



Bridging the Gap: Categorizing Gravitational-wave Events at the Transition between Neutron Stars and Black Holes

Amanda Farah¹ , Maya Fishbach^{2,7} , Reed Essick³ , Daniel E. Holz^{1,4} , and Shanika Galaudage^{5,6}

¹ Department of Physics, University of Chicago, Chicago, IL 60637, USA; afarah@uchicago.edu
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, N2L 2Y5, Canada

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

Received 2021 November 19; revised 2022 February 18; accepted 2022 March 16; published 2022 May 31

Abstract

We search for features in the mass distribution of detected compact binary coalescences which signify the transition between neutron stars (NSs) and black holes (BHs). We analyze all gravitational-wave (GW) detections by the LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration (LVK) 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 GW data. We confidently (99.3%) find a steepening relative to a single power law describing NSs and low-mass BHs below $2.4_{-0.5}^{+0.5} M_{\odot}$, which is consistent with many predictions for the maximum NS mass. We find suggestions of the purported lower mass gap between the most massive NSs and the least massive BHs, 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 reexamine events that have been deemed “exceptional” by the LVK 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 to be either a NS or a lower mass gap object, consistent with previous claims. Our models are the first to accommodate this event, which is an outlier with respect to the binary BH population. We find that GW200105 and GW200115 probe the edges of, and may have components within, the lower mass gap. As future data improve global population models, the classification of these events will also improve.

Unified Astronomy Thesaurus concepts: Gravitational wave astronomy (675); Gravitational waves (678); Gravitational wave sources (677); Neutron stars (1108); Compact objects (288); Nuclear astrophysics (1129); Astrophysical black holes (98); Black holes (162); LIGO (920)

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,⁸ there have been 46 unambiguous detections of binary black hole (BBH) mergers (Abbott et al. 2019b, 2021d), 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). 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. GW190917_114630, presented for the first time in Gravitational Wave Transient Catalog 2.1 (GWTC-2.1; Abbott et al. 2021d),

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 the subpopulation to which they belong.

However, their classification remains elusive, and in the absence of detectable tides or electromagnetic counterparts that definitively identify their secondaries as 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 (Abbott et al. 2020b, 2021b; Essick & Landry 2020). In low latency, the LVK searches used a hard cutoff in component masses of $3 M_{\odot}$ in the third observing run⁹ (Chatterjee et al. 2020) 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 NS mass, typically ranging anywhere between 2.0 and $2.5 M_{\odot}$ (Margalit & Metzger 2017; Ruiz et al. 2018; Shibata et al. 2019; Legred et al. 2021). See, for example, Chatzilioannou (2020) for a review of recent observational constraints. Additionally, the maximum mass of the galactic population of NSs is currently estimated to be $\sim 2\text{--}2.6 M_{\odot}$ (Antoniadis et al. 2016; Alsing et al. 2018;

⁷ NASA Hubble Fellowship Program Einstein Postdoctoral Fellow.

⁸ These numbers do not include the events reported in the most recent deep extended catalog (GWTC-2.1; Abbott et al. 2021d), nor do they include events from GWTC-3 (Abbott et al. 2021f), as it was not released at the time of this writing.

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

⁹ See the LIGO/Virgo Public Alerts User Guide: <https://emfollow.docs.ligo.org/userguide/content.html>.

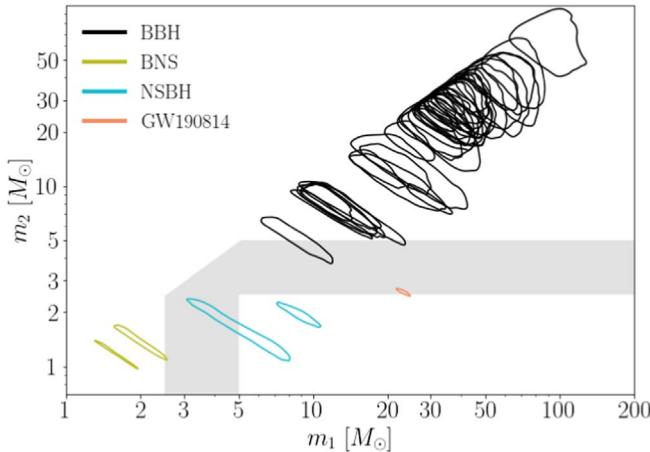


Figure 1. 90% posterior credible intervals for the component masses for all CBCs included in this population study assuming uniform priors in detector-frame masses. Events classified by the LVK as BBHs, BNSs, and NSBHs are shown in black, green, and blue, respectively. The ambiguously classified event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

Farr & Chatzioannou 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 $\sim 5 M_{\odot}$.

However, this inferred lack of low-mass BHs may be due to systematic errors in inclination measurements of X-ray binaries (Kreidberg et al. 2012). Correcting for such systematic errors populates the lower mass gap, but a relative paucity of low-mass BHs—as opposed to their complete absence—is still likely given the available observations. Additionally, some candidate mass gap objects have been identified in detached binaries (Giesers et al. 2018; Thompson et al. 2019; Jayasinghe et al. 2021) and by microlensing measurements (Wyrzykowski & Mandel 2020), though Mroz & Wyrzykowski (2021) argue that the objects identified by microlensing actually have masses below the lower edge of the mass gap. It is therefore possible that there may not be a gap in the mass spectrum of compact objects despite the existence of one in the X-ray binary mass spectrum. This may shed light on the evolutionary processes of compact object binaries (Mandel et al. 2021). Regardless, finding evidence for or against a lower mass gap in the binary merger population will add useful insight.

Interestingly, there have been very few GW detections in the lower mass gap range, as shown in Figure 1. However, 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. (2021c) found that if the BBH mass spectrum extends below $\sim 4 M_{\odot}$, a deviation from a power law starting at $\sim 8 M_{\odot}$

necessary to accommodate the BBH observations made in GWTC-2.

We pursue a different approach than Abbott et al. (2021c), simultaneously modeling the overall distribution of all compact binaries without first dividing them into subpopulations. In so doing, 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 complement 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. Fitting a single mass model across the entire spectrum of detected CBCs has the added benefit of allowing 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 Section 2. We show that the transition between NSs and BHs cannot be described by a single power law and explore other morphologies in Section 3. In Section 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 Section 4.3. We then classify all low-mass events with respect to the inferred subpopulation delineations in Section 4.1, concluding in Section 5. The appendices provide additional details on model fits.

2. Methods

We first describe our selection criteria for which events to include within our analysis in Section 2.1. We then describe our parametric models for the mass distribution in Section 2.2 before describing metrics that can be used to describe features in the mass spectrum, sometimes without the need for a particular parameterization, in Section 2.3. Section 2.4 describes our inference procedure.

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 (S/N) 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 (O3b; Abbott et al. 2021b). None of the events newly presented in GWTC-2.1 pass this threshold (Abbott et al. 2021d), and we do not consider events that are new in GWTC-3. This results in

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)

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}	$U(-4, 12)$
		γ_{low}	Roll-off mass for lower edge of the mass gap	$U(1.4 M_{\odot}, 3.0 M_{\odot})$
	Notch filter	η_{low}	Sharpness of lower edge of the mass gap	$\eta_{\text{low}} = 50$
		γ_{high}	Roll-off mass for the upper edge of the mass gap	$U(3 M_{\odot}, 9 M_{\odot})$
		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	$A = 0.99$ $U(0, 1)$ $U(0, 2)$ ^a $A = 0$
	High-pass filter	m_{min}	Roll-off mass for high-pass filter at the lowest masses allowed	$U(1.0 M_{\odot}, 1.4 M_{\odot})$
Pairing function		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	$U(-4, 12)$
		β_1	Spectral index of pairing function if $m_2 < 5 M_{\odot}$	$U(-4, 12)$
		β_2	Spectral index of pairing function if $5 M_{\odot} \leq m_2$	$U(-4, 12)$

Notes. 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.

^a 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.

46 events, all of which are shown in Figure 1. Parameter estimation for all events can be found at the GW Open Science Center (GWOSC; Abbott et al. 2021e).¹⁰

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 were not publicly available at the time of this writing. This choice introduces a modest bias in the inferred rate of NSBH-like events, which is discussed in Appendix A.

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 Table 1. Following Fishbach & Holz (2020), FEH20, and Doctor et al. (2020), we parameterize the joint component mass distribution

with separate draws for each component mass from a one-dimensional mass distribution and a pairing function. Our one-dimensional 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

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

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),

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

¹⁰ Events observed during the first and second LVK observing runs can be found at <https://dcc.ligo.org/LIGO-P1800370/public>, and this work uses the “Overall” samples for BBH events observed during this time and the samples from the IMRPhenomPv2NRT_highSpin waveform for the BNS event GW170817. Events observed during the first half of the third LVK observing runs can be found at <https://dcc.ligo.org/LIGO-P2000223/public>, and this work uses the samples from the PrecessingSpinIMRHM waveform for BBH events observed during this time and the samples from the AlignedSpinInspiralTidal_HS waveform for the BNS-like event GW190425. For events GW200105 and GW200115, this work uses the “Combined_PHM_high_spin” samples found at <https://dcc.ligo.org/LIGO-P2100143/public>.

is a corresponding high-pass filter, and

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

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 instead of a sharp cutoff at m_{min} . [FEH20](#)'s model is therefore nested in ours in the limit $\eta_{\text{min}} \gg 1$.

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

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

where $\Theta(\cdot)$ is a Heaviside function enforcing 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. The break point in this power law is fixed at $5 M_{\odot}$ and is not fit as a free parameter. This was done to avoid a scenario in which the large number of BBHs dominate the inference on the location of this break point. For example, contamination from relatively high total mass (but low mass ratio) BBHs could dominate the part of the pairing function parameterized by β_{low} . The choice of $5 M_{\odot}$ enforces a separation between the pairing function of BBHs and binaries that contain NS or lower mass gap objects. 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 comoving 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

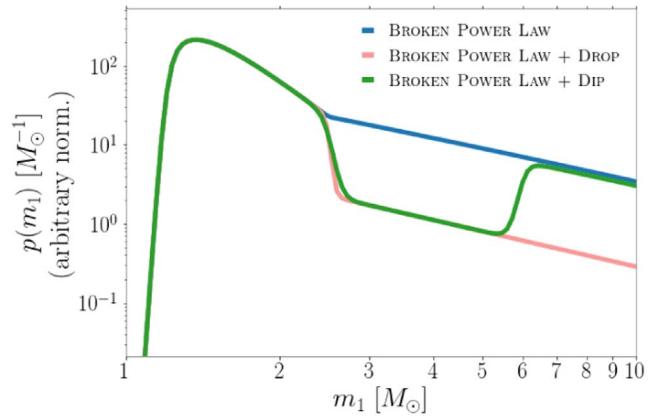


Figure 2. Illustration of the primary mass distribution for the three mass models used in this work. All models are depicted with the same values for m_{break} , α_1 , and α_2 .

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 the 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](#)), and it is unlikely that the spin magnitude distribution is uniform or that the spin tilt distribution is isotropic ([Abbott et al. 2021c](#)), we do not expect these choices to affect the mass inference within current statistical uncertainties.

Equations (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. We here present the three models in order of most to least model freedom.

1. 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 ($A = 0.99$ and fits η_{high}).
2. 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} .
3. 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 Figure 2, and Table 1 summarizes the hyperprior adopted for each.

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 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 feature prominence which can be used regardless of the exact parameterization adopted.

The first is the relative probability of events between γ_{low} and γ_{high} versus that of events between m_{min} and γ_{low} :

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

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

The second is similar: it is the relative probability of events between γ_{low} and γ_{high} versus that in an equally wide mass interval starting at γ_{high} :

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

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

To fully characterize a mass gap, one must examine both $r_{\text{gap}}^{\text{NS}}$ and $r_{\text{gap}}^{\text{BH}}$. The former only quantifies the extent of the initial drop in the mass distribution and the latter only quantifies the extent of a subsequent rise. For example, if the mass distribution were to monotonically decrease without a gap, $r_{\text{gap}}^{\text{NS}}$ would still be greater than unity. Just using this metric to identify a mass gap would be misleading. Similar effects would occur with $r_{\text{gap}}^{\text{BH}}$ if the mass distribution monotonically increased.

Furthermore, if the upper edge of the mass gap is sharp but the mass distribution falls off steeply after the gap, then $r_{\text{gap}}^{\text{BH}}$ may still be ~ 1 (within this scenario, one may not be able to distinguish between a “gap with sharp edges and rapid roll-off” or a “gentle roll-off with a narrow peak”). It is therefore the combination of $r_{\text{gap}}^{\text{NS}} > 1$, $r_{\text{gap}}^{\text{BH}} > 1$, and A that signals the presence of a gap.

We further discuss potential metrics of feature prominence in Appendix B.

2.4. Statistical Framework

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

hypervolume) 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 hyperparameters, assuming a prior on the overall rate of mergers $p(\mathcal{R}) \sim 1/\mathcal{R}$ and marginalizing, is

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

where

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

is the marginal likelihood for the j th event;

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

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.,

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

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 Equations (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

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

where $p_{\text{ref}}(D_j)$, the marginal likelihood for D_j under the reference prior, does not depend on Λ , and therefore need not be included when determining the population fit. Similarly, by simulating a large set of M signals drawn from an injected population, p_{inj} , we can approximate Equation (15) with a sum

over the subset of m detected signals:

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

We approximate the detectable set of signals as those that yield an optimal network S/N $\rho \geq 11$ based on networks of the LIGO Hanford and Livingston detectors during O1 and O2 as well as the LIGOs and Virgo during O3a. This S/N threshold was chosen to approximately match the S/N at which our chosen set of events is complete: the events included in this analysis follow the expected ρ^{-4} distribution above $\rho \sim 11$.

During each observing run, we use a single representative power spectral density (PSD) to separately model each detector's sensitivity (LVK 2015a, 2015b, 2017a, 2017b, 2020) and estimate the network S/N by coherently projecting simulated signals into all detectors. While this semianalytic 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 S/N 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 Equation (13) using the approximations in Equations (17) and (18) to determine the shape of the mass distribution using `gwpopulation` (Talbot et al. 2019) with the nested sampling algorithm `dynesty` (Speagle 2020). Furthermore, where needed, we estimate Bayes factors via Savage–Dickey density ratios (Dickey & Lientz 1970; Wagenmakers et al. 2010) using the hyperpostiors and the hyperpriors (described in Table 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 features by asking whether a single power law can describe the CBC mass spectrum between 1 and $10 M_{\odot}$ in Section 3.1. We find that it cannot. We then describe the nature of this deviation from a power law in Section 3.3 through the features provided by the parametrized models from Section 2.2. We provide the inferred posteriors for all hyperparameters in these models in Appendix C.

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

Neither FEH20 nor Abbott et al. (2021c) 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 $1\text{--}10 M_{\odot}$ 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. (2021c).

The first mass model we fit is the 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 Figure 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 a cold, nonspinning NS, which are typically between 2.0 and $2.5 M_{\odot}$ at 90% confidence (see, e.g., Chatzioannou 2020; Essick et al. 2020; Farr & Chatzioannou 2020; Legred et al. 2021), 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. The median and 90% credible limits on α_1 and α_2 are $-5.4^{+2.5}_{-2.2}$ and $-1.0^{+0.36}_{-0.4}$, respectively. This is consistent with findings in FEH20 and Abbott et al. (2021c), 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. The median and 90% credible limits on β_1 and β_2 are $0.79^{+1.2}_{-0.66}$ and $3.5^{+2.0}_{-1.6}$, respectively. 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.

Again, full posteriors on all hyperparameters in the BROKEN POWER LAW model can be found in Appendix C.

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 O3a, many other studies have found the BNS merger rate to be several times higher than the BBH merger rate (Abbott et al. 2019a, 2019c, 2021a, 2021c; Fishbach et al. 2020).

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 Figure 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 (Equation (10)). We find $r_{\text{gap}}^{\text{NS}} = 4.1^{+10}_{-3.0}$, which excludes 1 to $>90\%$ credibility.

Figure 4 shows the inferred posteriors for A and α_1 under the BROKEN POWER LAW + DROP framework. The region of parameter space with both small A and low-magnitude α_1 is ruled out, confirming the need for a sharp drop-off in the mass distribution at m_{break} . This feature can be achieved either by a

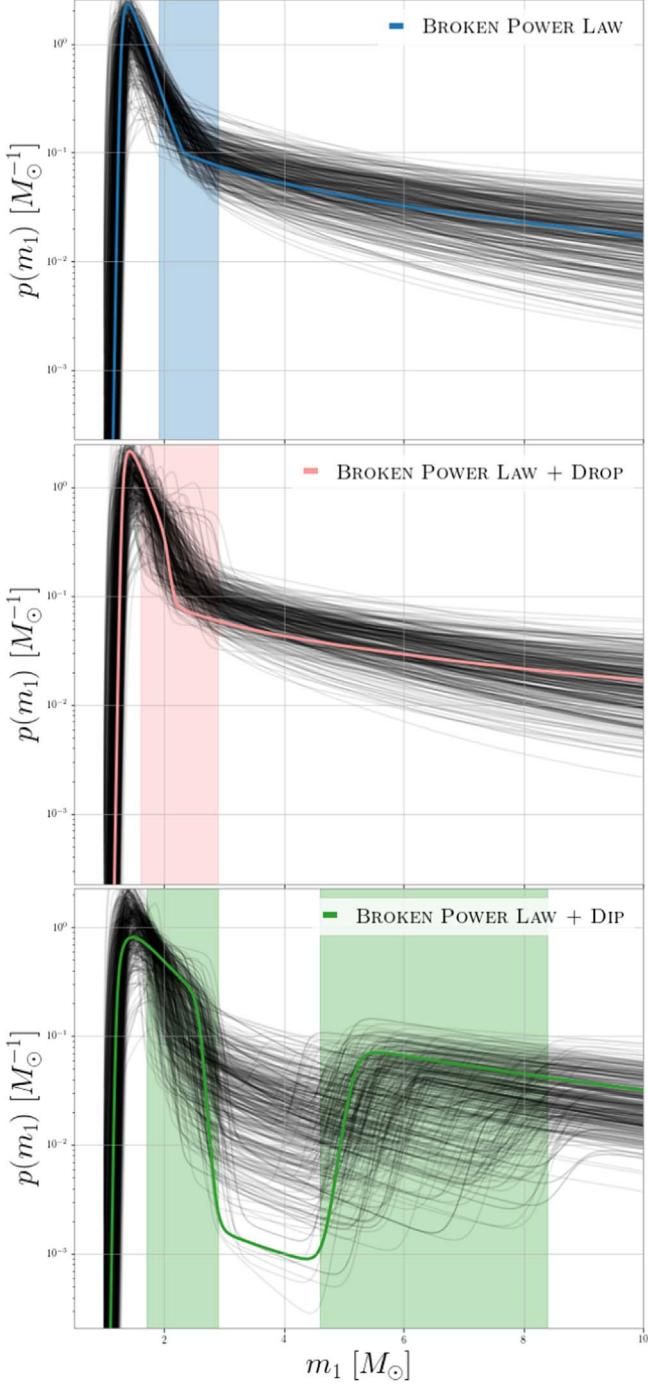


Figure 3. Draws from the inferred mass distributions between 1 and $10 M_{\odot}$ using the BROKEN POWER LAW (top), BROKEN POWER LAW + DROP (middle), and BROKEN POWER LAW + DIP (bottom) models. Each panel has 200 random draws from the hyperposterior (thin black lines). As an illustration, the draw that maximizes the overall 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 ~ 2.2 and $3 M_{\odot}$ and BROKEN POWER LAW + DIP finds a subsequent rise at $\sim 6 M_{\odot}$.

large negative α_1 , as in BROKEN POWER LAW, or by a large A . However, at the largest values of A , α_1 is unconstrained, indicating 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. We therefore conclude that BROKEN POWER LAW is sufficient to describe the limited data available.

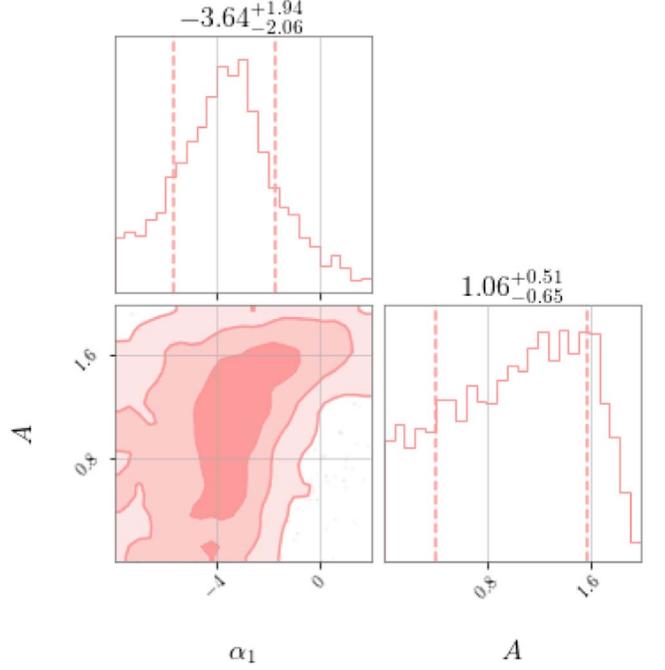


Figure 4. Values of the spectral index governing the low-mass part of the spectrum α_1 and drop depth A inferred by the BROKEN POWER LAW + DROP model. Contours indicate 1σ , 2σ , and 3σ credible regions. The two parameters are correlated and the lower-right region is excluded: mild power-law slopes are 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.

This is consistent with the findings of Landry & Read (2021), who show that the mass distribution of NSs in binaries cannot be determined from GW data alone (yet).

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 ($2.4^{+0.5}_{-0.5} M_{\odot}$). This is due to the different modeling 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. This conclusion is corroborated by the Bayes factors presented in Appendix C, which are similar between the two models. We therefore utilize BROKEN POWER LAW for the event-level analyses in this work due to its simplicity and fewer free parameters.

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 Section 2.2. The fit to this model is shown in the bottom panel of Figure 3 assuming $\eta_{\text{high}} = 50$. Note that the maximum-posterior draw (highlighted in green in Figure 3) is not necessarily required to lie in the bulk of other draws. This is true for all three models considered but is especially apparent

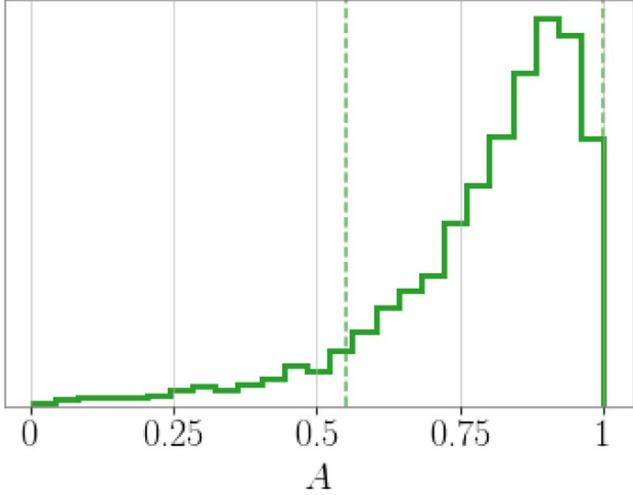


Figure 5. Posterior on the gap depth parameter A under the BROKEN POWER LAW + DIP model including all events. $A = 1$ indicates a maximal gap, and $A = 0$ indicates no mass gap. A peaks at 0.90, with 90% highest density interval between 0.55 and 1.0 (vertical dashed lines), indicating support for both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

for BROKEN POWER LAW + DIP. Values that maximize the hyperposterior for all models can be found in Appendix C.

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 $\gamma_{\text{high}} = 6.0^{+2.4512}_{-1.4} M_{\odot}$. 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 Figure 5. It peaks at large values but does not rule out the possibility that some events occupy the gap. Nevertheless, a model with a maximal gap ($A = 1$) is favored over no gap ($A = 0$) with a Bayes factor of 55.0.

However, $r_{\text{gap}}^{\text{BH}} = 1.84^{+5.74}_{-1.14}$, which suggests the existence of a lower mass gap but does not exclude $r_{\text{gap}}^{\text{BH}} = 1$. This may be due to the fact that the power-law slope governing the BH part of the spectrum is negative, which can lead to $r_{\text{gap}}^{\text{BH}} \sim 1$ even in the presence of a gap. However, it is also possible that the assumption of very sharp gap edges may be driving the preference for large A within BROKEN POWER LAW + DIP.

Due to the 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 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

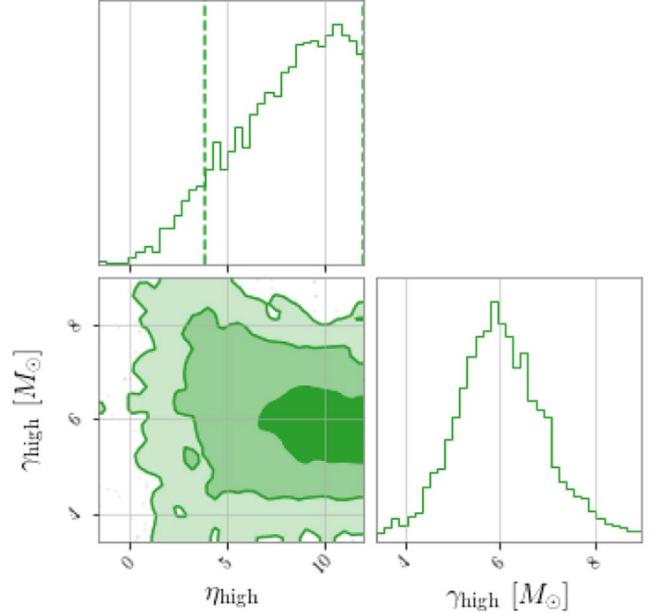


Figure 6. Corner plot of the parameters governing the location (γ_{high}) and steepness (η_{high}) of the upper edge of the mass gap in the case where a steep decline in the mass distribution is enforced after NS masses ($A = 0.99$). The 90% highest probability density interval for η_{high} is indicated 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} .

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 $A = 0.99$. We are therefore asking, given the existence of a drop in the mass distribution at γ_{low} , do the data necessitate a subsequent rise?

Figure 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 a 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 ($\eta_{\text{high}} = 0$).

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. (2021c). 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.

To determine if the assumption of $A = 0.99$ causes a subsequent rise in the mass distribution, we repeat the analysis but set $A = 0.5$. We find nearly identical results, indicating that a subsequent rise is preferred even with a moderate drop at γ_{low} .

As a final test of our results, we perform a fit where all hyperparameters are allowed to vary in Appendix B. This fit tends to prefer the existence of a lower mass gap, though the results are 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

Table 2
Events Classified in the BROKEN POWER LAW Formalism

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

Note. 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.

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 Section 4.

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. Section 4.1 discusses how 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 (Section 4.2), and the two NSBH-like events GW200105 and GW200115 (Section 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.

Table 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, marginalized over the uncertainties in each event’s component masses and the uncertainties in the relevant model hyperparameters. 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 Table 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 .

These events’ classifications in the BROKEN POWER LAW + DROP formalism are similar, though the lower inferred value of γ_{low} (and therefore m_{break}) decreases the probability that each component is below m_{break} by $\lesssim 10\%$ for the BNS and NSBH-like events, and by $\sim 20\%$ for GW190814.

Table 3 uses the BROKEN POWER LAW + DIP model fit to characterize events in relation to γ_{low} and γ_{high} . Most events that have been previously identified by the LVK as containing a NS (GW170817, GW190425, 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, and NSBH-like event GW200105 has $\sim 10\%$ support of its secondary being in the mass gap rather than an NS. This is likely due to the relatively large uncertainties on the component masses of GW190425 and GW200105 and on γ_{low} . Events previously identified by the LVK as being consistent with NSBHs (GW200105, GW200115; Abbott et al. 2021b) 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 nonnegligible 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 filter of BROKEN POWER LAW + DIP. This results in bimodality in γ_{low} ’s posterior distribution, which we discuss further in Section 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 a more definitive classification of this event.

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

4.2. GW190814

GW190814 was observed in O3a (Abbott et al. 2020b) and has the most asymmetric mass ratio of any GW event to date. Its secondary has a mass $m_2 = 2.59_{-0.09}^{+0.08} M_{\odot}$, 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. 2021c; Essick et al. 2022) 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 NSs 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.

Table 3
Events Classified in the BROKEN POWER LAW + DIP Formalism

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

Note. 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.

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 cases of BROKEN POWER LAW, BROKEN POWER LAW + DROP, and BROKEN POWER LAW + DIP, the inferred distributions without GW190814 are consistent with those with GW190814. In the frameworks of BROKEN POWER LAW and BROKEN POWER LAW + DROP, the posterior of m_{break} is largely unchanged with the exclusion of this event, but the posterior distribution of β_1 broadens and 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. Nonetheless, the inferred values with and without GW190814 are still consistent with one another, as can be seen in Figure 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. 2022). 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 Figure 7 will remain the same with the use of the methods presented in Essick et al. (2022).

4.3. GW200105 and GW200115

GW200105 and GW200115 are the most definitive detections of NSBH-like systems via GWs to date, observed by the LVK during the second half of the third observing run (Abbott et al. 2021b). GW200105 has component masses $m_1 = 8.9_{-1.5}^{+1.2} M_{\odot}$ and $m_2 = 1.9_{-0.2}^{+0.3} M_{\odot}$, and GW200115 has component masses $m_1 = 5.7_{-2.1}^{+1.8} M_{\odot}$ and $m_2 = 1.5_{-0.3}^{+0.7} M_{\odot}$. 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 A shows these two fits to the data. We find that the two inferred populations are consistent with one another, and that the NSBH events have the effect of constraining γ_{low} and γ_{high} toward 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 BHs below $5 M_{\odot}$ (Bailyn et al. 1998; Özel et al. 2010; Farr et al. 2011) despite predictions for the maximum allowed NS mass by the NS EOS (Margalit & Metzger 2017; Ruiz et al. 2018; Shibata et al. 2019; Chatzioannou 2020; Legred et al. 2021) of $\sim 2\text{--}2.5 M_{\odot}$. Such a dearth of systems just above the maximum NS 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; Fryer et al. 2012; Kushnir 2015; Drozda et al. 2020; Zevin et al. 2020).

In this work, we explore whether there is evidence for a lower mass gap in the existing GW data. Inferences on the edges of such a gap have previously been carried out by examining the extrema of separate BH 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 NS mass and smallest BH 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*

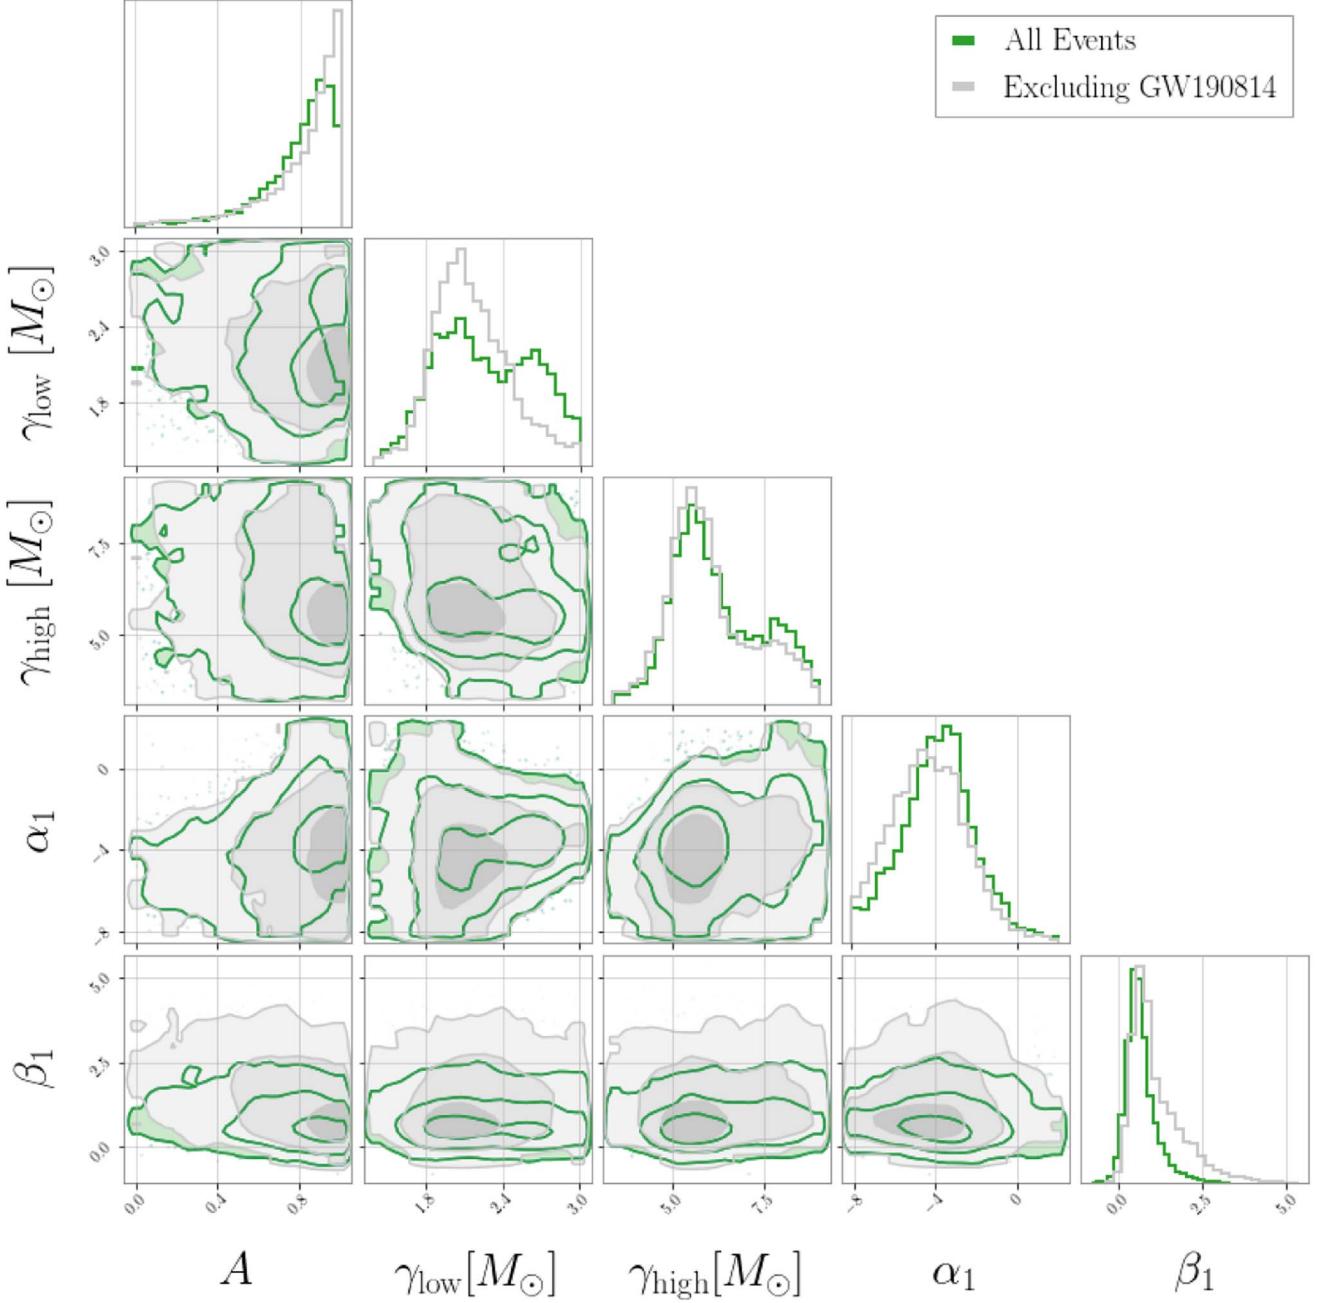


Figure 7. Effects of leaving out event GW190814 under the BROKEN POWER LAW + DIP model. Green contours correspond to the inference when GW190814 is included, and gray contours correspond to the inference when it is excluded. Almost all hyperparameters remain unchanged with the inclusion vs. exclusion 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.

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 Figure 3. All models considered find a sharp drop in the mass distribution at $\sim 2\text{--}3 M_{\odot}$, 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 NSs in GW binaries. It is suggestive that this break occurs precisely in the mass range

predicted from EOS 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 $\sim 4.5\text{--}8.5 M_{\odot}$ (bottom panel of Figure 3). While this hints at a potential mass gap between NSs and BHs, 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 reemphasize 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 to which subpopulation it belongs. In this work we focused on the component mass distribution to identify the component objects of binaries. However, future investigations of the mass ratio distribution may help us distinguish between NSBH and BBH populations, as recent population synthesis work shows that NSBH systems tend to prefer mass ratios <0.5 (Broekgaarden et al. 2021).

The delineation between NSs and BHs is a fundamental observable of GW binary populations. Resolving the nature of this transition will inform the astrophysics of compact object formation, NS EOS studies, and the formation channels of GW binaries (e.g., Fryer et al. 2002, 2012; Belczynski et al. 2012; Müller et al. 2016; Breivik et al. 2019; Liu et al. 2021; Shao & Li 2021). For example, the confident identification of a feature in the overall mass distribution that is distinct from the maximum NS mass inferred from nuclear theory would have important implications for specific supernova explosion mechanisms (Fryer et al. 2012; Kiziltan et al. 2013). 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 NSs. There is insufficient data to confidently constrain the existence of a subsequent mass gap between NSs and BHs. 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.

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 No. DGE-1735359, and by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1746045. M.F. is supported by NASA through NASA Hubble Fellowship grant No. 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 the 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 grant Nos. 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. This material is based upon work supported by the 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 grant Nos. PHY-0757058 and PHY-0823459.

Software: `gwpopulation` (Talbot et al. 2019), `dynesty` (Speagle 2020), `bilby` (Ashton et al. 2019), `numpy` (Harris et al. 2020), `xarray` (Hoyer & Hamman 2017), `matplotlib` (Hunter 2007).

Appendix A

Hyperposteriors Inferred with and without NSBH Events

In Figure 8 we present a corner plot of the posteriors for relevant hyperparameters inferred by including and excluding the NSBH events GW200105 and GW200115. We find these events to be 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 (O3a). 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.

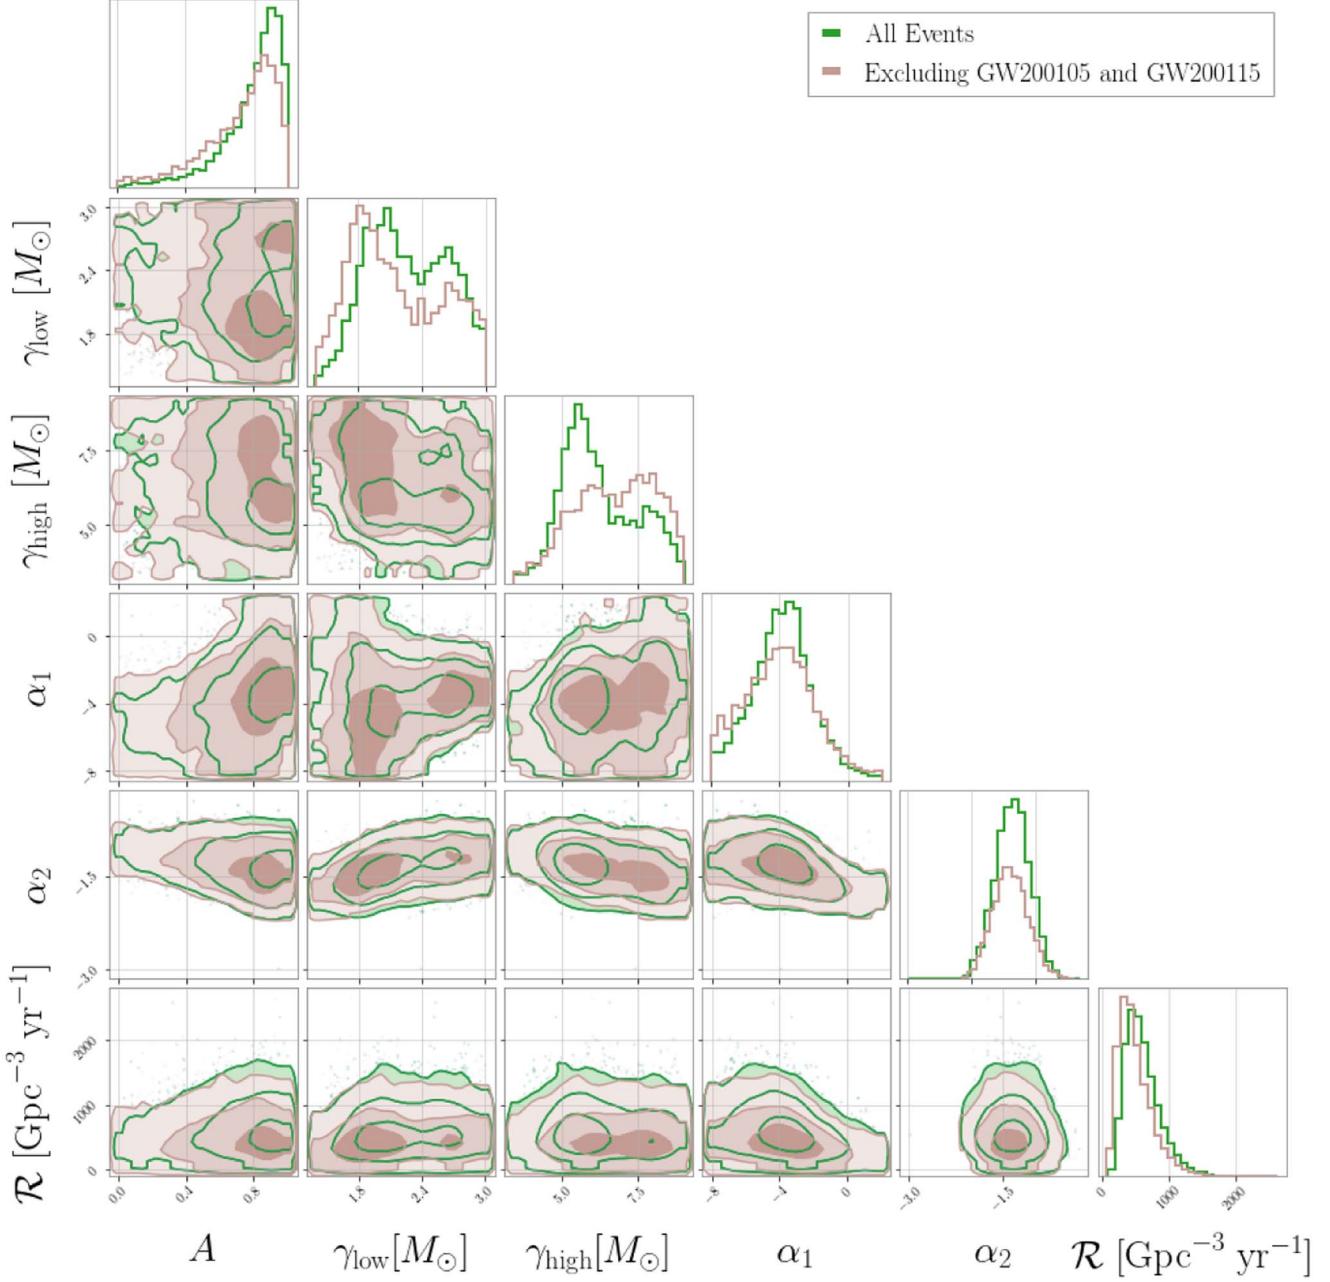


Figure 8. A subset of hyperposterior for the BROKEN POWER LAW + DIP model inferred with (green) and without (brown) the events GW200105 and GW200115. The mass distribution does not change significantly with the exclusion of these events, indicating that the NSBHs are consistent with the rest of the detected 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.

Appendix B Most Flexible Model Fit

As a final check of our results, we perform an analysis where all hyperparameters are simultaneously fit. A corner plot for this analysis is shown in Figure 9. Notably, we find $\gamma_{\text{low}} = 2.23^{+0.43}_{-0.38} M_{\odot}$, $\gamma_{\text{high}} = 6.34^{+1.18}_{-1.03} M_{\odot}$, and $A = 0.89^{+0.08}_{-0.26}$, all of which are broadly consistent with the corresponding hyperposterior 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} : $\eta_{\text{low}} = 7.94^{+2.70}_{-3.63}$ and $\eta_{\text{high}} = 8.54^{+2.36}_{-3.52}$. 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 hyperposterior are approximately uniform above ~ 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

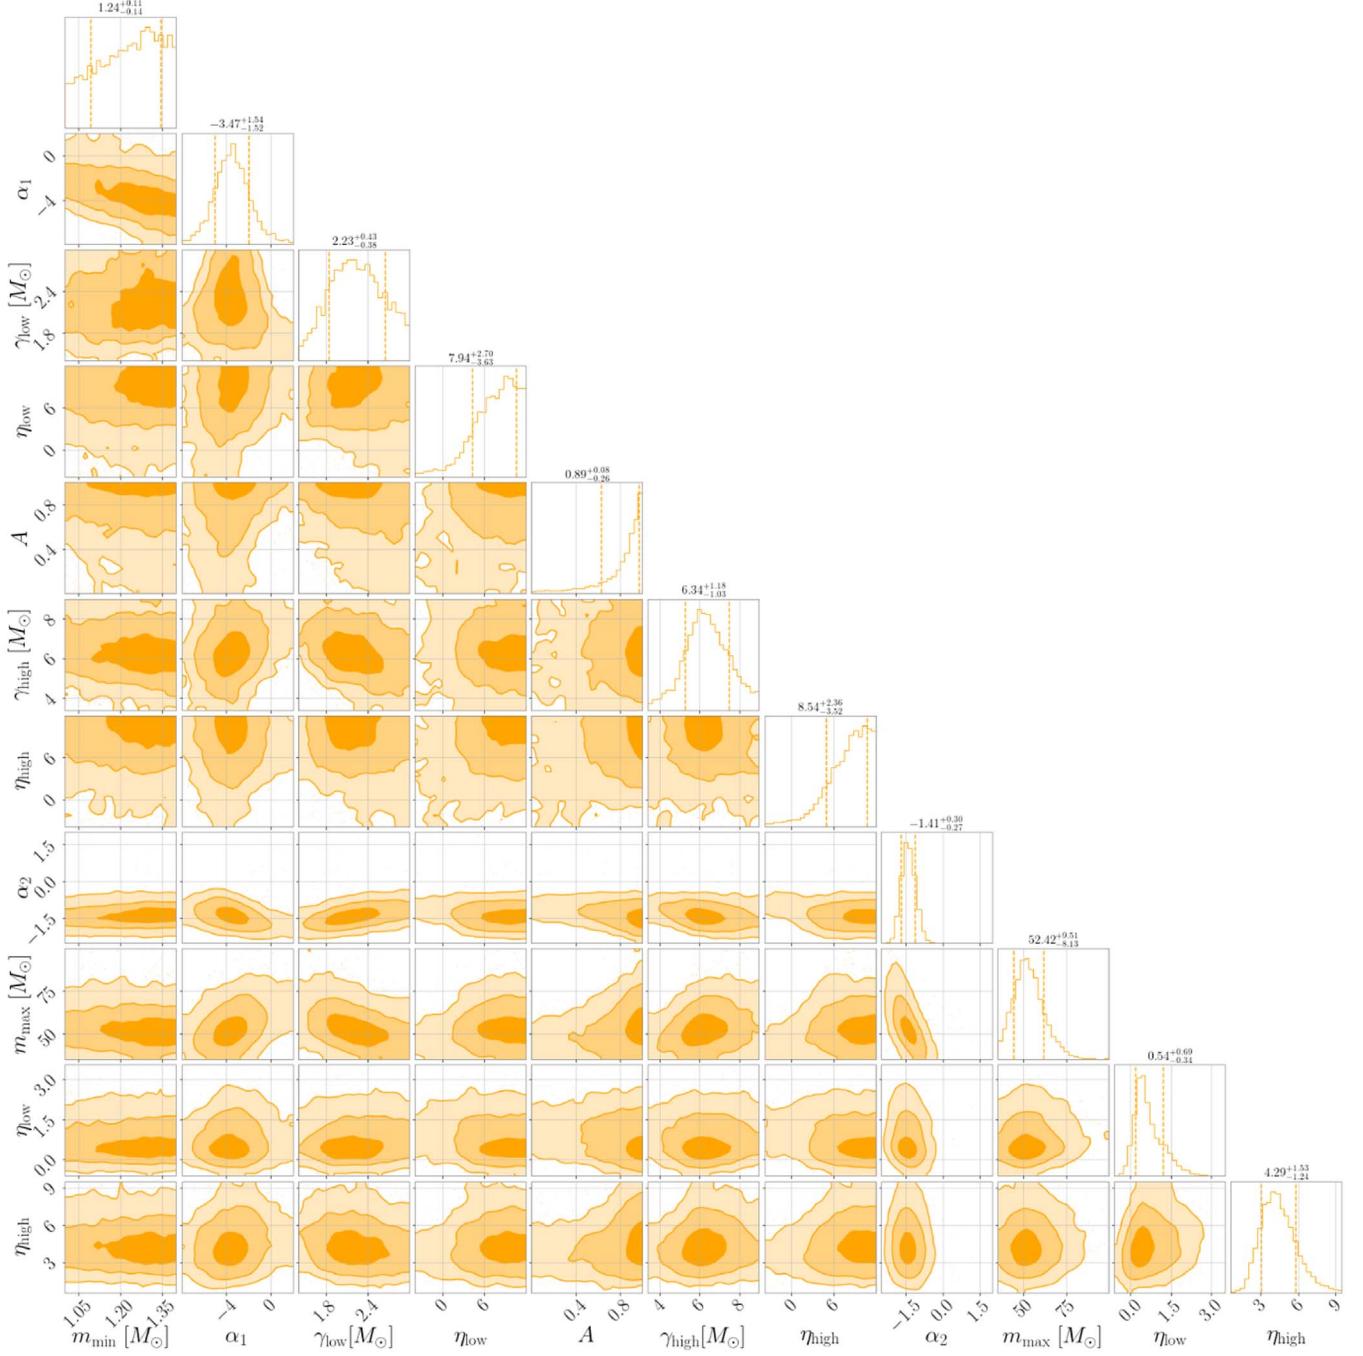


Figure 9. Corner plot for an analysis where all hyperparameters are simultaneously fit rather than set to the values in Table 1. On the diagonal are the one-dimensional marginal distributions for each hyperposterior. Above each column are written the median and 68% credible intervals for the corresponding hyperposterior.

edges to be relatively sharp ($\eta_{\{\text{high,low}\}} = 50$) is allowed by the data. We additionally find that η_{\min} peaks near the upper end of the prior bound, validating that a steep turn-on at the lowest masses is preferred by the data.

As discussed in Section 2.3, the depth and location of the mass gap are not well defined in terms of one hyperparameter when the edges of the gap are allowed to be smooth. It is therefore desirable to employ metrics for the existence of a mass gap that are both easily interpreted and independent of model parameters. These could be used for any mass model, not just the ones presented here.

Potentially the simplest such metric is the existence of a local minimum in the mass distribution between 1 and $10 M_{\odot}$. A lack of support for such a dip in the mass distribution in this region may indicate either a lack of a lower mass gap or an inability to resolve one if it does exist. We can therefore construct 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 minimum to the percentage of draws from the hyperprior that find a local minimum.

Within this most flexible model fit, we find that 80% of draws from the hyperposterior have support for a local

minimum between $1 M_{\odot}$ and $10 M_{\odot}$, 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.

An important caveat to this statistic, however, is that the existence of a local minimum does not always map to the existence of a mass gap. For example, a shallow-but-declining power law meeting a shallow-but-increasing power law would result in a local minimum, but not a feature that one would expect if there were an astrophysically induced dearth of events within a certain range. This phenomenon inflates the number of prior draws that produce a local minimum relative to those that produce a mass gap feature,

thus decreasing the resulting Bayes factor if the posterior does in fact resemble a mass gap.

More complex metrics can be constructed to mitigate this problem, but most are subject to similar issues and are less easily interpreted. Therefore, we caution against using such statistics as the single assessment of feature significance, and instead recommend they be paired with complementary probes that are more well suited to specific models.

Appendix C Model Comparison

In Figures 10–12 we show corner plots of BROKEN POWER LAW, BROKEN POWER LAW + DROP, and BROKEN POWER

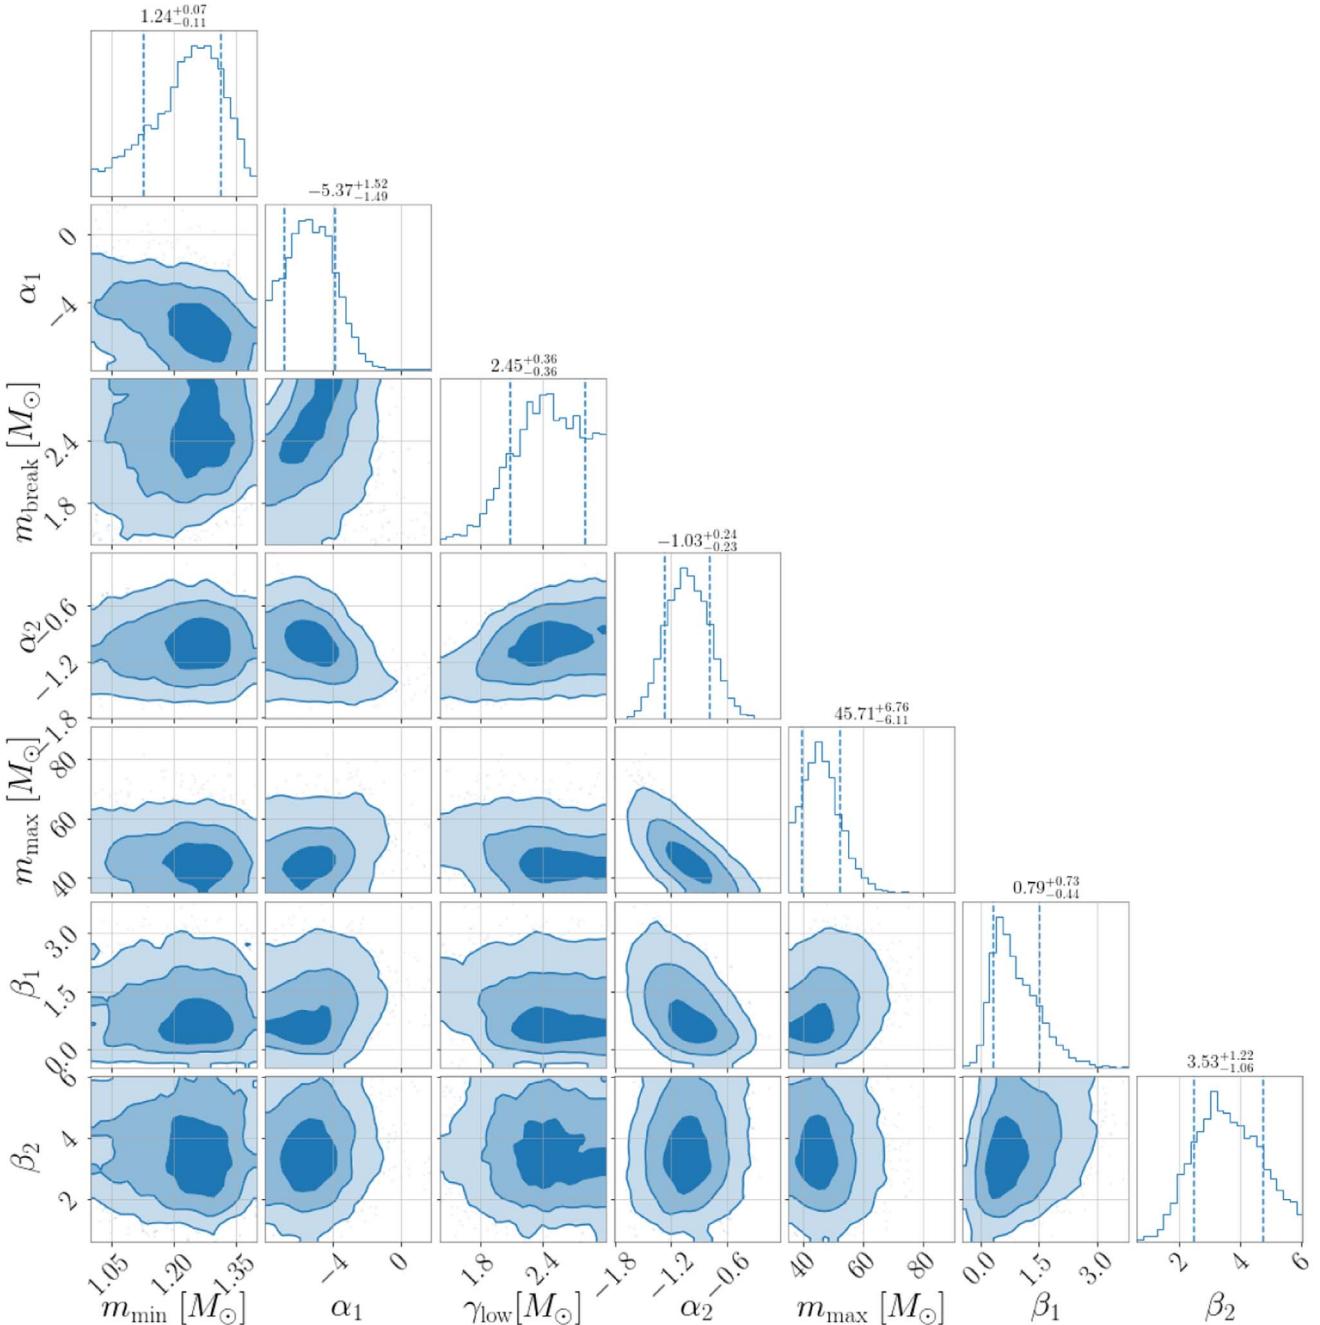


Figure 10. Corner plot for the BROKEN POWER LAW model. On the diagonal are the one-dimensional marginal distributions for each hyperposterior. Above each column are written the median and 68% credible intervals for the corresponding hyperposterior.

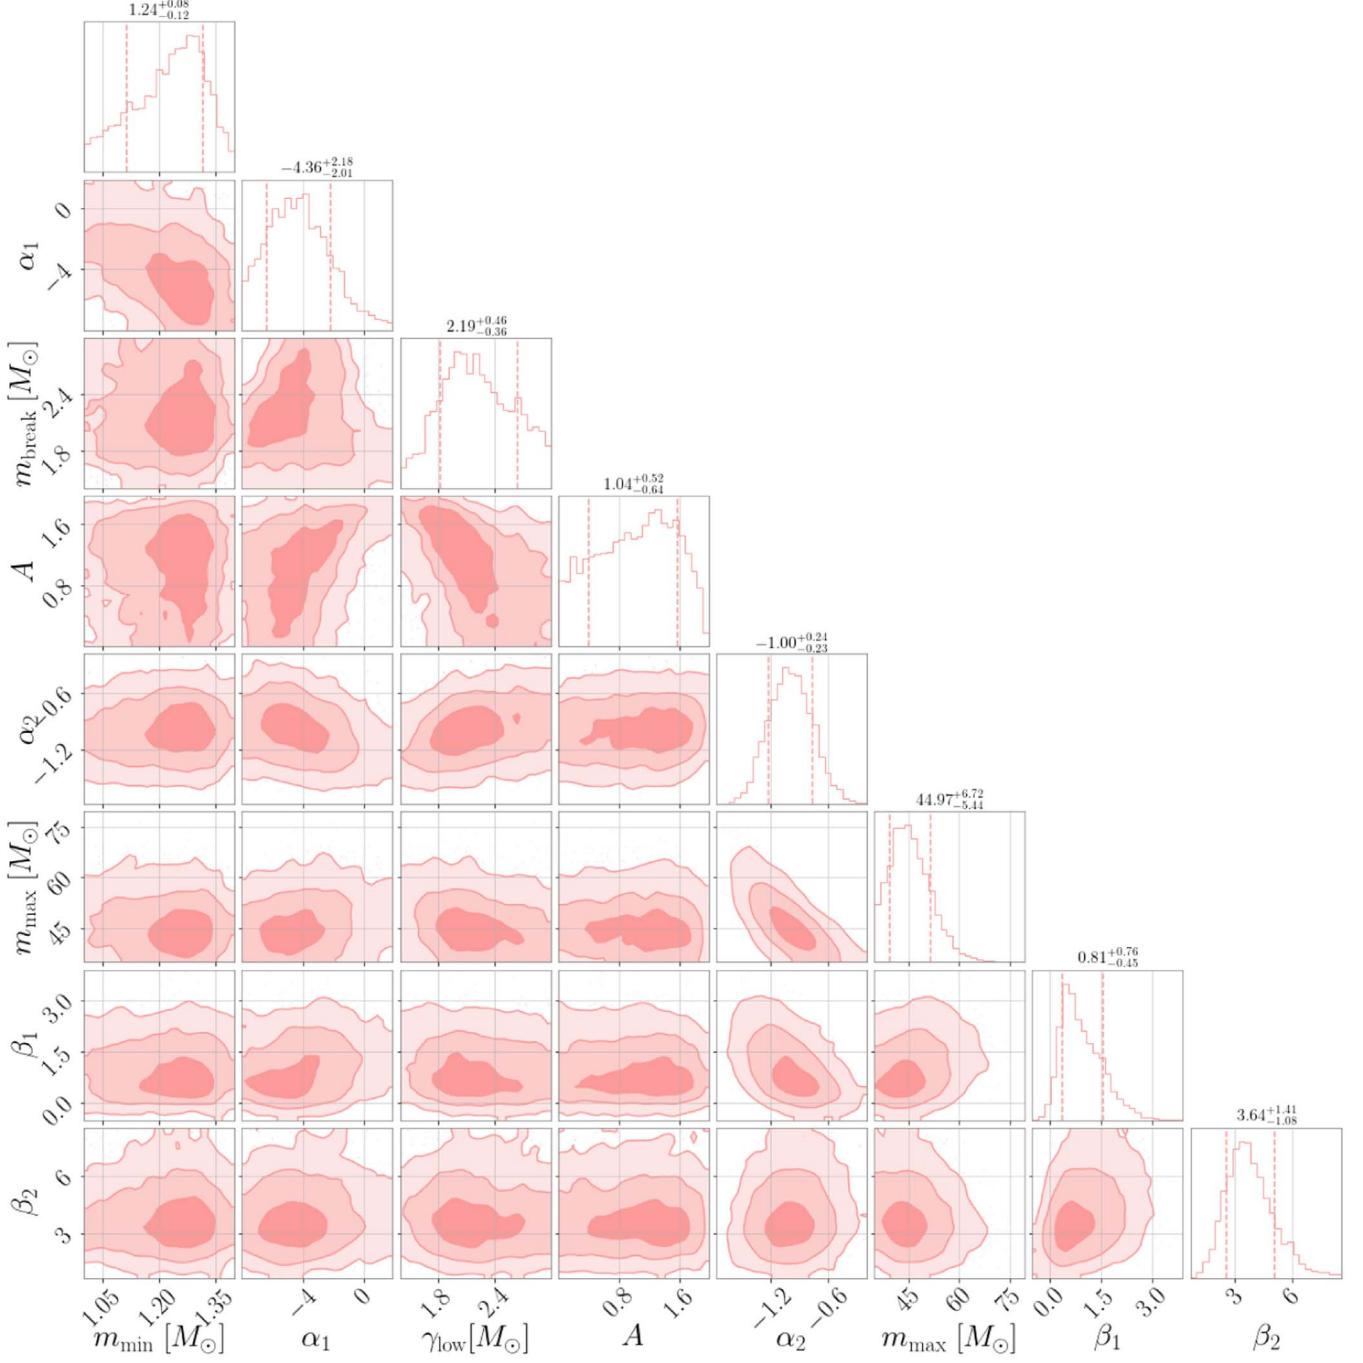


Figure 11. Same as Figure 10 but for BROKEN POWER LAW + DROP.

LAW + DIP, respectively. At the top of each column in these corner plots is listed the median and 68% credible intervals for each one-dimensional marginal distribution. Note that this is not the same as the set of overall median values, as there are significant correlations between the hyperparameters. We therefore provide the maximum hyperposterior value for each model in Table 4. These are the values used to produce the colorful line in each panel of Figure 3. We additionally provide Bayes factors (ratios of marginal likelihoods) and maximum-likelihood ratios for each model relative to BROKEN POWER LAW in Table 4. Three of the models have similar Bayes factors, \mathcal{B} , with a spread in \mathcal{B} of $\lesssim 20\%$. However, BROKEN POWER LAW + DIP is favored by a factor of ~ 10 . While not

conclusive (see rules of thumb in Jeffreys 1998 and Kass & Raftery 1995), this difference between Bayes factors suggests that a mass gap with sharp edges may be preferred. The ratio of maximum likelihoods for each model is provided to help disentangle the effect of the priors on the Bayes factors. The two models which allow for a lower mass gap seem to be preferred over those that do not, with maximum likelihoods that are larger by factors of $\sim 20\text{--}50$. Therefore, the low Bayes factor of the most flexible model may partially be due to its large prior volume, i.e., an “Occam penalty.” Nonetheless, the level of discrepancy seen in Table 4 is not large enough to rule out any of the models considered based on their predictive accuracy.

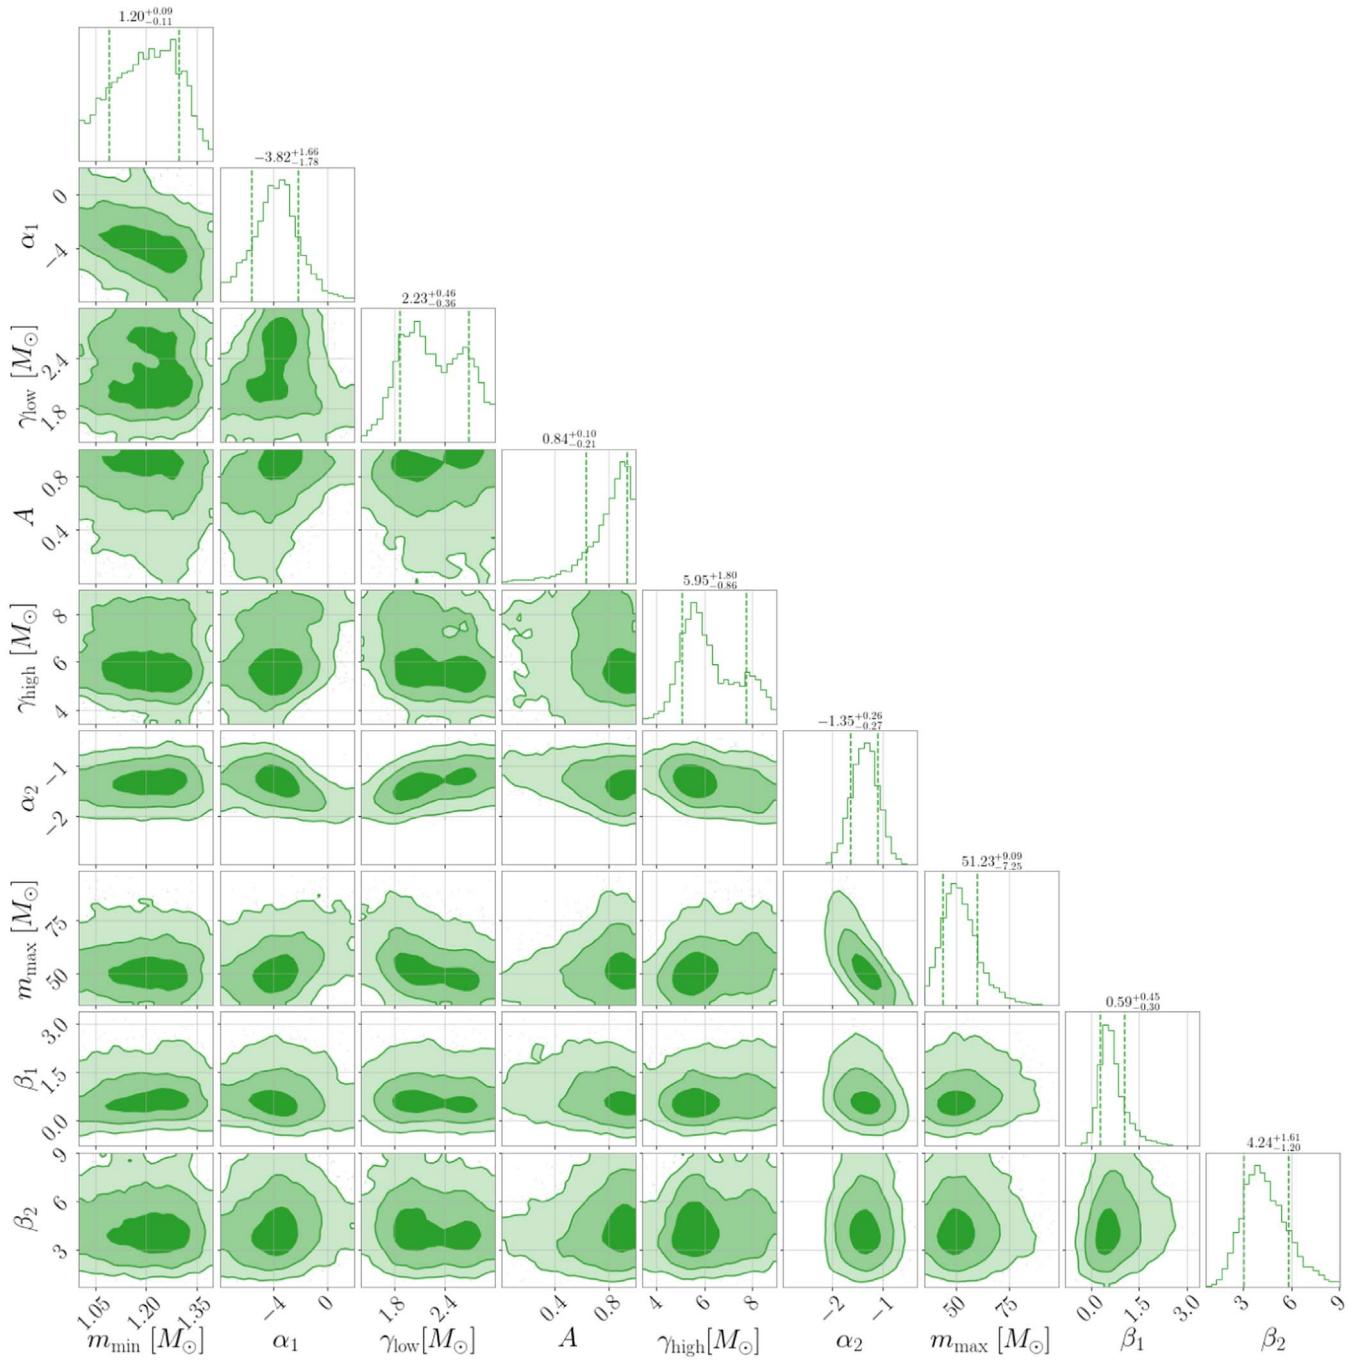


Figure 12. Same as Figure 10 but for BROKEN POWER LAW + DIP.

Table 4
Model Comparison

Parameter	Description	Value				
		MFM	BPL+Dip	BPL+Drop	BPL	
Broken power law	m_{break} Dividing mass between high-mass and low-mass power laws	2.32 [†] M_{\odot}	2.59 [†] M_{\odot}	2.07 [†] M_{\odot}	2.32 [†] M_{\odot}	
	α_1 Spectral index of power law below m_{break}	-4.25	-3.28	-6.79	-7.23	
	α_2 Spectral index of power law above m_{break}	-1.30	-1.15	0.65	-0.85	
	γ_{low} Roll-off mass for lower edge of the mass gap	2.32 [†] M_{\odot}	2.59 [†] M_{\odot}	2.07 [†] M_{\odot}	2.32 [†] M_{\odot}	
1D mass distribution	Notch filter	η_{low} Sharpness of lower edge of the mass gap	7.37	50 [*]	50 [*]	N/A
		γ_{high} Roll-off mass for the upper edge of the mass gap	5.87 M_{\odot}	5.15 M_{\odot}	N/A	N/A
		η_{high} Sharpness of the upper edge of the mass gap	9.14	50 [*]	N/A	N/A
		A Depth of the notch filter	0.99	0.99	1.44	0 [*]
High-pass filter	m_{min} Roll-off mass for high-pass filter at the lowest masses allowed	1.24 M_{\odot}	1.18 M_{\odot}	1.28 M_{\odot}	1.27 M_{\odot}	
		η_{min} Sharpness of high-pass filter at the lowest masses allowed	11.23	50 [*]	50 [*]	50 [*]
	Low-pass filter	m_{max} Roll-off mass for low-pass filter at the highest masses allowed	51.4 M_{\odot}	40.74 M_{\odot}	38.46 M_{\odot}	40.16 M_{\odot}
		η_{max} Sharpness of low-pass filter at the highest masses allowed	5.94	3.73	4.13	4.41
Pairing function	β_1 Spectral index of pairing function if $m_2 < 5 M_{\odot}$	0.18	0.41	0.57	0.54	
	β_2 Spectral index of pairing function if $5 M_{\odot} \leq m_2$	3.15	4.83	4.01	4.01	
Model selection	$\mathcal{B}_{\text{BPL}}^{\text{M}} \equiv p(\text{data} \text{M})/p(\text{data} \text{BPL})$	1.0	11.1	1.2	1.0	
	$\mathcal{L}_{\text{BPL}}^{\text{M}} \equiv \max_{\theta} \{p(\text{data} \theta, \text{M})\} / \max_{\theta} \{p(\text{data} \theta, \text{BPL})\}$	23.7	55.4	1.1	1.0	

Note. In the rightmost columns are hyperparameter values that maximize the overall hyperposterior for each of the three parameterizations. These may be different than the values that maximize the one-dimensional marginal posterior of each hyperparameter. Values labeled with a [†] denote those that are fixed to each other (i.e., $\gamma_{\text{low}} = m_{\text{break}}$ for all parameterizations used in this work), and values labeled with a ^{*} denote values that are fixed a priori and are therefore not fit as free parameters. In the second to last row we list Bayes factors relative to BROKEN POWER LAW. There is no clear preference between BROKEN POWER LAW, BROKEN POWER LAW + DROP, and the most flexible model (described in Appendix B and denoted “MFM” here). BROKEN POWER LAW + DIP is slightly favored, though this level of discrepancy is not definitive evidence against the other models’ ability to predict the data. In the final row are the maximum likelihoods of each model relative to BROKEN POWER LAW. The high maximum-likelihood ratios of the mass-gap-containing models further support the possibility that the data prefer a mass gap.

ORCID iDs

Amanda Farah <https://orcid.org/0000-0002-6121-0285>
 Maya Fishbach <https://orcid.org/0000-0002-1980-5293>
 Reed Essick <https://orcid.org/0000-0001-8196-9267>
 Daniel E. Holz <https://orcid.org/0000-0002-0175-5064>
 Shanika Galaudage <https://orcid.org/0000-0002-1819-0215>

References

Abbott, B., Abbott, R., Abbott, T., et al. 2017, *PhRvL*, **119**, 161101
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, *PhRvX*, **9**, 031040
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, *ApJ*, **875**, 161
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019c, *ApJL*, **882**, L24
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020a, *ApJL*, **892**, L3
 Abbott, R., Abbott, T. D., Abraham, S., et al. 2020b, *ApJL*, **896**, L44
 Abbott, R., Abbott, T. D., Abraham, S., et al. 2021a, *PhRvX*, **11**, 021053
 Abbott, R., Abbott, T. D., Abraham, S., et al. 2021b, *ApJL*, **915**, L5
 Abbott, R., Abbott, T. D., Abraham, S., et al. 2021c, *ApJL*, **913**, L7
 Abbott, R., Abbott, T. D., Acernese, F., et al. 2021d, arXiv:2108.01045
 Abbott, R., Abbott, T. D., Sheelu, A., et al. 2021e, *SoftX*, **13**, 100658
 Abbott, R., Abbott, T. D., Acernese, F., et al. 2021f, arXiv:2111.03606v2
 Alsing, J., Silva, H. O., & Berti, E. 2018, *MNRAS*, **478**, 1377
 Antoniadis, J., Tauris, T. M., Ozel, F., et al. 2016, arXiv:1605.01665
 Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, *ApJS*, **241**, 27
 Baily, C. D., Jain, R. K., Coppi, P., & Orosz, J. A. 1998, *ApJ*, **499**, 367
 Belczynski, K., Wiktorowicz, G., Fryer, C. L., Holz, D. E., & Kalogera, V. 2012, *ApJ*, **757**, 91
 Breivik, K., Chatterjee, S., & Andrews, J. J. 2019, *ApJL*, **878**, L4
 Broekgaarden, F. S., Berger, E., Neijssel, C. J., et al. 2021, *MNRAS*, **508**, 5028
 Chatterjee, D., Ghosh, S., Brady, P. R., et al. 2020, *ApJ*, **896**, 54
 Chatzioannou, K. 2020, *GR&Gr*, **52**, 109
 Dickey, J. M., & Lientz, B. P. 1970, *Ann. Math. Stat.*, **41**, 214
 Doctor, Z., Wysocki, D., O’Shaughnessy, R., Holz, D. E., & Farr, B. 2020, *ApJ*, **893**, 35
 Drozda, P., Belczynski, K., O’Shaughnessy, R., Bulik, T., & Fryer, C. L. 2020, arXiv:2009.06655
 Essick, R., Farah, A., Galaudage, S., et al. 2022, *ApJ*, **926**, 34
 Essick, R., & Landry, P. 2020, *ApJ*, **904**, 80
 Essick, R., Tews, I., Landry, P., Reddy, S., & Holz, D. E. 2020, *PhRvC*, **102**, 055803
 Farr, W. M., & Chatzioannou, K. 2020, *RNAAS*, **4**, 65
 Farr, W. M., Sravan, N., Cantrell, A., et al. 2011, *ApJ*, **741**, 103
 Fishbach, M., Doctor, Z., Callister, T., et al. 2021, *ApJ*, **912**, 98
 Fishbach, M., Essick, R., & Holz, D. E. 2020, *ApJL*, **899**, L8
 Fishbach, M., & Holz, D. E. 2020, *ApJL*, **891**, L27
 Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021, *ApJL*, **915**, L12
 Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, *ApJ*, **749**, 91
 Fryer, C. L., Heger, A., Langer, N., & Wellstein, S. 2002, *ApJ*, **578**, 335
 Fryer, C. L., & Kalogera, V. 2001, *ApJ*, **554**, 548
 Giesers, B., Dreizler, S., Husser, T.-O., et al. 2018, *MNRAS*, **475**, L15
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
 Hoyer, S., & Hamman, J. 2017, *JORS*, **5**, 10
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Jayasinghe, T., Stanek, K. Z., Thompson, T. A., et al. 2021, *MNRAS*, **504**, 2577
 Jeffreys, H. 1998, Theory of Probability (3rd ed.; Oxford: Oxford Univ. Press)
 Kass, R. E., & Raftery, A. E. 1995, *J. Am. Stat. Assoc.*, **90**, 773
 Kiziltan, B., Kottas, A., Yoreo, M. D., & Thorsett, S. E. 2013, *ApJ*, **778**, 66
 Kreidberg, L., Baily, C. D., Farr, W. M., & Kalogera, V. 2012, *ApJ*, **757**, 36

Kushnir, D. 2015, arXiv:[1502.03111](https://arxiv.org/abs/1502.03111)

Landry, P., & Read, J. S. 2021, *ApJL*, **921**, L25

Legred, I., Chatzioannou, K., Essick, R., Han, S., & Landry, P. 2021, *PhRvD*, **104**, 063003

Liu, T., Wei, Y.-F., Xue, L., & Sun, M.-Y. 2021, *ApJ*, **908**, 106

Loredo, T. J. 2004, in 24th Int. Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering. AIP Conf. Proc. 735 (College Park, MD: AIP), **195**

LVK 2015a, L1 Calibrated Sensitivity Spectra Oct 24 2015 (Representative for O1), <https://dcc.ligo.org/LIGO-G1600151/public>

LVK 2015b, H1 Calibrated Sensitivity Spectra Oct 24 2015 (Representative for O1), <https://dcc.ligo.org/LIGO-G1600150/public>

LVK 2017a, L1 Calibrated Sensitivity Spectra Aug 06 2017 (Representative Best of O2—C02, With Cleaning/Subtraction), <https://dcc.ligo.org/LIGO-G1801952/public>

LVK 2017b, H1 Calibrated Sensitivity Spectra Jun 10 2017 (Representative Best of O2—C02, With Cleaning/Subtraction), <https://dcc.ligo.org/LIGO-G1801950/public>

LVK 2020, aLIGO, CAL, Official Advanced LIGO Sensitivity Plots, <https://dcc.ligo.org/LIGO-G1500623/public>

Mandel, I., Farr, W. M., & Gair, J. R. 2019, *MNRAS*, **486**, 1086

Mandel, I., Müller, B., Riley, J., et al. 2021, *MNRAS*, **500**, 1380

Margalit, B., & Metzger, B. D. 2017, *ApJL*, **850**, L19

Mroz, P., & Wyrzykowski, L. 2021, *AcA*, **71**, 89

Müller, B., Heger, A., Liptai, D., & Cameron, J. B. 2016, *MNRAS*, **460**, 742

Özel, F., Psaltis, D., Narayan, R., & McClintock, J. E. 2010, *ApJ*, **725**, 1918

Ruiz, M., Shapiro, S. L., & Tsokaros, A. 2018, *PhRvD*, **97**, 021501

Shao, Y., & Li, X.-D. 2021, *ApJ*, **920**, 81

Shibata, M., Zhou, E., Kiuchi, K., & Fujibayashi, S. 2019, *PhRvD*, **100**, 023015

Speagle, J. S. 2020, *MNRAS*, **493**, 3132

Talbot, C., Smith, R., Thrane, E., & Poole, G. B. 2019, *PhRvD*, **100**, 043030

Thompson, T. A., Kochanek, C. S., Stanek, K. Z., et al. 2019, *Sci*, **366**, 637

Thrane, E., & Talbot, C. 2019, *PASA*, **36**, e010

Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., & Grasman, R. 2010, *Cog. Psychol.*, **60**, 158

Wyrzykowski, u., & Mandel, I. 2020, *A&A*, **636**, A20

Zevin, M., Spera, M., Berry, C. P. L., & Kalogera, V. 2020, *ApJL*, **899**, L1