2021a), two detections of binary neutron star (BNS) mergers (Abbott et al. 2017, 2020a), and, most recently, two neutron star-black hole (NSBH) merger candidates (Abbott et al. 2021b), not including

the events reported in the most recent deep extended catalog (Abbott et al. 2021c, GWTC-2.1). However, some of the detected sources cannot obviously be ascribed to one of these source categories. For example, GW190814 (Abbott et al. 2020b) has a secondary mass of $m_2 = 2.59_{-0.09}^{+0.08} M_{\odot}$, making it either the most massive neutron star (NS) or the lowest mass black hole (BH) detected to date. GW190719_215514, presented for the first time in GWTC-2.1 (Abbott et al. 2021c), is in a similar position. These events lie at the edges of both the NSBH and BBH populations and therefore have the potential to probe the extremes of whichever subpopulation they belong.

However, their classification remains elusive, and in the absence of detectable tides or electromagnetic counterparts that definitively identify their secondaries as

afarah@uchicago.edu

* NASA Hubble Fellowship Program Einstein Postdoctoral Fellow

NSs, other classification schemes are necessary. Some analyses have directly used NS equation of state (EOS) constraints to compute the probability that a given component mass is below the maximum allowed NS mass (Essick & Landry 2020; Abbott et al. 2020b, 2021b). In low latency, the LVK searches used a hard cutoff in component masses of $3 M_{\odot}$ in the third observing run (Chatterjee et al. 2020)¹ and $2.83 M_{\odot}$ in the second observing run (Abbott et al. 2019b), both motivated by EOS predictions for the maximum possible NS mass.

However, modern models for the EOS predict different values for the maximum allowed neutron star mass, typically ranging anywhere between 2.0 - $2.5 M_{\odot}$ (Legred et al. 2021). See, e.g., Chatziioannou (2020) for a review of recent observational constraints. Additionally, the maximum mass of the galactic population of NSs is currently estimated to be $\sim 2 - 2.6 M_{\odot}$ (Antoniadis et al. 2016; Alsing et al. 2018; Farr & Chatziioannou 2020; Fonseca et al. 2021) by electromagnetic observations of pulsars.

Meanwhile, electromagnetic observations of BHs in X-ray binaries suggest an absence of objects below $5 M_{\odot}$ (Bailyn et al. 1998; Özel et al. 2010; Farr et al. 2011). This mismatch between the maximum observed NS mass and minimum observed BH mass points to the possibility of a “lower mass gap”: a dearth of compact objects between ~ 2.5 and $\sim 5 M_{\odot}$. GW detections have followed suit, with a lack of BBH detections in this range. The low number of GW detections in the purported lower mass gap may be due to lower detector sensitivity to events with low masses (which are intrinsically quieter), but studies that take such selection effects into account find that the BBH distribution cannot be trivially extended to NS mass ranges. Fishbach et al. (2020) found that a single power law cannot fit the mass distribution of both BBHs and BNSs via an analysis of events in GWTC-1 (Abbott et al. 2019a). Similarly, Abbott et al. (2021d) find a deviation from a power law in the BBH spectrum below $4 M_{\odot}$ using BBHs in GWTC-2.

We pursue a different approach, simultaneously modeling the overall distribution of all compact binaries without first dividing them into subpopulations. In doing so, we may be less susceptible to issues arising from the placement of *ad hoc* boundaries. By examining the total mass distribution, we can first ask whether there is a need to subdivide the observed events into separate categories at all, such as events that fall above, below,

or within the lower mass gap, instead of asserting such subpopulations exist *a priori*.

A full description of the lower mass gap has been previously stymied by the small number of detections available. Characterizing the nature of this feature in the overall mass spectrum of compact objects can reveal the true delineation between NSs and BHs in merging binaries, providing a useful compliment to event classification based on cutoffs that are either arbitrary, motivated by external observations of arguably different populations, or reliant on theoretical predictions. With the recent influx of GW detections, enough data is available to make such classification schemes possible.

To resolve the transition between NSs and BHs, we follow the procedure of Fishbach et al. (2020) (herein FEH20) and model the full spectrum of compact binary coalescences (CBCs). We do not separate events based on previous classifications, instead using all GW events available to determine if a global fit to component masses alone is able to find two or more distinct subpopulations. Additionally, fitting a single mass model across the entire spectrum of detected CBCs allows for the inclusion of GW events that are not definitively classified into a specific source category.

We further describe our methodology and present the population models used in Sec. 2. We show that the transition between NSs and BHs cannot be described by a single power law and explore other morphologies in Sec. 3. In Sec. 4.2, we demonstrate that the models considered in this work are the first to be able to accommodate GW190814. We describe insights on the NSBH events in Sec. 4.3. We then classify all low-mass events with respect to the inferred subpopulation delineations in Sec. 4.1, concluding in Sec. 5

2. METHODS

We first describe our selection criteria for which events to include within our analysis in Sec. 2.1. We then describe our parametric models for the mass distribution in Sec. 2.2 before describing metrics that can be used to describe features in the mass spectrum, sometimes without the need for a particular parametrization, in Sec. 2.3.

2.1. Event List

We consider all published CBC detections made by the LVK to date that have a matched-filter signal-to-noise ratio (SNR) greater than 11. That is, all events in GWTC-2 (Abbott et al. 2021a) that pass this threshold, as well as the two NSBH detections from the second half of the third observing run (Abbott et al. 2021b, O3b)). None of the events newly presented in GWTC-2.1 pass

¹ See the LIGO/Virgo Public Alerts User Guide: <https://emfollow.docs.ligo.org/userguide/content.html>




Figure 1. 90% posterior credible intervals for the component masses for all CBCs included in this population study assuming uniform priors in detector-frame masses. Events classified by the LVK as BBHs, BNSs and NSBHs are shown in black, green, and blue, respectively. The ambiguously classified event GW190814 is shown in orange. The grey band indicates the approximate location of the purported “lower mass gap.” GW190814 is the only event within this region at more than 90% credibility.

this threshold (Abbott et al. 2021c). This results in

$$p(m|\lambda_{1D}) \propto h(m|m_{\min}, \eta_{\min}) \times n(m|\gamma_{\text{low}}, \gamma_{\text{high}}, \eta_{\text{low}}, \eta_{\text{high}}, A) \times \ell(m|m_{\max}, \eta_{\max}) \begin{cases} (m/m_{\text{break}})^{\alpha_1} & \text{if } m < m_{\text{break}} \\ (m/m_{\text{break}})^{\alpha_2} & \text{if } m_{\text{break}} \leq m \end{cases} \quad (1)$$

where

$$\ell(m|m_0, \eta) = \left(1 + \left(\frac{m}{m_0}\right)^\eta\right)^{-1} \quad (2)$$

is a low-pass Butterworth filter with roll-off mass m_0 and the sharpness of the roll-off set by η (larger η imply sharper roll-offs),

$$h(m|m_0, \eta) = 1 - \ell(m|m_0, \eta) = \left(1 + \left(\frac{m_0}{m}\right)^\eta\right)^{-1} \quad (3)$$

is a corresponding high-pass filter, and

$$n(m|\gamma_{\text{low}}, \gamma_{\text{high}}, \eta_{\text{low}}, \eta_{\text{high}}, A) = 1 - A \times h(m|\gamma_{\text{low}}, \eta_{\text{low}}) \times \ell(m|\gamma_{\text{high}}, \eta_{\text{high}}) \quad (4)$$

is a notch filter that subtracts a fraction of the signal (set by A) between γ_{low} and γ_{high} . Our one-dimensional mass model builds upon a basic broken power law by adding a high-pass filter at the lowest masses to allow for a fixed but smooth turn-on of the mass distribution, a low-pass filter at the highest masses to allow for a variable tapering of the mass distribution, and a notch filter in between to model a potential lower mass gap. This model has 12 parameters, which we collectively refer to as λ_{1D} . It was first presented in FEH20. However, we include an additional high-pass filter at low masses

46 events, all of which are shown in Fig. 1. For the purposes of estimating sensitivity, we treat the NSBH events as if they were detected in the first half of the third observing run because data from other detections during O3b are not yet publicly available. This choice introduces a modest bias in the inferred rate of NSBH-like events, which is discussed in Appendix B

2.2. Population Models

We consider several nested mass models. That is, we describe our conclusions in the context of different sets of assumptions about the mass distribution, each of which corresponds to a specific hyperprior as described in Tab. 1. Following Fishbach & Holz (2020), FEH20, and Doctor et al. (2020) we parametrize the joint component mass distribution with separate draws for each component mass from a one-dimensional mass distribution and a pairing function. Our one-dimension mass model can be described by a two-piece broken power law adorned with several multiplicative features that chip off parts of the mass distribution to produce dips and additional features. That is, the one-dimensional distribution is described by

instead of a sharp cut-off at m_{\min} . FEH20’s model is therefore nested in ours in the limit $\eta_{\min} \gg 1$.

To construct population models for both component masses, we employ a pairing function f_p that relates primary and secondary masses:

$$p(m_1, m_2|\lambda_{1D}, \beta_{\text{low}}, \beta_{\text{high}}) \propto \Theta(m_2 \leq m_1) \times p(m_1|\lambda_{1D}) \times p(m_2|\lambda_{1D}) \times f_p(m_1, m_2|\beta_{\text{low}}, \beta_{\text{high}}), \quad (5)$$

where $\Theta(\cdot)$ is the Heaviside function that enforces the convention $m_2 \leq m_1$.

Note that the one-dimensional distribution's hyperparameters λ_{1D} are shared between the primary and secondary mass distributions. Note also that the one-dimensional distribution does not, in general, correspond to the marginal distribution for either the primary or secondary mass. We use the pairing function

$$f_p(m_1, m_2 | \beta_{\text{low}}, \beta_{\text{high}}) = \begin{cases} (m_2/m_1)^{\beta_{\text{low}}} & \text{if } m_2 < 5 M_\odot \\ (m_2/m_1)^{\beta_{\text{high}}} & \text{if } 5 M_\odot \leq m_2 \end{cases} \quad (6)$$

to allow for the possibility that most BBHs have different formation channels than NS-containing or lower-mass-gap binaries. We note that this is a different pairing function than that used in FEH20, which employed a single power law in mass ratio for all sources. Again, their model is nested within ours in the limit $\beta_{\text{low}} = \beta_{\text{high}}$.

For simplicity, we assume fixed, independent spin and redshift distributions. We assume sources are uniformly distributed in co-moving volume (V_c) and source frame time so that

$$p(z) \propto \frac{dV_c}{dz} \left(\frac{1}{1+z} \right) \quad (7)$$

We also assume spin distributions where each component is independent and follows a distribution that is uniform in magnitude $\chi_{1,2}$ and isotropic in orientation so that

$$p(|\chi|) = U(0, 1), \quad (8)$$

$$p(z) = U(z) \quad (9)$$

where $U(x, y)$ is the uniform distribution between x and y and $z_{1,2} = \cos \theta_{1,2}$ is the cosine of the tilt angle between component spin and a binary's angular momentum.

Although some authors have shown evidence for the evolution of the mass distribution and/or the merger rate with redshift (Fishbach et al. 2021; Abbott et al. 2021d), we do not expect these choices to affect the mass inference within current statistical uncertainties.

Eqs. 1 and 6 intentionally provide a lot of flexibility in the functional form of the mass distribution via 14 free parameters while encoding several features predicted by theory. Because the current data may not be able to constrain all the hyperparameters, we consider three nested models in order to demonstrate several of our conclusions more clearly. In order of most to least model freedom,




Figure 2. Illustration of the primary mass distribution for the three mass models used in this work. All models are depicted with the same values for m_{break} , α_1 , and α_2 .

- **BROKEN POWER LAW + DIP (BPL+Dip):** assumes the edges of low-mass features are steep but leaves the locations as free parameters, and fixes the break in the power law to occur at the lower edge of the gap. Sets $\eta_{\text{min}} = \eta_{\text{high}} = 50 \gg 1$ and assumes $m_{\text{break}} = \gamma_{\text{low}}$. Alternatively, at times we assume the gap is deep with a sharp lower edge but allow the upper edge to be fit as a free parameter ($A = 0.98$ and fits η_{high}).
- **BROKEN POWER LAW + DROP (BPL+Drop):** makes similar assumptions as BPL+Dip, except for $\eta_{\text{high}} = 0$. This *de facto* models a sharp drop-off above γ_{low} and removes the feature at γ_{high} .
- **BROKEN POWER LAW (BPL):** assumes a steep turn-on at m_{min} but removes the dip/discontinuity by setting $\eta_{\text{low}} = \eta_{\text{high}} = 0$.

An illustration of these mass models is provided in Fig. 2, and Tab. 1 summarizes the hyperprior adopted for each of these models.

2.3. Characterizing the Prominence of Features in the Mass Distribution

In the case of a mass gap with sharp edges ($\eta_{\text{low,high}} \gg 1$), the notch filter strength A is an unambiguous measure of the gap depth, and the parameters γ_{low} and γ_{high} are the precise locations of the beginning and end of the gap, respectively. With this in mind, we can quantify the evidence in favor of a gap with a Bayes factor between models where $A \neq 0$ and $A = 0$: $\mathcal{B}_{A=0}^{A \neq 0}$.

However, this assumes the shape of the gap is well known *a priori*. This is not the case, particularly for the upper edge of the gap. In the case of smooth gap

Parameter		Description	Prior					
			BPL+Dip		BPL+Drop	BPL		
1D mass distribution	broken power law	m_{break}	dividing mass between high-mass and low-mass power laws		$m_{\text{break}} = \gamma_{\text{low}}$			
		α_1	spectral index of power law below m_{break}		$U(-4, 12)$			
		α_2	spectral index of power law above m_{break}		$U(-4, 12)$			
	notch filter	γ_{low}	roll-off mass for lower edge of the mass gap		$U(1.4 M_{\odot}, 3.0 M_{\odot})$			
		η_{low}	sharpness of lower edge of the mass gap		$\eta_{\text{low}} = 50$		N/A	
		γ_{high}	roll-off mass for the upper edge of the mass gap		$U(3 M_{\odot}, 9 M_{\odot})$		N/A	
		η_{high}	sharpness of the upper edge of the mass gap		$U(-4, 12)$	$\eta_{\text{high}} = 50$	$\eta_{\text{high}} = 0$	N/A
		A	depth of the notch filter		$A = 0.98$	$U(0, 1)$	$U(0, 2)^{\dagger}$	$A = 0$
	high-pass filter	m_{min}	roll-off mass for high-pass filter at the lowest masses allowed		$U(1.0 M_{\odot}, 1.4 M_{\odot})$			
		η_{min}	sharpness of high-pass filter at the lowest masses allowed		$\eta_{\text{min}} = 50$			
	low-pass filter	m_{max}	roll-off mass for low-pass filter at the highest masses allowed		$U(30 M_{\odot}, 100 M_{\odot})$			
		η_{max}	sharpness of low-pass filter at the highest masses allowed		$U(-4, 12)$			
pairing function	β_1	spectral index of pairing function if $m_2 < 5 M_{\odot}$		$U(-4, 12)$				
	β_2	spectral index of pairing function if $5 M_{\odot} \leq m_2$		$U(-4, 12)$				

Table 1. Hyperparameters of our mass model and hyperpriors corresponding to specific nested models: BROKEN POWER LAW + DIP (BPL+Dip), BROKEN POWER LAW + DROP (BPL+Drop), and BROKEN POWER LAW (BPL). We denote the uniform distribution between x and y as $U(x, y)$, list specific values that are fixed in some priors, and denote when hyperparameters are irrelevant to a specific nested model with “N/A”. Note that there are two priors that we refer to as BPL+Dip since they explore similar phenomenology: one assumes a sharp upper edge of the gap but variable gap depth and the other assumes a deep gap but allows for smooth edges. We specify which priors were used in context within the text. [†]The allowed values of A within BROKEN POWER LAW + DIP are between 0 and 1, whereas the elimination of the upper edge of the notch filter in BROKEN POWER LAW + DROP allows the values of A to range between 0 and 2 in that case.

edges, the depth and location of the gap are dictated by a more complex combination of hyperparameters. We therefore define a few straightforward measures of gap prominence which can be used regardless of the exact parametrization adopted.

The first is the existence of a local minimum in the mass distribution between 1–10 M_{\odot} . A lack of support for such a dip in the mass distribution in this region would indicate either a lack of a lower mass gap or an inability to resolve one if it does exist. Throughout this text, we report Bayes factors for the existence of a local minimum that are defined as the ratio of the percentage of draws from the hyperposterior that find a local min-

imum to the percentage of draws from the hyperprior that find a local minimum.

The second is the relative probability of events between γ_{low} and γ_{high} vs. that between m_{min} and γ_{low} :

$$r_{\text{gap}}^{\text{NS}} = \frac{p(m \in [m_{\text{min}}, \gamma_{\text{low}}] | \lambda_{1\text{D}})}{p(m \in [\gamma_{\text{low}}, \gamma_{\text{high}}] | \lambda_{1\text{D}})}. \quad (10)$$

With this definition, $r_{\text{gap}}^{\text{NS}} \gg 1$ indicates a prominent drop-off in the mass distribution at γ_{low} , whereas $r \sim 1$ is indicative of a smoother transition above NS masses.

The third is similar to the second: it is the relative probability of events between γ_{low} and γ_{high} vs. that in

an equally wide mass interval starting at γ_{high} :

$$r_{\text{gap}}^{\text{BH}} = \frac{p(m \in [\gamma_{\text{high}}, 2\gamma_{\text{high}} - \gamma_{\text{low}}] | \lambda_{\text{1D}})}{p(m \in [\gamma_{\text{low}}, \gamma_{\text{high}}] | \lambda_{\text{1D}})}. \quad (11)$$

A value of $r_{\text{gap}}^{\text{BH}} \gg 1$ indicates a very prominent lower mass gap, whereas r consistent with unity indicates a flat transition between NS and BH masses.

2.4. Statistical Framework

Using the parametrized distributions from Sec. 2.2, we construct a hierarchical Bayesian inference to determine the appropriate population-level parameters, $\Lambda = \{\lambda_{\text{1D}}, \beta_{\text{low}}, \beta_{\text{high}}\}$ given the observed set of data $\{D_j\}$ for N observed events (see, e.g., [Loredo 2004](#); [Thrane & Talbot 2019](#); [Mandel et al. 2019](#), for more

details). We model the data as an inhomogeneous Poisson process with the rate density (expected number of events per unit time per single-event-parameter hyper-volume) given by

$$\frac{d\mathcal{N}}{dm_1 dm_2, ds_1 ds_2 dz} = \mathcal{R} p(z) p(s_1) p(s_2) p(m_1, m_2 | \Lambda) \quad (12)$$

where \mathcal{R} acts as a normalizing constant that sets the overall magnitude of the rate. The posterior on the population hyper-parameters, assuming a prior on the overall rate of mergers $p(\mathcal{R}) \sim 1/\mathcal{R}$ and marginalizing, is

$$p(\Lambda | \{D_j\}) = p(\Lambda) \prod_j \frac{p(D_j | \Lambda)}{\mathcal{E}(\Lambda)}, \quad (13)$$

where

$$p(D_j | \Lambda) = \int dm_1 dm_2 ds_1 ds_2 dz p(z) p(s_1) p(s_2) p(m_1, m_2 | \Lambda) p(D_j | m_1, m_2, s_1, s_2, z) \quad (14)$$

is the marginal likelihood for the j^{th} event,

$$\mathcal{E}(\Lambda) = \int dm_1 dm_2 ds_1 ds_2 dz p(z) p(s_1) p(s_2) p(m_1, m_2 | \Lambda) P(\text{det} | m_1, m_2, s_1, s_2, z) \quad (15)$$

is the fraction of detectable events in a population described by Λ , and $P(\text{det} | m_1, m_2, s_1, s_2, z)$ is the probability that any individual event with parameters m_1, m_2, s_1, s_2 , and z would be detected, averaged over the duration of the experiment.

If we wish to simultaneously infer the properties of individual events along with the population hyperparameters, we simply do not marginalize over those events' parameters and obtain, e.g.,

$$p(m_1^{(i)}, m_2^{(i)}, \Lambda | \{D_j\}) = p(D_i | m_1^{(i)}, m_2^{(i)}) p(m_1^{(i)}, m_2^{(i)} | \Lambda) \times \frac{p(\Lambda)}{\mathcal{E}(\Lambda)} \prod_{j \neq i} \frac{p(D_j | \Lambda)}{\mathcal{E}(\Lambda)} \quad (16)$$

where $p(D_i | m_1^{(i)}, m_2^{(i)})$ is still marginalized over $s_1^{(i)}, s_2^{(i)}$, and $z^{(i)}$. We typically investigate the different levels of our hierarchical inference separately, though. That is, we examine the hyperposterior for Λ in the population-level inference after marginalizing over uncertainty in individual-event parameters, or we examine the event-level population-informed posteriors for individual event parameters after marginalizing over the uncertainty in the inferred population.

In practice, the high-dimensional integrals in Eqs. 14 and 15 are approximated via importance sampling. That is, given a set of N_j event-level posterior samples for the j^{th} event drawn with a reference prior $p_{\text{ref}}(m_1, m_2, s_1, s_2, z)$, we approximate

$$p(D_j | \Lambda) \approx \frac{1}{N_j} \sum_{\alpha} \frac{p(m_1^{(\alpha)}, m_2^{(\alpha)}, s_1^{(\alpha)}, s_2^{(\alpha)}, z^{(\alpha)} | \Lambda)}{p_{\text{ref}}(m_1^{(\alpha)}, m_2^{(\alpha)}, s_1^{(\alpha)}, s_2^{(\alpha)}, z^{(\alpha)})} \quad (17)$$

Similarly, by simulating a large set of M signals drawn from an injected population p_{inj} , we can approximate Eq. 15 with a sum over the subset of m detected signals

$$\mathcal{E}(\Lambda) \approx \frac{1}{M} \sum_{\alpha} \frac{p(m_1^{(\alpha)}, m_2^{(\alpha)}, s_1^{(\alpha)}, s_2^{(\alpha)}, z^{(\alpha)} | \Lambda)}{p_{\text{inj}}(m_1^{(\alpha)}, m_2^{(\alpha)}, s_1^{(\alpha)}, s_2^{(\alpha)}, z^{(\alpha)})} \quad (18)$$

We approximate the detectable set of signals as those that yield an optimal network signal-to-noise ratio (SNR) ≥ 11 based on networks of the LIGO Hanford and Livingston detectors during O1 and O2 as well as the LIGOs and Virgo during O3a. During each observing run, we use a single representative power spectral density (PSD) to model each detector's sensitivity separately ([LVK 2015a,b, 2017a,b, 2020](#)) and estimate the network SNR by coherently projecting simulated signals

into all detectors. While this semi-analytic approach to sensitivity estimation is commonplace in the literature, we note that it makes several assumptions that may not be perfectly accurate (a single stationary PSD for each detector, the same SNR threshold for all systems, etc.). The size of the possible systematic errors these assumptions could introduce is not known precisely, but we believe it is less than $\mathcal{O}(10\%)$ based on comparisons of our semianalytic rate estimates with fixed populations to those previously published in the literature using real searches (Abbott et al. 2019a, 2021a).

We sample from the posterior distribution in Eq. 13 using the approximations in Eqs. 17 and 18 to determine the shape of the mass distribution using `gwpopulation` (Talbot et al. 2019). Furthermore, where needed, we estimate Bayes factors via Savage-Dickey Density Ratios (Dickey & Lientz 1970; Wagenmakers et al. 2010) using the hyperposteriors and the hyperpriors described in Tab. 1.

3. POPULATION-LEVEL INSIGHTS

In this section, we search for and characterize features between NS-like and BH-like masses. Following FEH20, we first determine the existence of such a feature by asking whether a single power law can describe the CBC mass spectrum between $1 - 10M_{\odot}$ in Sec. 3.1. We find that it cannot. We then describe the nature of this deviation from a power law in Sec. 3.3 through the features provided by the parametrized models from Sec. 2.2.

3.1. Deviation from a Single Power Law Describing BH and NS Masses

Neither FEH20 nor Abbott et al. (2021d) include GW190814, nor do they include GW190425, a BNS detected in the first half of the third LVK observing run (Abbott et al. 2020a). NSBH-like events GW200105_162426 and GW200115_042309 (abbreviated GW200105 and GW200115), which contain secondaries near the upper limits of what can be considered a NS and primaries at lower BH masses than have been previously detected in definitive BBHs, had not been published at the time of either study. These events appear to bridge the NS and BH subpopulations, so the deviations from a single power law in the range $1 - 10M_{\odot}$ may no longer be present in the most recent data. We therefore aim to determine if the inclusion of these events removes the feature found by FEH20 and Abbott et al. (2021d).

The first mass model we fit is BROKEN POWER LAW. It has one power law spectral index α_1 at NS-like masses and another α_2 at BH-like masses, with a transition between the two at m_{break} . Draws from the hyperposterior for BROKEN POWER LAW are shown in the top panel

of Fig. 3. There is a clear transition in the mass distribution at $m_{\text{break}} = 2.4_{-0.5}^{+0.5} M_{\odot}$. This value is consistent with current constraints on the maximum mass of cold, non-spinning NS, which are typically between $2.0 - 2.5 M_{\odot}$ at 90% confidence (see, e.g., Essick et al. 2020; Legred et al. 2021; Farr & Chatziioannou 2020; Chatziioannou 2020), although larger values are still consistent with low-density nuclear physics.

We find $\alpha_2 - \alpha_1 > 0$ at 99.3% credibility, indicating that a single power law is unable to fit the mass distribution down to low masses. This is consistent with findings in FEH20 and Abbott et al. (2021d), and therefore the low mass events GW190814, GW200105, and GW200115 do not fully bridge the BH and NS populations.

As an aside, we also note that $\beta_2 - \beta_1 > 0$ at 99.6% credibility, indicating that the pairing functions between BBHs and CBCs with low secondary mass are distinct. However, because of the fixed location of the break between β_1 and β_2 , we do not use these parameters to locate delineations between subpopulations.

3.2. Steep Decline After NS Masses

The steepness of the NS part of the power law in BROKEN POWER LAW (i.e. large magnitude of α_1) may be driven primarily by a difference in the merger rates of BNS and BBH systems, rather than by the shape of the NS distribution itself. While our particular analysis has a slight bias in the relative heights of the low mass and high mass parts of the mass distribution because of the choice to treat the NSBH events as detected in the second half of the third observing run, many other studies have found the BNS merger rate to be several times higher than the BBH merger rate (Abbott et al. 2019a,c; Fishbach et al. 2020; Abbott et al. 2021a,d).

To decouple the shape of the NS distribution from the difference in merger rates between BNS and BBH systems, we turn to BROKEN POWER LAW + DROP. This model allows for a drop-off in the mass distribution at m_{break} . The extent of the discontinuity is parameterized by A , which is a free parameter. $A = 2$ corresponds to a maximal discontinuity: in this case, the rate of BH-containing events drops to zero after m_{break} . $A = 0$ corresponds to no discontinuity — in this case, BROKEN POWER LAW + DROP reduces to BROKEN POWER LAW. The resulting fit to this model is shown in the middle panel of Fig. 3. A clear drop-off is found after NS masses, which can be quantified by $r_{\text{gap}}^{\text{NS}}$, the ratio of probability mass below m_{break} to that in a comparable interval just above it (Eq. 10). We find $r_{\text{gap}}^{\text{NS}} = 4.1_{-3.0}^{+10}$, which excludes 1 to $> 90\%$ credibility.

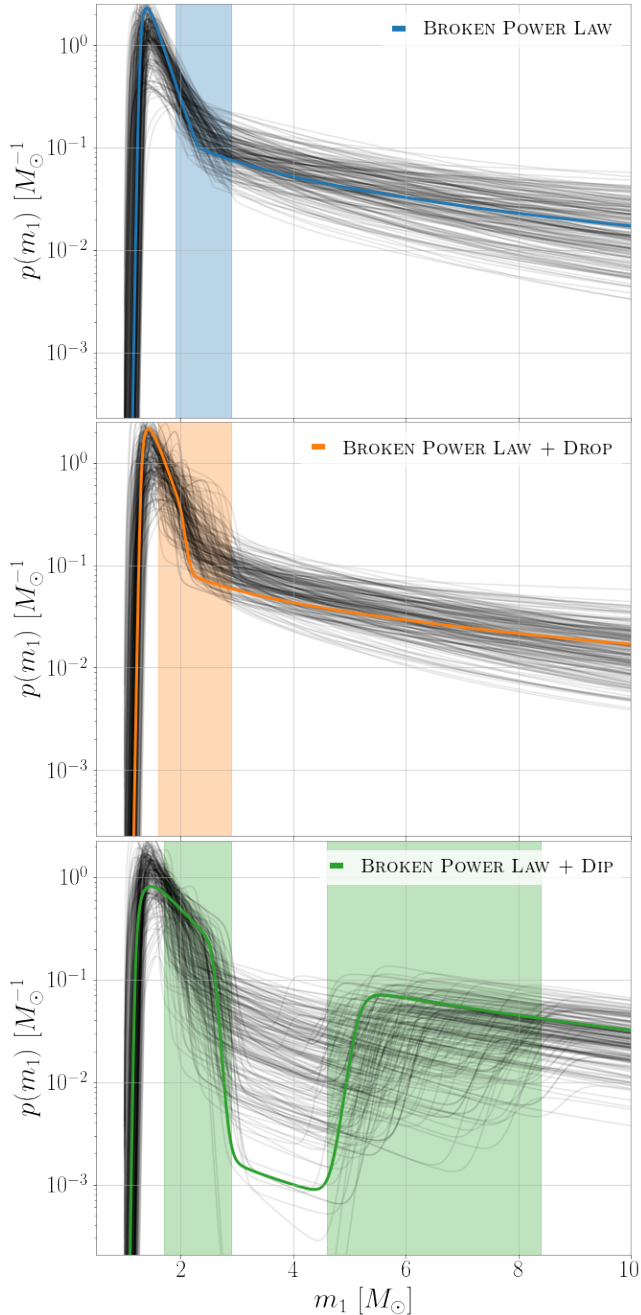


Figure 3. Draws from the inferred mass distributions between $1 - 10M_{\odot}$ using the BROKEN POWER LAW (*top*), BROKEN POWER LAW + DROP (*middle*), and BROKEN POWER LAW + DIP (*bottom*) models. Each panel has 200 random draws from the hyperposterior (*thin black lines*). As an illustration, the draw that maximizes the posterior is highlighted in color. 90% credible intervals on the locations of features in the mass distribution are indicated by colored bands. All three models find a steep falloff in the mass distribution between $\sim 2.2 - 3M_{\odot}$ and BROKEN POWER LAW + DIP finds a subsequent rise at $\sim 6M_{\odot}$.

Fig. 4 shows the inferred posteriors for A and α_1 under the BROKEN POWER LAW + DROP framework. These confirm the need for a sharp drop-off in the mass distribution at m_{break} , either by a large negative α_1 or by a large A . However, at the largest values of A , α_1 is unconstrained. We therefore conclude that the steep drop-off in the merger rate after NS masses is driving the constraints on α_1 , rather than the shape of the NS distribution itself. This is consistent with the findings of Landry & Read (2021), who show that the mass distribution of NSs in binaries is not yet able to be determined from GW data alone.

Within BROKEN POWER LAW + DROP, the location of the break in the power law (and therefore the location of the discontinuity) is inferred to be at $m_{\text{break}} = 2.20^{+0.51}_{-0.41}M_{\odot}$, which is consistent with but shifted lower than the corresponding value inferred by BROKEN POWER LAW ($2.4^{+0.5}$). This is due to the different modelling assumptions inherent in both models, and both values are equally plausible. Furthermore, given the one-dimensional posterior on A , we find that the data cannot distinguish between $A = 0$ and $A \neq 0$ and therefore have no clear preference for BROKEN POWER LAW or BROKEN POWER LAW + DROP.

3.3. Searching for a Lower Mass Gap

Having confidently established a deviation from a single power law in the form of a steep decrease in the mass distribution after NS-like masses, we wish to identify other features. In particular, we search for signatures of the purported lower mass gap.

To do this, we employ the BROKEN POWER LAW + DIP model, described in Sec. 2.2. The fit to this model is shown in the bottom panel of Fig. 3 assuming $\eta_{\text{high}} = 50$. We infer the depth of a potential mass gap through the hyperparameter A , where $A = 0$ reduces to the BROKEN POWER LAW mass model and $A = 1$ corresponds to a maximal gap. The lower and upper edges of this gap are represented by γ_{low} and γ_{high} , respectively. We infer $\gamma_{\text{low}} = 2.2^{+0.7}_{-0.5}M_{\odot}$ and $\gamma_{\text{high}} = 6.0^{+2.4}_{-1.4}M_{\odot}$. These are consistent with the values found in FEH20 ($\gamma_{\text{low}} = 2.2^{+0.6}_{-0.5}M_{\odot}$, $\gamma_{\text{high}} = 6.7^{+1.0}_{-1.5}M_{\odot}$). This agreement with a GWTC-1 analysis suggests that the addition of events in GWTC-2 and the NSBH-like events are consistent with this previously inferred population and do not fully fill the lower mass gap, if one exists. These events may, however, indicate that a lower mass gap might not be completely empty.

Within the framework of BROKEN POWER LAW + DIP, we find suggestions of a lower mass gap, but are unable to conclusively detect or rule out its existence. The posterior on A is shown in Fig. 5. It peaks at large

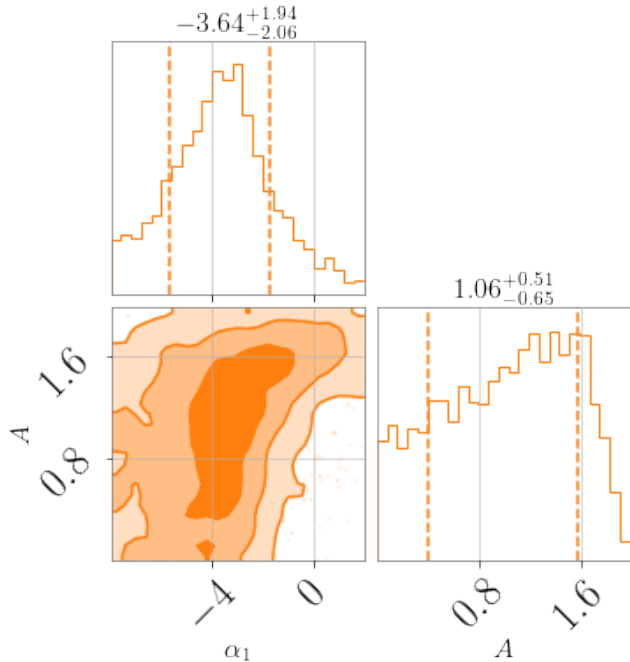


Figure 4. Values of the spectral index governing the low mass part of the spectrum α_1 and drop depth A inferred by the BROKEN POWER LAW + DROP model. Contours indicate 1- σ , 2- σ , and 3- σ credible regions. The two parameters are correlated and the lower right region is excluded: mild power law slopes are only allowed when A is large. Thus, the inference finds a dramatic drop-off in the merger rate, either through a steep negative initial power-law slope (α_1), or a discontinuity (large A). The two scenarios are both allowed by the data due to the inability to resolve the shape of the NS distribution.

values but does not rule out the possibility that some events occupy the gap. A model with a maximal gap ($A = 1$) is favored over no gap ($A = 0$) with a Bayes factor of 55.0, and $r_{\text{gap}}^{\text{BH}} = 1.84^{+5.74}_{-1.14}$, which is suggestive of the existence of a lower mass gap ($r_{\text{gap}}^{\text{BH}}$ is defined in Eq. 11).

However, we find a Bayes factor of only 1.1 in favor of a local minimum between $1 - 10 M_{\odot}$, indicating that the assumption of sharp gap edges may be driving the other metrics of gap prominence².

Due to the symmetric nature of the notch filter $n(m|A, \gamma_{\text{low}}, \gamma_{\text{high}}, \eta_{\text{low}}, \eta_{\text{high}})$ used in BROKEN POWER LAW + DIP, a sharp drop-off in the mass distribution after γ_{low} will always be accompanied by a subsequent rise at γ_{high} when we set $\eta_{\text{low}} = \eta_{\text{high}} \gg 1$. Thus, if the

² 99.5% of hyperposterior samples produce a local minimum between $1 M_{\odot}$ and $10 M_{\odot}$, but this must be compared with the percentage of prior draws that find the same feature, which is 90.5%.

data prefer such a sharp drop-off that is not achieved through a steep α_1 , using BROKEN POWER LAW + DIP may artificially inflate our metrics of gap significance. This systematic is due in part to the choice to set $\eta_{\text{low}} = \eta_{\text{high}} = 50$, rather than leaving the steepness of the gap edges as free parameters. We therefore seek to determine if the data actually prefer a subsequent rise in the mass distribution after the initial drop at γ_{low} by allowing η_{high} to be a free parameter. Positive values of η_{high} indicate the existence of a subsequent rise, and larger positive values imply steeper rises. To lessen the number of free parameters being fit to a region with relatively few detections, we enforce a drop in the mass distribution after NS masses by fixing $\eta_{\text{low}} = 50$ and $A = 0.98$. We are therefore asking, given the existence of a drop in the mass distribution at γ_{low} , do the data necessitate a subsequent rise?

Fig. 6 shows the two dimensional posteriors on γ_{high} and η_{high} , the parameters determining the location and steepness of the upper edge of the mass gap. We find $\eta_{\text{high}} > 0$ at 99.8% credibility, with lower 90% credible bound at $\eta_{\text{high}} = 2.82$, indicating that if there is a large drop in the mass distribution at γ_{low} , the data prefer a subsequent rise at γ_{high} rather than an additional drop ($\eta_{\text{high}} < 0$) or a featureless transition ($\eta_{\text{high}} = 0$). We additionally find a Bayes factor of 1.14 in favor of a local minimum between $1 - 10 M_{\odot}$.

In this model fit, $\gamma_{\text{high}} = 5.98^{+0.99}_{-0.87} M_{\odot}$, which is completely consistent with the location of the global maximum of the BBH distribution found in Abbott et al. (2021d). They find $7.8^{+1.8}_{-2.0} M_{\odot}$ for the POWER LAW + PEAK model and $6.02^{+0.78}_{-1.96} M_{\odot}$ for the BROKEN POWER LAW model, the latter of which is more morphologically similar to the high-mass behavior of the models considered in this work.

As a final test of our results, we perform a fit where all hyperparameters are allowed to vary. The results are described in Appendix A, and we find a Bayes factor of 1.63 in favor of a local minimum. This indicates slightly more evidence for a lower mass gap, but is far from conclusive. However, due to the small number of detections at low masses, we cannot meaningfully constrain more than two hyperparameters in this region. We therefore fix $\eta_{\text{high}} = 50$ and allow A to vary for the remaining analyses, noting that a steep upper edge and variable gap depth allows for more intuitive interpretations of the event classification presented in Sec. 4.1.

4. EVENT-LEVEL INSIGHTS

Building upon our population-level insights, we now place individual events in the context of the inferred features in the mass distribution. Sec. 4.1 discusses how




Figure 5. Posterior on the gap depth parameter A under the BROKEN POWER LAW + DIP model including all events. $A = 1$ indicates a maximal gap, and $A = 0$ indicates no mass gap. A peaks at 0.90, with 90% highest density interval between 0.55 and 1.0 (vertical dashed lines), indicating support for both a partially-filled and totally empty mass gap.




Figure 6. Corner plot of the parameters governing the location (γ_{high}) and steepness (η_{high}) of the upper edge of the mass gap in the case where a steep decline in the mass distribution is enforced after NS masses ($A = 0.98$). The 90% highest probability density interval for η_{high} is indicated by vertical dashed lines. η_{high} is found to be positive, indicating a rise in the mass distribution near γ_{high} after the drop enforced at γ_{low} .

we do this in general, marginalizing over both the uncertainty in individual events’ parameters as well as the population hyperparameters. We then discuss the implications for specific events that previous models have not been able to accommodate: GW190814 (Sec. 4.2) and the two NSBH-like events GW200105 and GW200115 (Sec. 4.3).

4.1. Event Classification

Fits to the full spectrum of CBCs allow us to classify events with respect to features inferred from the detected population.

Tab. 2 uses the BROKEN POWER LAW model fit to characterize events in relation to m_{break} . It lists the probabilities that each event’s component masses lie in regions of interest. For example, if one interprets m_{break} to be the maximum NS mass, the probability that an event contains a NS is $P(m_2 \leq m_{\text{break}})$ and the probability that an event is a BNS is $P(m_1 \leq m_{\text{break}})$, since $m_2 \leq m_1$ by definition. Similarly, if we consider objects with masses larger than m_{break} to be BHs, the probability that an event is a NSBH is $P(m_2 \leq m_{\text{break}} < m_1)$. Using these assumptions, we find that all probabilities presented in Tab. 2 are qualitatively consistent with the classifications given to each event by the LVK.

GW170817 is found here to have $P(m_1 \leq m_{\text{break}}) = 0.99$, which is consistent with the LVK’s classification of this event as a BNS. GW190425 has $P(m_1 \leq m_{\text{break}}) = 0.9$, with $< 10\%$ support for a NSBH classification. This is consistent with this event’s classification as a BNS. NSBH events GW200105 and GW200115 both have larger $P(m_2 \leq m_{\text{break}} < m_1)$ than their respective $P(m_1 < m_{\text{break}})$ or $P(m_2 > m_{\text{break}})$, so under this framework they are more likely NSBHs than BNSs or BBHs. The secondary mass of GW190814 is neither definitively a NS nor a BH under this classification scheme: the posterior on m_{break} overlaps almost completely with that of GW190814’s m_2 .

Tab. 3 uses the BROKEN POWER LAW + DIP model fit to characterize events in relation to γ_{low} and γ_{high} . Events that have been previously identified by the LVK as containing a NS (GW170817, GW190425, GW200105, GW200115) all have $P(m_2 \leq \gamma_{\text{low}}) > 0.95$. BNS event GW170817 has $P(m_1 \leq \gamma_{\text{low}}) > 0.95$ as well, but BNS event GW190425 has $\sim 10\%$ support of its primary being in the mass gap rather than a NS. This is likely due to the relatively large uncertainties on both the component masses of GW and on γ_{low} . Events previously identified by the LVK as being consistent with NSBHs (Abbott et al. 2021b, GW200105, GW200115) receive a similar classification in this analysis, with large $P(m_2 \leq \gamma_{\text{low}})$, and near-vanishing $P(m_1 \leq \gamma_{\text{low}})$.

Additionally, the component masses of three events have considerable support in the mass gap: the secondary of GW190814, the secondary of GW190924_021846, and the primary of GW200115. While a totally empty gap is not ruled out — none of the three have more than 75% support in the gap — the combined effect of these three events is to decrease the inferred depth of the lower mass gap. The secondary mass of GW190814 remains ambiguously classified: it has non-negligible support of being either in the lower mass gap or being a NS. As the posterior on m_2 for GW190814 is well constrained, the ambiguous classification is instead due to the inference’s inability to determine whether the event should be positioned in the first power law or in the notch filer of BROKEN POWER LAW + DIP. This results in bimodality in γ_{low} ’s posterior distribution, which we discuss further in Sec. 4.2. As more detections of GW events constrain the global mass distribution, uncertainties in the locations of γ_{low} and γ_{high} will decrease, allowing for more a more definitive classification of this event.

The probabilities in Tabs. 2 and 3 are marginalized over the uncertainties in each event’s component masses and the uncertainties in the relevant model hyperparameters.

We further examine three events that were published as “exceptional” by the LVK.

4.2. *GW190814*

GW190814 was observed in the first half of the third observing run (Abbott et al. 2020b) and has the most asymmetric mass ratio of any gravitational wave event to date. Its secondary has a mass of $m_2 = 2.59^{+0.08}_{-0.09} M_{\odot}$, leaving the nature of this component ambiguous: depending on the choice of maximum allowed NS mass, it can either be classified as a low-mass BH or the most massive NS observed to date. Previous “leave-one-out” studies (Abbott et al. 2021d; Essick et al. 2021) have definitively identified GW190814 as an outlier with respect to the detected BBH population. Additionally, Landry & Read (2021) find that including the secondary of GW190814 in the population of neutron stars in merging binaries considerably changes the inferred mass distribution when compared to a fit that excludes this event, though uncertainties in both fits are large. It is therefore desirable to find a model that can account for such an extreme event.

We perform a similar study with GW190814 by comparing the mass distributions inferred with the full event list to those inferred by excluding GW190814. We find that in the case of both BROKEN POWER LAW and BROKEN POWER LAW + DIP, the inferred distribu-

tions without GW190814 are consistent with those with GW190814. In the framework of BROKEN POWER LAW, the posterior of m_{break} is largely unchanged with the exclusion of this event, but the posterior distribution of β_1 shifts to higher values when excluding GW190814. In the framework of BROKEN POWER LAW + DIP, the posterior of γ_{low} when GW190814 is included in the inference is in some disagreement with that inferred without GW190814. When GW190814 is included, the posterior on γ_{low} is bimodal, with the lower mode corresponding to putting GW190814’s secondary in the gap and the higher mode corresponding to locating the gap at higher masses than GW190814’s secondary. Excluding GW190814 keeps only the lower mode. Additionally, the posterior on A peaks at a slightly higher value when GW190814 is excluded. However, the inferred values with and without GW190814 are still consistent with one another, as can be seen in Fig. 7. Therefore, both BROKEN POWER LAW and BROKEN POWER LAW + DIP successfully account for this event.

Note that such leave-one-out analyses can introduce biases that may lead to the false classification of the left-out event as an outlier unless modifications to the likelihood are used to mitigate this effect (Essick et al. 2021). However, since these modifications tend to bring the hyperposterior inferred without the event in question closer to that inferred with all events, we assume the conclusions derived from Fig. 7 will remain the same with the use of the methods presented in Essick et al. (2021).

4.3. *GW200105 and GW200115*

GW200105 and GW200115 are the most definitive detections of NSBH-like systems via gravitational waves to date, observed by the LVK during the second half of the third observing run (Abbott et al. 2021b). GW200105 has component masses $m_1 = 8.9^{+1.2}_{-1.5} M_{\odot}$ and $m_2 = 1.9^{+0.3}_{-0.2} M_{\odot}$, and GW200115 has component masses $m_1 = 5.7^{+1.8}_{-2.1} M_{\odot}$ and $m_2 = 1.5^{+0.7}_{-0.3} M_{\odot}$. Both events therefore have components near the edges of the purported lower mass gap, and analyzing them in the context of the BROKEN POWER LAW + DIP model is helpful in understanding their classification. Likewise, these events play a large role in probing the values of γ_{low} and γ_{high} inferred in this analysis.

We first determine if GW200105 and GW200115 are consistent with the rest of the population of CBCs by performing a model fit to BROKEN POWER LAW + DIP without these events and comparing it to the population inferred with these events. Appendix B shows these two fits to the data. We find that the two inferred populations are consistent with one another, and that the

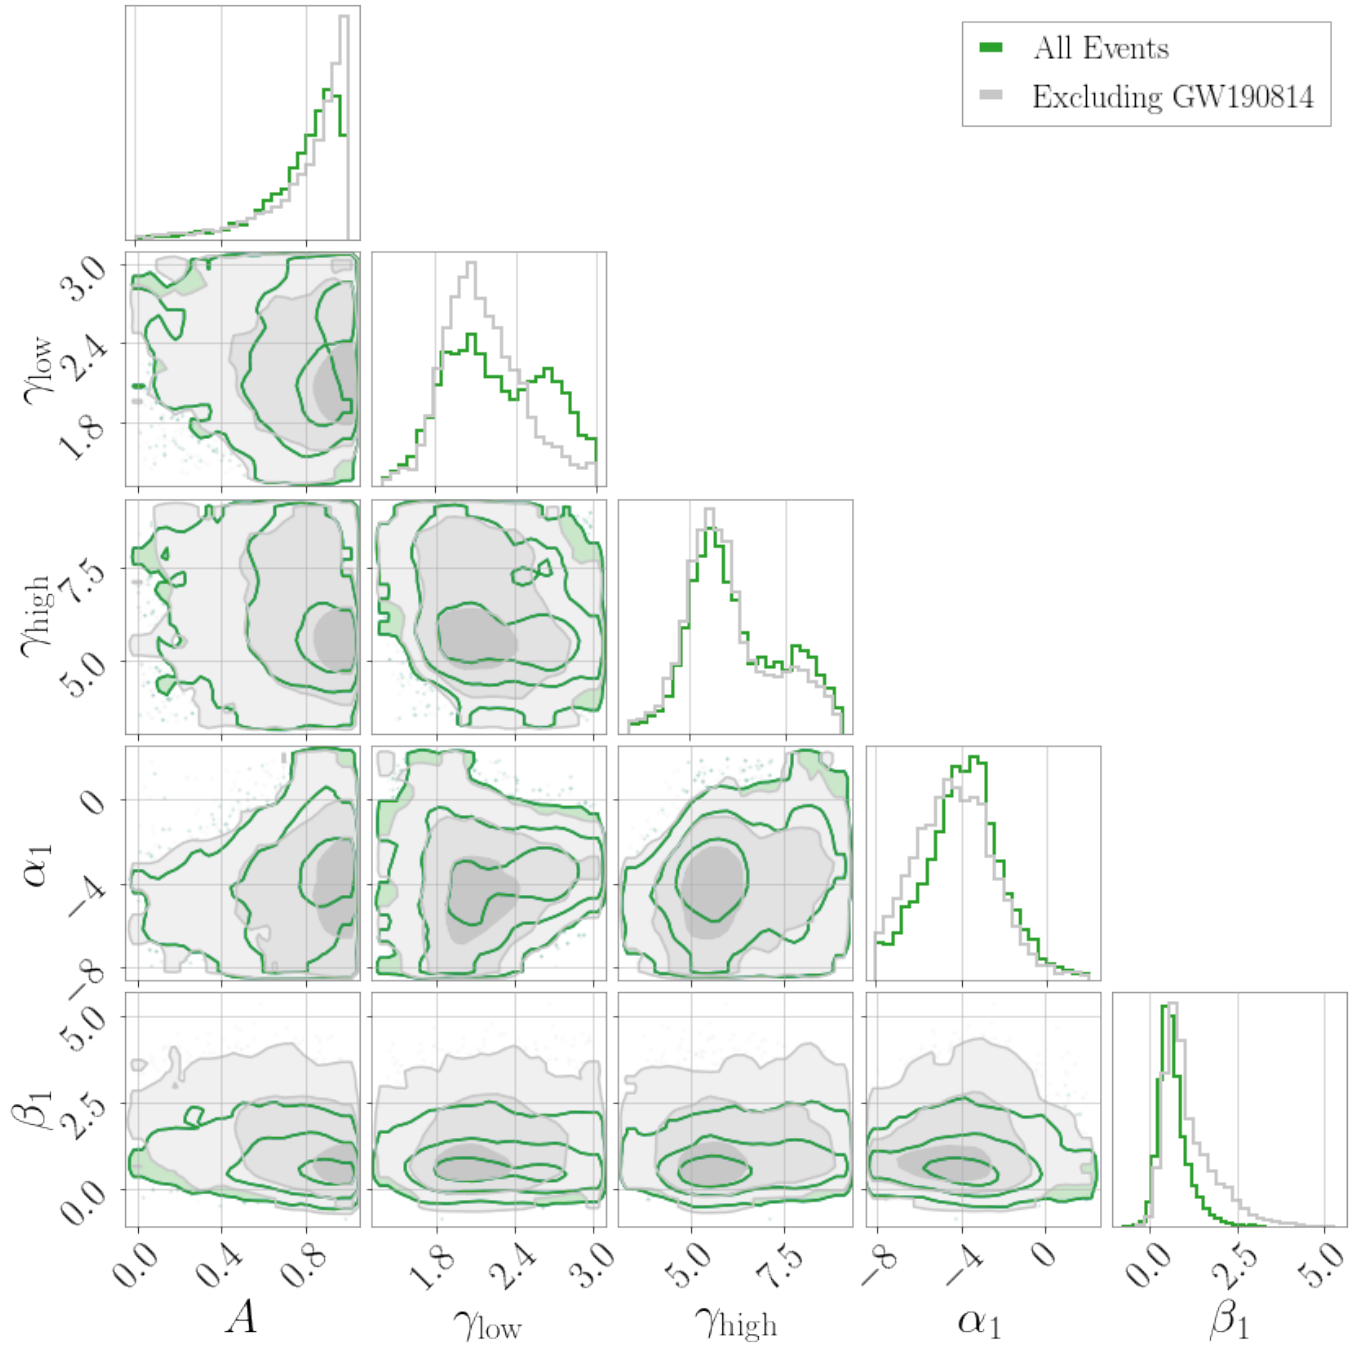


Figure 7. Effects of leaving out event GW190814 under the BROKEN POWER LAW + DIP model. Green contours correspond to the inference when GW190814 is included, and grey contours correspond to the inference when it is excluded. Almost all hyperparameters remain unchanged with the inclusion vs exclusion of GW190814, indicating that this event is consistent with the inferred population using this model. The only parameter that exhibits a noticeable change is γ_{low} , where the bimodality of the posterior distribution disappears with the exclusion of GW190814.

Probability	$P(m_2 \leq m_{\text{break}})$	$P(m_1 \leq m_{\text{break}})$	$P(m_2 \leq m_{\text{break}} < m_1)$
Interpretation	$P(\text{contains a NS})$	$P(\text{is a BNS})$	$P(\text{is a NSBH})$
GW170817	> 0.99	0.99	0.01
GW190425	> 0.99	0.90	0.09
GW190814	0.35	< 0.001	0.35
GW200105	0.88	< 0.001	0.88
GW200115	0.94	< 0.001	0.94

Table 2. Events classified in the BROKEN POWER LAW formalism. Values are the probabilities that each event’s component masses lie in a given region of parameter space, considering uncertainty in the component masses as well as in the inferred values of m_{break} . If m_{break} is interpreted as the delimiter between NS and BH masses, the probabilities can be understood as labeled in the second row. The uncertainties in these probabilities are estimated to be less than 1%. Only events with > 5% probability of containing a NS according to this classification are included in this table.

Probability	$P(m_2 \leq \gamma_{\text{low}})$	$P(m_1 \leq \gamma_{\text{low}})$	$P(\gamma_{\text{low}} \leq m_2 \leq \gamma_{\text{high}})$	$P(\gamma_{\text{low}} \leq m_1 \leq \gamma_{\text{high}})$
Interpretation	$P(\text{contains a NS})$	$P(\text{is a BNS})$	$P(m_2 \text{ in mass gap})$	$P(m_1 \text{ in mass gap})$
GW170817	> 0.99	> 0.99	< 0.001	< 0.001
GW190425	> 0.99	0.87	0.01	0.13
GW190814	0.26	< 0.001	0.74	< 0.001
GW190924_021846	< 0.001	< 0.001	0.49	0.05
GW200105	0.81	< 0.001	0.19	0.02
GW200115	0.96	< 0.001	0.04	0.44

Table 3. Events classified in the BROKEN POWER LAW + DIP formalism. Values are the probabilities that each event’s component masses lie in a given region of parameter space, considering uncertainty in the component masses as well as in the inferred values of γ_{low} and γ_{high} . If γ_{low} and γ_{high} are interpreted as the lower and upper edges of the mass gap and γ_{low} is therefore interpreted as the maximum observed NS mass in merging binaries, the probabilities can be understood as labeled in the second row. The uncertainties in these probabilities are estimated to be less than 1%. Only events with > 5% probability of being in one of these categories are included in this table.

NSBH events have the effect of constraining γ_{low} and γ_{high} towards a narrower gap. This effect, as well as the fact that both events have $P(\gamma_{\text{low}} \leq m_i \leq \gamma_{\text{high}}) < 0.5$, indicate that GW200105 and GW200115 are both consistent with being “gap-straddling binaries,” though there is still a definite possibility that either or both events have one component in the lower mass gap.

5. DISCUSSION

A lower mass gap between NSs and BHs has been suggested by observations of X-ray binaries that find a lack of black holes below $5 M_{\odot}$ (Bailyn et al. 1998; Özel et al. 2010; Farr et al. 2011) despite predictions for the maximum allowed neutron star mass by the neutron star equation of state (Legred et al. 2021; Chatziioannou 2020) of $\sim 2 - 2.5 M_{\odot}$. Such a dearth of systems just above the maximum neutron star mass is in contrast with theoretical studies of compact object formation scenarios which predict a continuous distribution of supernova remnant masses that have a smooth transition from NSs to BHs (e.g., Fryer & Kalogera 2001).

In this work, we explore whether there is evidence for a lower mass gap in the existing gravitational wave data. Inferences on the edges of such a gap have previously been carried out by examining the extrema of separate black hole and NS distributions. However, this approach lessens the amount of information available to both analyses, and can lead to biased inferences. It is insufficient to simply look at the largest neutron star mass and smallest black hole mass in the catalog, and make statements about a mass gap. *The best estimates for individual compact object masses can only be achieved through an analysis of the full population, including systems on both sides of the gap as well as those in the gap (if any).* Fitting a single model across the entire spectrum of compact binary coalescences allows us to find features that are difficult to infer from fits to the BBH, BNS, or NSBH mass distributions alone.

We present global fits to the mass distribution of NSs and BHs in merging binary systems in Fig. 3. All models considered find a sharp drop in the mass distribution at $\sim 2-3 M_{\odot}$, either via a steep power law (*top panel*) or a sharper “discontinuity” (*middle and bottom panels*). Although this feature has been determined solely from the mass distribution of GW systems, it presumably relates to the maximum mass of neutron stars in GW binaries. It is suggestive that this break occurs precisely in the mass range predicted from equation-of-state calculations Legred et al. (2021). More detections at NS masses will inform the shape of the mass spectrum below $3 M_{\odot}$, allowing us to better resolve the nature of this drop-off.

For models that allow for a feature in the mass distribution after the initial “neutron star” drop, we find a subsequent rise at $\sim 4.5-8.5 M_{\odot}$ (bottom panel of Fig. 3). While this hints at a potential mass gap between neutron stars and black holes, we cannot confidently detect nor rule out the existence of such a feature with the limited detections and low sensitivity in that region.

As we emphasize above, understanding a single GW event can only be done in the context of the full underlying population. Given our population model which bridges the gap between NSs and BHs, we are able to analyze individual events in the context of this model, and thereby classify them as BNS, NSBH, BBH, or binaries with one component in the mass gap.

Using this population-based classification scheme, we find that GW200105 likely straddles the mass gap, with its primary component confidently higher than the inferred upper edge, and the secondary likely falling below the edge, though uncertainties on the lower edge location make the exact classification uncertain. We similarly reaffirm that GW200115 is an NSBH, with its secondary component below the lower edge of the mass gap and its primary having roughly equal support below and above the upper edge. Our inference is unable to determine if GW190814’s secondary mass lies below or above the lower edge of the mass gap, but slightly prefers it to lie within the mass gap. This ambiguity is driven by uncertainty in the location of the lower edge of the gap. Thus, the nature of GW190814 may emerge as future data constrain global population models. We re-emphasize that these results are based solely on GW data; the existence of baryons has not been assumed in this analysis. Nonetheless, the conclusions are broadly consistent with expectations from NS EOS studies as well as analyses of galactic X-ray binaries.

All of the models considered in this work can accommodate the full list of observed events without any obvious outliers. In particular, we find that GW190814 is consistent with the full set of compact binary coalescences, even if we are not yet able to identify which subpopulation it belongs to.

The delineation between neutron stars and black holes is a fundamental observable of GW binary populations. Observable constraints will inform the astrophysics of compact object formation, neutron star equation-of-state studies, and the formation channels of GW binaries (e.g. Belczynski et al. 2012; Breivik et al. 2019). For example, the confident identification of a feature in the overall mass distribution that is distinct from the maximum neutron star mass inferred from nuclear theory would have important implications for specific supernova explosion mechanisms. From GW observations

alone, we have identified a sharp drop in the mass distribution of compact objects at $\sim 2\text{--}3 M_{\odot}$, as might be expected from the maximum allowed mass of neutron stars. There is insufficient data to confidently constrain the existence of a subsequent mass gap between neutron stars and black holes. However, with the second half of LVK’s third observing run completed, and a fourth observing run planned, the GW community is positioned to definitively resolve or disprove this mass gap.

ACKNOWLEDGMENTS

The authors graciously thank Colm Talbot for significant assistance with modifying and running `gwpopulation`, Mike Zevin for several helpful comments on the manuscript, and Daniel Wysocki, Richard O’Shaughnessy, and Will Farr for useful insights on early results. A.F. is supported by the NSF Research Traineeship program under grant DGE-1735359. M.F. is supported by NASA through NASA Hubble Fellowship grant HST-HF2-51455.001-A awarded by the Space Telescope Science Institute. R.E. thanks the Canadian Institute for Advanced Research (CIFAR) for support. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. D.E.H is supported by NSF grants PHY-2006645 and PHY-2110507, as well as by the Kavli Institute for Cosmological Physics through an endowment from the Kavli Foundation and its founder Fred Kavli. D.E.H also gratefully acknowledges the Marion and Stuart Rice Award. S.G. is supported through Australian Research Council (ARC) Centre of Excellence CE170100004. This material is based upon work supported by NSF LIGO Laboratory which is a major facility fully funded by the National Science Foundation. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. List of software: `gwpopulation` (Talbot et al. 2019) `bilby` (Ashton et al. 2019), `numpy` (Harris et al. 2020), `xarray` (Hoyer & Hamman 2017), `matplotlib` (Hunter 2007).

REFERENCES

- Abbott, B., Abbott, R., Abbott, T., et al. 2017, *Physical Review Letters*, 119, 161101, doi: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101)
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, *Physical Review X*, 9, 031040, doi: [10.1103/PhysRevX.9.031040](https://doi.org/10.1103/PhysRevX.9.031040)
- . 2019b, 875, 161, doi: [10.3847/1538-4357/ab0e8f](https://doi.org/10.3847/1538-4357/ab0e8f)

- . 2019c, *The Astrophysical Journal*, 882, L24, doi: [10.3847/2041-8213/ab3800](https://doi.org/10.3847/2041-8213/ab3800)
- . 2020a, *The Astrophysical Journal*, 892, L3, doi: [10.3847/2041-8213/ab75f5](https://doi.org/10.3847/2041-8213/ab75f5)
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2020b, 896, L44, doi: [10.3847/2041-8213/ab960f](https://doi.org/10.3847/2041-8213/ab960f)
- . 2021a, *Physical Review X*, 11, 021053, doi: [10.1103/PhysRevX.11.021053](https://doi.org/10.1103/PhysRevX.11.021053)
- . 2021b, *The Astrophysical Journal Letters*, 915, L5, doi: [10.3847/2041-8213/ac082e](https://doi.org/10.3847/2041-8213/ac082e)
- Abbott, R., Abbott, T. D., Acernese, F., et al. 2021c, arXiv:2108.01045 [gr-qc]. <http://arxiv.org/abs/2108.01045>
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021d, *The Astrophysical Journal Letters*, 913, L7, doi: [10.3847/2041-8213/abe949](https://doi.org/10.3847/2041-8213/abe949)
- Alsing, J., Silva, H. O., & Berti, E. 2018, *Monthly Notices of the Royal Astronomical Society*, 478, 1377, doi: [10.1093/mnras/sty1065](https://doi.org/10.1093/mnras/sty1065)
- Antoniadis, J., Tauris, T. M., Ozel, F., et al. 2016, arXiv:1605.01665 [astro-ph, physics:nucl-th]. <http://arxiv.org/abs/1605.01665>
- Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, 241, 27, doi: [10.3847/1538-4365/ab06fc](https://doi.org/10.3847/1538-4365/ab06fc)
- Bailyn, C. D., Jain, R. K., Coppi, P., & Orosz, J. A. 1998, *The Astrophysical Journal*, 499, 367, doi: [10.1086/305614](https://doi.org/10.1086/305614)
- Belczynski, K., Wiktorowicz, G., Fryer, C. L., Holz, D. E., & Kalogera, V. 2012, 757, 91, doi: [10.1088/0004-637X/757/1/91](https://doi.org/10.1088/0004-637X/757/1/91)
- Breivik, K., Chatterjee, S., & Andrews, J. J. 2019, 878, L4, doi: [10.3847/2041-8213/ab21d3](https://doi.org/10.3847/2041-8213/ab21d3)
- Chatterjee, D., Ghosh, S., Brady, P. R., et al. 2020, *The Astrophysical Journal*, 896, 54, doi: [10.3847/1538-4357/ab8dbe](https://doi.org/10.3847/1538-4357/ab8dbe)
- Chatziioannou, K. 2020, *General Relativity and Gravitation*, 52, 109, doi: [10.1007/s10714-020-02754-3](https://doi.org/10.1007/s10714-020-02754-3)
- Dickey, J. M., & Lientz, B. P. 1970, *The Annals of Mathematical Statistics*, 41, 214. <https://www.mendeley.com/catalogue/ff8e9252-ad6c-37f5-9e8f-8a48de0e2a06/>
- Doctor, Z., Wysocki, D., O’Shaughnessy, R., Holz, D. E., & Farr, B. 2020, 893, 35, doi: [10.3847/1538-4357/ab7fac](https://doi.org/10.3847/1538-4357/ab7fac)
- Essick, R., Farah, A., Galaudage, S., et al. 2021, arXiv:2109.00418 [astro-ph, physics:gr-qc]. <http://arxiv.org/abs/2109.00418>
- Essick, R., & Landry, P. 2020, *The Astrophysical Journal*, 904, 80, doi: [10.3847/1538-4357/abbd3b](https://doi.org/10.3847/1538-4357/abbd3b)
- Essick, R., Tews, I., Landry, P., Reddy, S., & Holz, D. E. 2020, *Physical Review C*, 102, 055803, doi: [10.1103/PhysRevC.102.055803](https://doi.org/10.1103/PhysRevC.102.055803)
- Farr, W. M., & Chatziioannou, K. 2020, 4, 65, doi: [10.3847/2515-5172/ab9088](https://doi.org/10.3847/2515-5172/ab9088)
- Farr, W. M., Sravan, N., Cantrell, A., et al. 2011, 741, 103, doi: [10.1088/0004-637X/741/2/103](https://doi.org/10.1088/0004-637X/741/2/103)
- Fishbach, M., Essick, R., & Holz, D. E. 2020, *The Astrophysical Journal*, 899, L8, doi: [10.3847/2041-8213/aba7b6](https://doi.org/10.3847/2041-8213/aba7b6)
- Fishbach, M., & Holz, D. E. 2020, *The Astrophysical Journal*, 891, L27, doi: [10.3847/2041-8213/ab7247](https://doi.org/10.3847/2041-8213/ab7247)
- Fishbach, M., Doctor, Z., Callister, T., et al. 2021, *The Astrophysical Journal*, 912, 98, doi: [10.3847/1538-4357/abee11](https://doi.org/10.3847/1538-4357/abee11)
- Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021, *The Astrophysical Journal Letters*, 915, L12, doi: [10.3847/2041-8213/ac03b8](https://doi.org/10.3847/2041-8213/ac03b8)
- Fryer, C. L., & Kalogera, V. 2001, *The Astrophysical Journal*, 554, 548, doi: [10.1086/321359](https://doi.org/10.1086/321359)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Hoyer, S., & Hamman, J. 2017, *Journal of Open Research Software*, 5, 10, doi: [10.5334/jors.148](https://doi.org/10.5334/jors.148)
- Hunter, J. D. 2007, *Computing in Science Engineering*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Landry, P., & Read, J. S. 2021, arXiv:2107.04559 [astro-ph, physics:gr-qc]. <http://arxiv.org/abs/2107.04559>
- Legred, I., Chatziioannou, K., Essick, R., Han, S., & Landry, P. 2021, *Physical Review D*, 104, 063003, doi: [10.1103/PhysRevD.104.063003](https://doi.org/10.1103/PhysRevD.104.063003)
- Loredo, T. J. 2004, 735, 195, doi: [10.1063/1.1835214](https://doi.org/10.1063/1.1835214)
- LVK. 2015a, L1 Calibrated Sensitivity Spectra Oct 24 2015 (Representative for O1). <https://dcc.ligo.org/LIGO-G1600151/public>
- . 2015b, H1 Calibrated Sensitivity Spectra Oct 24 2015 (Representative for O1). <https://dcc.ligo.org/LIGO-G1600150/public>
- . 2017a, L1 Calibrated Sensitivity Spectra Aug 06 2017 (Representative Best of O2 – C02, With Cleaning/Subtraction). <https://dcc.ligo.org/LIGO-G1801952/public>
- . 2017b, H1 Calibrated Sensitivity Spectra Jun 10 2017 (Representative Best of O2 – C02, With Cleaning/Subtraction). <https://dcc.ligo.org/LIGO-G1801950/public>
- . 2020, aLIGO, CAL, Official Advanced LIGO Sensitivity Plots. <https://dcc.ligo.org/LIGO-G1500623/public>
- Mandel, I., Farr, W. M., & Gair, J. R. 2019, *Monthly Notices of the Royal Astronomical Society*, 486, 1086, doi: [10.1093/mnras/stz896](https://doi.org/10.1093/mnras/stz896)

Talbot, C., Smith, R., Thrane, E., & Poole, G. B. 2019,
Physical Review D, 100, 043030,

doi: [10.1103/PhysRevD.100.043030](https://doi.org/10.1103/PhysRevD.100.043030)

Thrane, E., & Talbot, C. 2019, Publications of the
Astronomical Society of Australia, 36, e010,
doi: [10.1017/pasa.2019.2](https://doi.org/10.1017/pasa.2019.2)

Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., &
Grasman, R. 2010, Cognitive Psychology, 60, 158,
doi: [10.1016/j.cogpsych.2009.12.001](https://doi.org/10.1016/j.cogpsych.2009.12.001)

Özel, F., Psaltis, D., Narayan, R., & McClintock, J. E.
2010, 725, 1918, doi: [10.1088/0004-637X/725/2/1918](https://doi.org/10.1088/0004-637X/725/2/1918)

APPENDIX

A. MOST FLEXIBLE MODEL FIT

As a check that the parameters we choose to fix are set to values allowed by the data, we perform an analysis where all hyperparameters are simultaneously fit. Notably, we find $\gamma_{\text{low}} = 2.23_{-0.38}^{+0.43} M_{\odot}$, $\gamma_{\text{high}} = 6.34_{-1.03}^{+1.18} M_{\odot}$, and $A = 0.89_{-0.26}^{+0.08}$, all of which are broadly consistent with the corresponding hyperposteriors inferred using BROKEN POWER LAW + DIP.

A prior range for $\eta_{\{\text{high},\text{low}\}}$ that included both negative values and 50 was too large to allow convergence of our sampler given the relatively low number of detections, so we chose uniform priors between -4 and 12. Though neither hyperparameter is particularly well-measured, we find both η_{low} and η_{high} to be positive with $> 90\%$ credibility, indicating the existence of a dip in the mass distribution between γ_{low} and γ_{high} : $\eta_{\text{low}} = 7.94_{-3.63}^{+2.70}$ and $\eta_{\text{high}} = 8.54_{-3.52}^{+2.36}$. Similarly, we find a Bayes Factor of $\mathcal{B}_{\eta_{\text{high}}=0}^{\eta_{\text{high}}=12} = 13$ in favor of a rise in the mass distribution at γ_{high} . However, both hyperposteriors are approximately uniform after ~ 8 , indicating that the inference is unable to resolve the steepness of the gap edges beyond a power law that is increasing (or decreasing, in the case of η_{low}) faster than $\sim m^8$. We therefore conclude that fixing the gap edges to be relatively sharp ($\eta_{\{\text{high},\text{low}\}} = 50$) is allowed by the data.

Within this fit, we find that 80% of draws from the hyperposterior have support for a local minimum between $1M_{\odot}$ and $10M_{\odot}$, compared to 49.1% of prior draws. This gives a Bayes factor of 1.63 in favor of a local minimum, showing that a lower mass gap is slightly but not definitively preferred by the data.

B. HYPERPOSTERIOR INFERRED WITH AND WITHOUT NSBH EVENTS.

In Fig. 8 we present a corner plot of with the posteriors of relevant hyperparameters inferred by including and excluding the NSBH events GW200105 and GW200115. We find these events are broadly consistent with the rest of the CBC population. The bottom row shows the effect on the overall rate of mergers. It illustrates the effects of treating these two events as if they were detected in the first half of the third observing run. This choice introduces a systematic bias on the rate inference of 23%, which is well within the statistical uncertainties of the rate. Nonetheless, we do not quote merger rates of NSBH-like events, leaving this calculation to future analyses that are able to include the full set of GW detections in the third observing run.

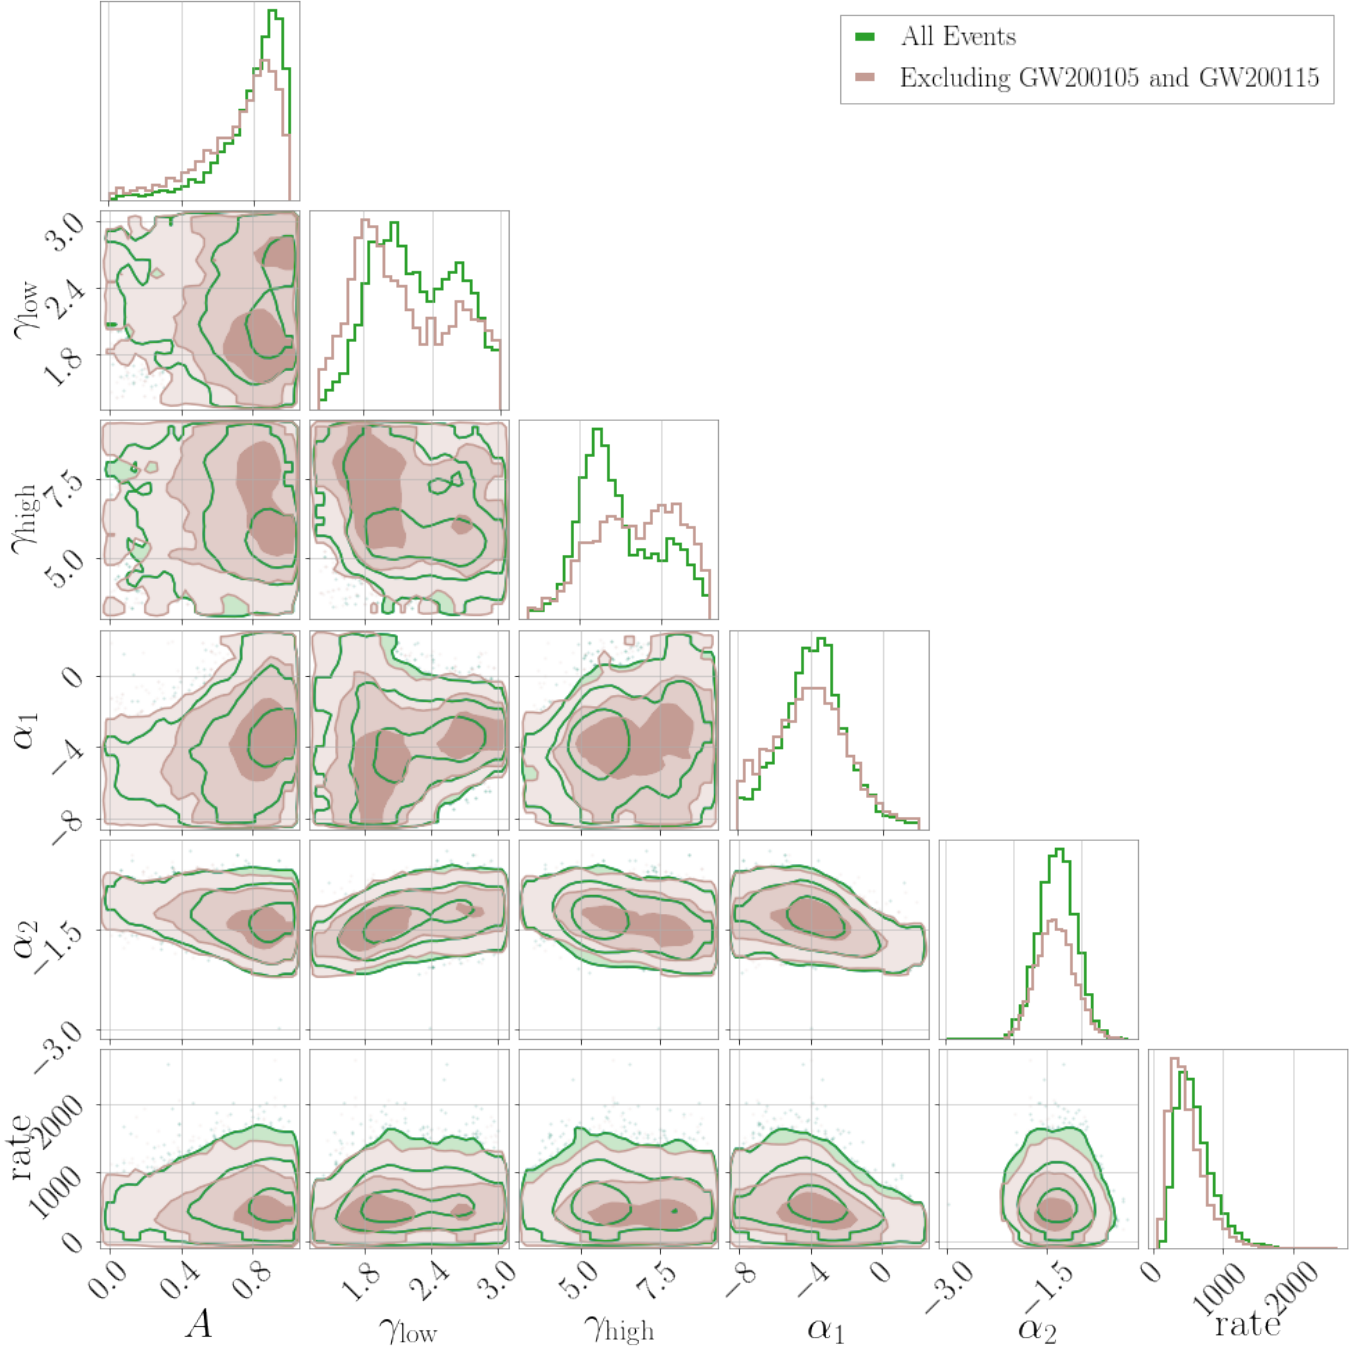


Figure 8. A subset of hyperposteriors for the BROKEN POWER LAW + DIP model inferred with (green) and without (brown) the events GW200105 and GW200115. The mass distribution does not change significantly with the exclusion of these events, indicating that the NSBHs are consistent with the rest of the detected population. When the NSBHs are included, the upper edge of the mass gap, γ_{high} , is slightly better constrained and support for a wider gap (low γ_{low} and high γ_{high}) diminishes. This suggests that these events probe the gap location. Hyperparameters on which the exclusion of the NSBH events has no effect are not included in this plot.