The Astrophysical Journal, 962:69 (17pp), 2024 February 10 https://doi.org/10.3847/1538-4357/ad0558

Two of a Kind: Comparing Big and Small Black Holes in Binaries with Gravitational Waves

Amanda M. Farah[1], Maya Fishbach[2], and Daniel E. Holz[1,3,4]

[1] Department of Physics, The University of Chicago, Chicago, IL 60637, USA; afarah@uchicago.edu

[2] Canadian Institute for Theoretical Astrophysics, David A. Dunlap Department of Astronomy and Astrophysics, and Department of Physics, 60 Saint George Street, University of Toronto, Toronto, ON M5S 3H8, Canada

[3] Kavli Institute for Cosmological Physics, The University of Chicago, 5640 South Ellis Avenue, Chicago, IL 60637, USA

[4] Enrico Fermi Institute, The University of Chicago, 933 East 56th Street, Chicago, IL 60637, USA

Received 2023 August 9; revised 2023 September 14; accepted 2023 October 16; published 2024 February 7

Abstract

When modeling the population of merging binary black holes, analyses have generally focused on characterizing the distribution of primary (i.e., more massive) black holes in the binary, while using simplistic prescriptions for the distribution of secondary masses. However, the secondary mass distribution and its relationship to the primary mass distribution provide a fundamental observational constraint on the formation history of coalescing binary black holes. If both black holes experience similar stellar evolutionary processes prior to collapse, as might be expected in dynamical formation channels, the primary and secondary mass distributions would show similar features. If they follow distinct evolutionary pathways (for example, due to binary interactions that break symmetry between the initially more massive and less massive stars), their mass distributions may differ. We present the first analysis of the binary black hole population that explicitly fits for the secondary mass distribution. We find that the data is consistent with a $\sim 30\,M_\odot$ peak existing only in the distribution of secondary rather than primary masses. This would have major implications for our understanding of the formation of these binaries. Alternatively, the data is consistent with the peak existing in both component mass distributions, a possibility not included in most previous studies. In either case, the peak is observed at $31.4^{+2.3}_{-2.6}\,M_\odot$, which is shifted lower than the value obtained in previous analyses of the marginal primary mass distribution, placing this feature in further tension with expectations from a pulsational pair-instability supernova pileup.

Unified Astronomy Thesaurus concepts: Gravitational wave sources (677); Gravitational waves (678); Gravitational wave astronomy (675); Binary stars (154); Globular star clusters (656); Common envelope evolution (2154); Hierarchical models (1925); Bayesian statistics (1900); A supergiant stars (8); Supernova remnants (1667)

1. Introduction

Dozens of gravitational-wave (GW) events have been observed by the LIGO (Aasi et al. 2015), Virgo (Acernese et al. 2014), and KAGRA (Akutsu et al. 2021) detector network (Abbott et al. 2023c), and many more detections are anticipated by the end of the fourth observing run. However, the formation mechanism of binary black hole (BBH) mergers, which are the source of the majority of detected GW events, is still largely uncertain. While it is not possible to know the formation history of any one BBH with certainty, the population of all merging BBHs encodes information about which astrophysical processes give rise to the bulk of these systems (e.g., Stevenson et al. 2015; Zevin et al. 2017).

The stellar-mass BBHs detectable by LIGO, Virgo, and KAGRA are likely created from the collapse of massive stars. These massive stars may be born as a binary system in the galactic field, which then evolves into a BBH system. Alternatively, they may be born in a dense stellar environment, such as a star cluster, in which the black hole (BH) stellar remnants dynamically assemble into tightly bound binaries. The population of GW sources contains signatures of the initial conditions of their progenitor stellar systems, as well as many of the evolutionary processes that occur between star formation and BBH merger. The initial conditions that impact the GW source population include the initial mass function (IMF) of either binary or single stars in different environments, the birth metallicities of the progenitor stars, and other aspects of their formation environments. Depending on the BBH formation scenario, the evolutionary processes that are imprinted on the BBH population include stellar mass loss, transfer of matter between the two component stars, the supernova process for massive stars, and dynamical interactions in star clusters or the disks of active galactic nuclei (see reviews by Mapelli 2020; Mandel & Farmer 2022, and references therein).

In general, these uncertain formation processes affect the masses, spins, redshifts, eccentricities, and merger rates of BBH systems. Here we focus on the mass distribution, since BH masses are well measured from the GW signal.

The BBH mass distribution is typically parameterized by the primary mass m_1, the larger of the two component masses in the binary, and either the secondary mass m_2, or the mass ratio $q = m_2/m_1$. Within these parameterizations, it is typically assumed either that the primary and secondary masses follow the same underlying distribution (Doctor et al. 2020; Fishbach & Holz 2020; Farah et al. 2022; Abbott et al. 2023a; Edelman et al. 2023; Sadiq et al. 2023), or that the primary mass distribution has distinct features from those of the secondary mass distribution. The latter assumption implies that primary and secondary mass are physically meaningful labels, and is typically achieved by modeling the secondary mass distribution as a single power law between some minimum mass and m_1 (Fishbach & Holz 2017;

Kovetz et al. 2017; Talbot & Thrane 2018; Abbott et al. 2021a, 2023a; Baxter et al. 2021; Tiwari 2021; Edelman et al. 2022a; Callister & Farr 2023; Godfrey et al. 2023). The former assumption—that m_1 and m_2 follow the same underlying distribution—implies that the two-dimensional mass distribution is symmetric in primary and secondary mass, meaning that m_1 and m_2 can be interchanged without changing the mass distribution. Neither of these assumptions have been explicitly verified with the available data. In this work, we aim to determine if both components in merging BBHs follow the same underlying distribution or if there is a physical distinction between primary and secondary masses.

Understanding whether the primary and secondary masses in BBH mergers follow the same distribution will provide insight into their formation histories. For BBH mergers that are dynamically assembled in dense environments, we expect that both component BHs are drawn from the same population of stellar remnants (i.e., m_1 and m_2 are not physically meaningful labels), and that the two-dimensional BBH mass distribution will therefore be symmetric in m_1 and m_2. For BBH mergers that originate from binary stars that formed and evolved in relative isolation (“field binaries”), there may be a physical distinction between primary and secondary BH masses. To start off, the progenitor stars in the binary are drawn from the binary IMF, which may not be symmetric between the two components (e.g., Grudić et al. 2023). Furthermore, the two stars exchange mass during binary stellar evolution. In each phase of mass transfer, one component acts as the donor and the other as the accretor, depending on their initial masses. Mass transfer affects the donor and accretor stars in different ways, which can impact the masses of the resulting BHs following stellar collapse (Laplace et al. 2021). On the other hand, both BH progenitors are expected to have undergone binary stripping and therefore experience similar supernova physics, potentially washing out any significant differences between the mass distributions of first- and secondborn BHs (van Son et al. 2022b; Schneider et al. 2023). However, the degree to which these distributions are similar or different depends on uncertain mass loss and accretion physics (van Son et al. 2022a).

Supernova kicks may also be different for firstborn versus secondborn components in *merging* binaries, because kicks determine whether or not the binary can merge within the age of the Universe (Kalogera 1996; Gallegos-Garcia et al. 2022). Because natal kicks are related to the remnant mass via the supernova prescription (Fryer et al. 2012; Mandel et al. 2021), different preferences for the natal kick magnitudes between firstborn and secondborn BHs may cause the primary and secondary mass distributions to differ in merging binaries formed through isolated binary evolution (e.g., Oh et al. 2023).

In short, binary stellar evolution consists of several processes that can break the symmetry between the population of initially more massive stars, which generally correspond to the firstborn and more massive (primary) BHs, and the population of initially less massive stars, which generally correspond to the secondborn and less massive (secondary) BHs. However, if mass inversion occurs in some systems, some initially less massive stars will end up as the more massive BHs by the time of merger, and the distribution of primary BH masses will have contributions from both the secondborn and firstborn BHs (Olejak & Belczynski 2021; Broekgaarden et al. 2022; Hu et al. 2022; Zevin & Bavera 2022). If mass inversion happens in exactly half of merging BBH systems, the primary and secondary component mass distributions may be indistinguishable even if the first- and secondborn distributions differ.

From a data analysis perspective, knowing that the primary and secondary mass distributions are the same allows us to measure a single set of model parameters: those of the shared distribution. This may allow features of the distribution to be measured with higher precision since both components in the binary will contribute to the measurement of each feature, rather than just the primary mass. Furthermore, disentangling the role of primary and secondary masses aids the physical interpretation of such features.

For example, Abbott et al. (2023a) found that the mass distribution exhibits a peak at primary masses $m_1 \sim 35\,M_{\odot}$. The astrophysical origin of this overdensity in the mass distribution is uncertain, although it may be related to (pulsational) pair-instability supernovae (e.g., Heger & Woosley 2002; Fishbach & Holz 2017; Talbot & Thrane 2018; Farmer et al. 2019). A necessary ingredient toward understanding the origin of the $m_1 \sim 35\,M_{\odot}$ peak is to first understand whether the peak appears exclusively among primary BBH masses, or whether secondary masses also display a peak, indicating that secondary BHs also experience the astrophysical process that leads to a mass pileup.

This paper is organized as follows. In Section 2 we describe the different population models considered in this work. In Section 3 we present the results of fitting each of the models to the Third Gravitational-wave Transient Catalog (GWTC-3; Abbott et al. 2023c), finding that while the primary and secondary masses of merging BBHs are consistent with following the same underlying distribution, it is also possible that the secondary mass distribution has a more prominent peak than does the primary mass distribution. In Section 4 we discuss possibilities for future observations. In Section 5 we summarize our conclusions and present possible astrophysical interpretations of our results. The appendices discuss the fundamental differences between commonly used parameterizations for the two-dimensional mass distribution and include details about the population models and statistical framework.

2. Population Models

Because our aim is to learn whether the primary and secondary masses are consistent with being drawn from the same distribution, we describe the primary and secondary mass distributions separately, but according to the same functional form. We model each of the one-dimensional mass distributions as a mixture model between a smoothed power-law component and a Gaussian component in order to make direct comparisons to the POWER LAW + PEAK model used by the LIGO–Virgo–KAGRA collaboration (LVK) to describe the distribution of primary masses (Talbot & Thrane 2018; Abbott et al. 2021a, 2023a).

There are several ways to construct a two-dimensional mass model for the binary from this one-dimensional mass model for the component masses. However, as we discuss in Appendix A, in order to explicitly compare primary and secondary mass distributions, it is necessary to use the pairing function framework first introduced for GW population modeling in Fishbach & Holz (2020). Explicitly,

$$p(m_1, m_2|\Lambda) = p_1(m_1|\Lambda_1)p_2(m_2|\Lambda_2)f(q;\, \beta_q), \tag{1}$$

where $p_1(m_1|\Lambda_1)$ is the underlying distribution of primary masses, $p_2(m_2|\Lambda_2)$ is the underlying distribution of secondary masses, $f(q)$ is a pairing function that depends on the mass ratio of the system[5], Λ_1 and Λ_2 are hyperparameters[6] describing the underlying primary and secondary mass distributions, respectively, and $\Lambda = \{\Lambda_1, \Lambda_2, \beta_q\}$ is the set of all mass model hyperparameters. In this work, we use a pairing function of the form $f(q; \beta_q) = q^{\beta_q}\Theta(q \leqslant 1)$, though other forms may provide a better fit to the data (e.g., Farah et al. 2022).

As illustrated in Figure 1, different choices for the relationship between Λ_1 and Λ_2 result in distinct morphologies for the two-dimensional mass distribution. Below, we list each of the variations we consider in this work, along with the panels in Figure 1 to which it corresponds. A table describing these variations in terms of choices for the hyperprior is given in Appendix B.3.

1. PAIRING:SYMMETRIC. This model sets $\Lambda_1 = \Lambda_2$, making the distribution symmetric under the transformation $m_1 \leftrightarrow m_2$. It corresponds to the first row of Figure 1: any feature in one of the distributions has to be present in both, so it always makes two bands that connect on the diagonal.
2. PAIRING:GENERIC. This model allows the hyperparameters describing the $\sim 35\,M_\odot$ peak to differ between $p_1(m_1|\Lambda_1)$ and $p_2(m_2|\Lambda_2)$. It can produce any of the panels in Figure 1, and is the only model that can produce a scenario such as that illustrated in the middle panels, where the feature in m_1 has a different amplitude and location from those of the feature in m_2.
3. PAIRING:NO PEAK IN $p_2(m_2)$. This model sets all $\Lambda_2 = \Lambda_1$, except for the parameter governing the height of the $\sim 35\,M_\odot$ peak, which we set to vanish for the secondary mass distribution but fit as a free parameter for the primary mass distribution. This model corresponds to the bottom row of Figure 1: it is only capable of having a peak in the primary mass distribution, so can only produce vertical bands in a two-dimensional mass distribution. In Appendix A.3, we show that PAIRING:NO PEAK IN $p_2(m_2)$ approximately mimics the behavior of the commonly used POWER LAW + PEAK model from, e.g., Abbott et al. (2023a), which we refer to as LVK 2023.

The different columns in Figure 1 correspond to different power-law spectral indices, β_q, for the mass-ratio-dependent pairing function. The leftmost panels show models where components in the binary are allowed to pair up independently of the mass ratio, the middle column shows a model where components have a slight preference to pair up with partners that are of equal mass, and the rightmost panels show the case where components are very "picky": they almost always pair up with equal-mass partners (Fishbach & Holz 2020). When components pair up nearly independently of the mass ratio, the asymmetric models produce noticeably different distributions from those produced by the symmetric models. However, in the case of very picky binaries, the two scenarios become difficult to tell apart. There is therefore a degeneracy between the steepness of the pairing function and the existence of distinct features in the two mass distributions (see Tiwari 2023 for a discussion of this phenomenon in terms of Jacobian transformations).

In all models considered in this work, the redshift distribution is modeled as a power law with spectral index κ (Fishbach et al. 2018). We use the DEFAULT spin model from Abbott et al. (2021a, 2023a) to describe the spin magnitudes and tilts of each component. These are the same redshift and spin distributions used with LVK 2023 in Abbott et al. (2023a). The explicit form of the full population model is given in Appendix B.

Using these parameterized models, we construct a hierarchical Bayesian inference (described in Appendix C) to determine the appropriate population-level parameters for the mass distribution, Λ, given the observed set of data $\{D_j\}$ for N observed events (Loredo 2004; Mandel et al. 2019). We model the data as an inhomogeneous Poisson process with the rate density (the number of events per unit time per single-event-parameter hypervolume) given by

$$\frac{d\mathcal{N}}{dm_1 dm_2,\, ds_1 ds_2 dz} = \mathcal{R}p(z)p(s_1, s_2)p(m_1, m_2|\Lambda), \tag{2}$$

where $\mathcal{R}$ acts as a normalizing constant that sets the overall magnitude of the rate.

3. Results

We fit each model described in Section 2 to the BBHs in GWTC-3. The resulting two-dimensional mass distributions for the PAIRING:SYMMETRIC and PAIRING:GENERIC models are shown in Figure 2. These plots represent an average over the hyperposterior for each model; this average is sometimes referred to as a posterior population distribution (PPD). The contours differ in morphology from those illustrated in Figure 1 because the actual distribution of BBH component masses exhibits two peaks, one at $\sim 10\,M_\odot$ and another at $\sim 35\,M_\odot$ (Abbott et al. 2021a; Tiwari 2021; Edelman et al. 2022a, 2023; Sadiq et al. 2022; Abbott et al. 2023a; Farah et al. 2023; Ray et al. 2023), whereas we only place one peak in the models shown in Figure 1. The peak at $\sim 10\,M_\odot$ creates bands in all panels that have very little vertical extent because the peak is proximal to the minimum BH mass. Nonetheless, fits using both models exhibit vertical and horizontal bands, indicating peaks in both $p_1(m_1)$ and $p_2(m_2)$. While the PAIRING:SYMMETRIC model requires this behavior, the PAIRING:GENERIC model does not, meaning that the secondary mass distribution appears to exhibit its own feature at $\sim 35\,M_\odot$. Furthermore, the fact that PAIRING:GENERIC produces a PPD similar to that of PAIRING:SYMMETRIC indicates that the data support consistent primary and secondary mass distributions.

The bands in all panels do not go to the full extent of the parameter space but rather taper off from the diagonal, indicating a preference for equal-mass binaries either through a pairing function or through a mass ratio distribution that favors $m_1 \approx m_2$.

3.1. Primary and Secondary Masses Are Consistent with Having the Same Underlying Distribution

The underlying distributions (i.e., before a pairing function is applied) of the primary and secondary masses are shown in Figure 3 for the two pairing function models. There is a region of overlap between the primary and secondary mass distributions under the PAIRING:GENERIC model, indicating that the primary and secondary mass distributions are consistent with each other. As expected, this region of overlap also coincides with $p(m)$, the distribution describing both component masses in the PAIRING:SYMMETRIC model.

---

[5] In principle, the pairing function can be parameterized in terms of any observable parameter (e.g., total mass).

[6] We use the term "hyperparameter" for a model parameter that describes the population of merging compact binaries, in contrast to parameters describing each individual GW detection, such as one system's masses and spins.

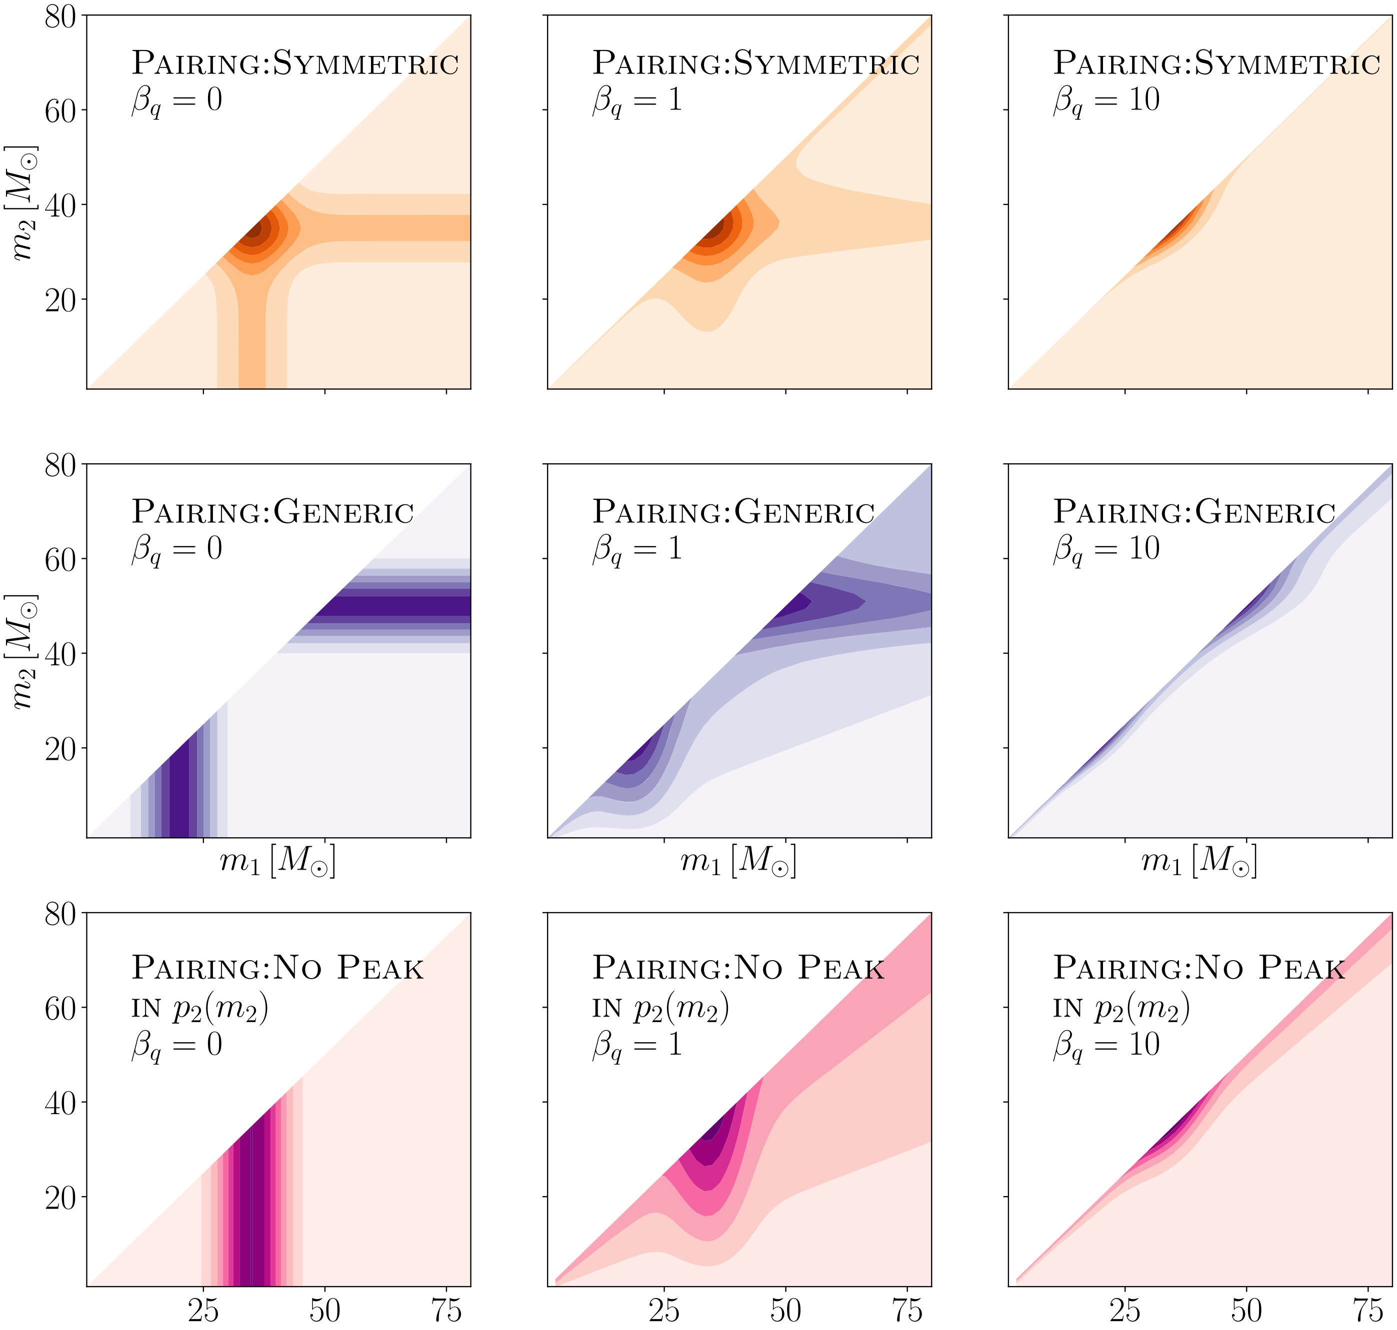

Figure 1. Illustration of some possible two-dimensional mass distributions under the models considered in this work. These distributions are not indicative of specific astrophysical predictions, but are instead meant to be illustrative of the morphologies accessible by current models. Overdensities/peaks in the mass distribution appear as darker filled contours in these figures. The different columns correspond to different power-law spectral indices for the mass-ratio-dependent pairing function, β_q. In the case where components strongly prefer to pair with nearly equal-mass partners, it becomes difficult to determine whether a feature appears only in one component mass distribution (as in the PAIRING:NO PEAK IN $p_2(m_2)$ model, bottom row) or in both (PAIRING:SYMMETRIC, top row). The diagonal contours in the middle and right columns are caused by a preference for equal-mass binaries and follow lines of constant mass ratio. The goal of this paper is to distinguish between the different scenarios illustrated in this figure.

Models not explicitly parameterized in terms of a pairing function are unable to produce underlying distributions such as those shown in Figure 3, so for the sake of comparison to previous work, we turn to conditional[7] m_2 distributions, $p(m_2|m_1 = C)$, where C can be any number in the domain of the m_1 distribution. Figure 4 shows these conditional distributions for the PAIRING:SYMMETRIC, PAIRING:GENERIC, and LVK 2023 models, the latter of which does not use a pairing function. The curves in Figure 4 are averaged over the hyperposterior for each model.

For all values of C, the inferred distribution under the PAIRING:GENERIC model behaves similarly to that of the PAIRING:SYMMETRIC model, indicating consistency between the primary and secondary mass distributions. In particular, under the PAIRING:SYMMETRIC and PAIRING:GENERIC models,

---

[7] While marginal distributions are typically used for the purpose of comparing PPDs from several models on a single plot, they are not sensitive to differences between two-dimensional mass distributions when equal masses are preferred (see Appendix B.3), so we use conditional distributions instead.

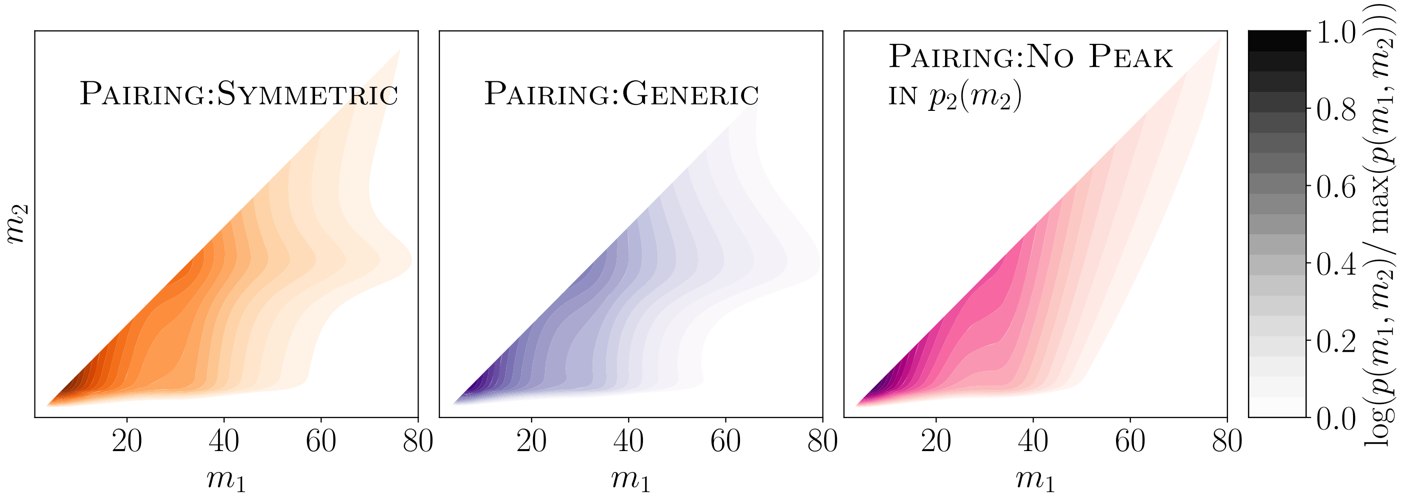

Figure 2. Two-dimensional PPDs for PAIRING:SYMMETRIC (left, orange), PAIRING:GENERIC (middle, purple), and PAIRING:NO PEAK IN $p_2(m_2)$ (right, pink), all fit to the BBHs in GWTC-3. Darker colors indicate a higher rate of sources, and each panel is individually normalized to its maximum value. All three models find a higher rate of events near the $m_1 = m_2$ diagonal, as well as peaks at $m_1 \sim 10\,M_\odot$ and $m_1 \sim 35\,M_\odot$. PAIRING:SYMMETRIC and PAIRING:GENERIC both find peaks in m_2 as well, as indicated by the horizontal bands in the two leftmost panels, whereas PAIRING:NO PEAK IN $p_2(m_2)$ is unable to model features in the m_2 direction. Additionally, PAIRING:SYMMETRIC and PAIRING:GENERIC display peaks in similar locations despite the fact that PAIRING:SYMMETRIC requires that both m_1 and m_2 follow the same underlying distribution but PAIRING:GENERIC is able to model each component separately. This suggests that both components may be drawn from the same underlying distribution, up to a pairing function.

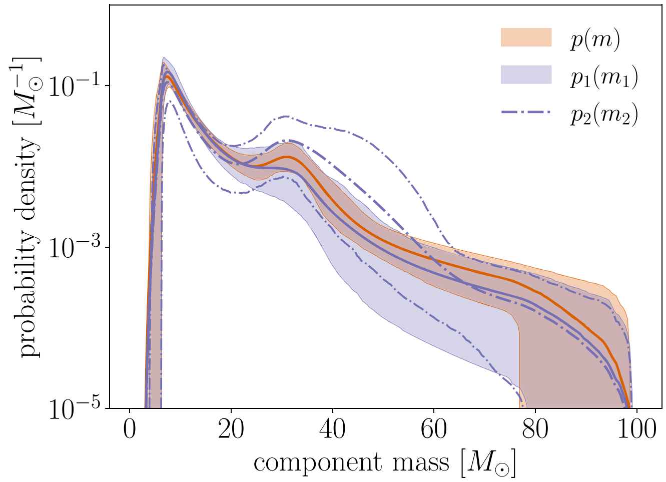

Figure 3. Underlying distribution of primary and secondary masses for the PAIRING:SYMMETRIC (orange) and PAIRING:GENERIC (purple) models. Under the PAIRING:SYMMETRIC model, $p_1(m_1) = p_2(m_2) \equiv p(m)$, so only $p(m)$ is shown. The $p(m)$ under PAIRING:SYMMETRIC is more tightly constrained than $p_1(m_1)$ (purple filled band) or $p_2(m_2)$ (dotted–dashed lines) under PAIRING:GENERIC as it has twice as many observations per free parameter. These distributions are constructed to describe the population of BHs *before* the function that pairs components into merging binaries is applied. All three underlying distributions are consistent with one another, though $p_2(m_2)$ appears to have a large peak at $\sim 35\,M_\odot$ while $p_1(m_1)$ has some support for no peak in that region; $p(m)$ does not have support for no peak, but its peak is constrained to be small, while the peak in $p_2(m_2)$ is less well constrained and may be larger in amplitude. This hints at the possibility that the peak identified in the primary mass distribution by the LVK 2023 formalism may have been driven by a peak in $p_2(m_2)$ rather than a peak in $p_1(m_1)$, though hyperparameter uncertainties within PAIRING:GENERIC are too large for us to definitively determine this, and the data is consistent with $p_1(m_1) = p_2(m_2)$.

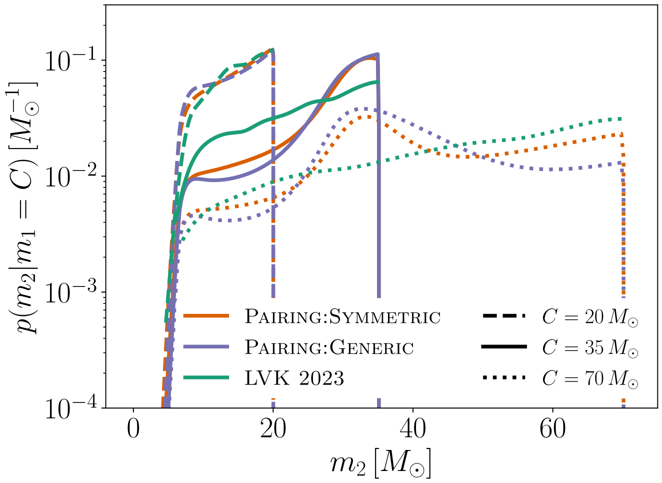

Figure 4. Conditional PPDs for the models considered in this work. We show the secondary mass distributions conditioned on $m_1 = 20\,M_\odot$ (dashed), $m_1 = 35\,M_\odot$ (solid), and $m_1 = 70\,M_\odot$ (dotted) to exemplify when the primary mass is below, inside, and above the Gaussian peak, respectively. Orange lines correspond to PAIRING:SYMMETRIC, purple lines to PAIRING:GENERIC, and green lines to LVK 2023. Lines denote the mean posterior probability, marginalized over the hyperparameter uncertainty, and credible intervals are omitted for clarity. When m_1 is below the peak, all models agree. When m_1 is above or within the peak, PAIRING:SYMMETRIC and PAIRING:GENERIC exhibit a stronger preference for m_2 to be in the peak than the LVK 2023 model does because the latter is constructed to behave like a power law in the range $[m_{\min}, m_1]$. The fact that PAIRING:GENERIC and PAIRING:SYMMETRIC approximately agree on the peak location and height indicates that the primary and secondary masses may follow the same underlying distribution up to a pairing function.

we see a peak in the conditional secondary mass distribution $p(m_2|m_1 = C)$ when $C \geqslant 35\,M_\odot$ (solid and dotted lines). While m_2 generally prefers to be near m_1, as indicated by the upward trend in all models, the two pairing function models have more support for m_2 being within the peak than for m_2 being nearly equal to m_1 when m_1 is large (dotted lines). This behavior is in contrast to that of the LVK 2023 model, which forces the conditional secondary mass distribution $p(m_2|m_1 = C)$ to monotonically increase for all values of C because it does not allow for a peak in m_2. While the PAIRING:GENERIC model can replicate this behavior, it instead recovers similar distributions to those of the PAIRING:SYMMETRIC model. The Bayes factor of PAIRING:GENERIC relative to PAIRING:SYMMETRIC is $\mathcal{B}(\text{PAIRING:GENERIC})/\mathcal{B}(\text{PAIRING:SYMMETRIC}) = 0.38$, indicating an inability

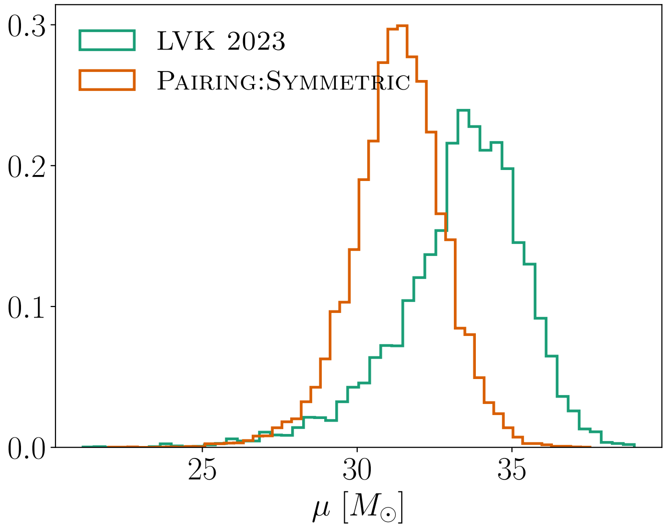

Figure 5. Marginal posteriors on μ, the location of the Gaussian peak, for the LVK 2023 and PAIRING:SYMMETRIC models. This parameter is more precisely measured under the PAIRING:SYMMETRIC formalism by $\sim 1/\sqrt{2}$, indicating that m_1 and m_2 independently contribute to this measurement. This parameter μ is the most correlated parameter with the Hubble constant when using "spectral sirens" (Ezquiaga & Holz 2022; Abbott et al. 2023b), so using the PAIRING:SYMMETRIC model will presumably improve cosmological measurements. The central values of the two distributions differ because the hyperparameter has slightly different effects on the resulting two-dimensional mass distribution of each model.

to distinguish between the two models, with perhaps a modest preference for the PAIRING:SYMMETRIC model.

3.2. Improved Constraints on Peak Location

If the primary and secondary mass distributions are identical, we can better measure the properties of features in that common distribution by using measurements of both m_1 and m_2. Figure 5 shows the one-dimensional hyperposterior for the location, μ, of the Gaussian peak in the mass distribution under the PAIRING:SYMMETRIC and LVK 2023 models. We measure μ with a standard deviation of $1.50\,M_\odot$ under the PAIRING:SYMMETRIC model compared to $2.08\,M_\odot$ under the LVK 2023 model, representing a 28% improvement. This increase in precision is similar to that expected from using twice the number of independent events ($1 - 1/\sqrt{2} = 0.29$), because now both m_1 and m_2 contribute to the inference, rather than just m_1.

Furthermore, the PAIRING:SYMMETRIC model recovers a lower value of $\mu = 31.4^{+2.3}_{-2.6}\,M_\odot$ (median and 90% credible interval) compared to the LVK 2023 result of $\mu = 33.6^{+2.6}_{-4.0}$, because it refers to a feature in the underlying $p(m)$ distribution rather than one in the marginal m_1 distribution. Features in $p(m)$ appear at larger masses in the marginal m_1 distribution and at lower masses in the marginal m_2 distribution due to the constraint that $m_1 > m_2$.

If the $\sim 35\,M_\odot$ peak is believed to be a feature of the supernova remnant mass distribution, it is best to use its location in the underlying $p(m)$ distribution rather than its location in the marginal distribution. For example, analyses wishing to compare this feature with expectations from a pair-instability supernova pileup (e.g., Stevenson et al. 2019; Farmer et al. 2020) should use the value presented here ($\mu = 31.4^{+2.3}_{-2.6}\,M_\odot$). Interestingly, this lower peak location is in further tension with the latest theoretical predictions for a pair-instability pileup (Farag et al. 2022).

When possible, though, it is best to compare theoretical predictions directly to the two-dimensional mass distribution (such as those in Figure 2) rather than to the values of specific hyperparameters, since hyperparameters have different meanings in different models. Correspondingly, it is important that models used to fit the data are intentionally constructed to allow for features in both the primary and secondary mass distribution. Many of the current parametric (Fishbach & Holz 2017; Talbot & Thrane 2018; Abbott et al. 2021a, 2023a; Baxter et al. 2021) and nonparametric (Tiwari 2021; Edelman et al. 2022a; Callister & Farr 2023; Godfrey et al. 2023) BBH mass distribution models enforce asymmetry between m_1 and m_2, excluding the possibility that the primary and secondary mass distributions share the same features[8] (with a few exceptions, e.g., Edelman et al. 2023; Ray et al. 2023; Sadiq et al. 2023). It is possible to construct a mass distribution that allows for features in both the primary and secondary mass distribution without using a pairing function (e.g., Ray et al. 2023), but directly comparing the primary and secondary mass distributions using these parameterizations is less straightforward.

Another application of our improved measurement of the peak location is "spectral siren" cosmology, which uses such features in the mass distribution to infer the expansion history of the Universe (Chernoff & Finn 1993; Messenger & Read 2012; Taylor et al. 2012; Farr et al. 2019; Ezquiaga & Holz 2021, 2022; Abbott et al. 2023b). Previous analyses have found the location of this peak to be the parameter most correlated with the Hubble constant (see Figures 5 and 13 of Abbott et al. 2023b), so an improved constraint on μ should have a relatively large effect on the constraints of cosmological parameters if all other mass distribution parameters remain equally well constrained.

3.3. Evidence for a Peak in the Secondary Mass Distribution

Consistency between the primary and secondary mass distributions implies a peak in the secondary mass distribution at $m_2 \sim 30\,M_\odot$, as this feature has already been robustly identified in the primary mass distribution (Abbott et al. 2021a, 2023a; Tiwari 2021; Callister & Farr 2023; Edelman et al. 2023; Farah et al. 2023). However, the data is also consistent with differing primary and secondary mass distributions, as shown by the regions in which the filled and dashed bands do not overlap in Figure 3. In this case, it is worthwhile to explicitly determine whether there exists a peak in the secondary mass distribution.

The left panel of Figure 6 shows the posterior distribution for the parameters governing the height of the peak in m_1 and m_2 under the PAIRING:GENERIC model, λ_1 and λ_2, respectively. Setting $\lambda_{\{1,2\}} = 0$ means that no binaries in the astrophysical population are in the Gaussian peak[9], while $\lambda_{\{1,2\}} = 1$ means all binaries are in the Gaussian peak. The lower left and upper right regions of the two-dimensional posterior are excluded, meaning that λ_1 and λ_2 cannot both be zero, nor can they both

---

[8] Callister & Farr (2023) and Godfrey et al. (2023) use flexible models to fit $p(q)$ rather than using power laws, but they do not use a pairing function and instead fit the m_1 and q distributions independently, therefore assuming the product of the marginal distributions equals the two-dimensional distribution. This is the same fundamental choice as is described in Equation (A1), and therefore assumes that m_1 and m_2 follow different distributions.

[9] The integrated fraction of events in the region near the peak is higher than the fraction indicated by λ because both the Gaussian peak and the underlying power law contribute to the rate in that region. Therefore, $\lambda_{\{1,2\}} = 0$ does not mean that there are no events with masses $\sim 35\,M_\odot$.

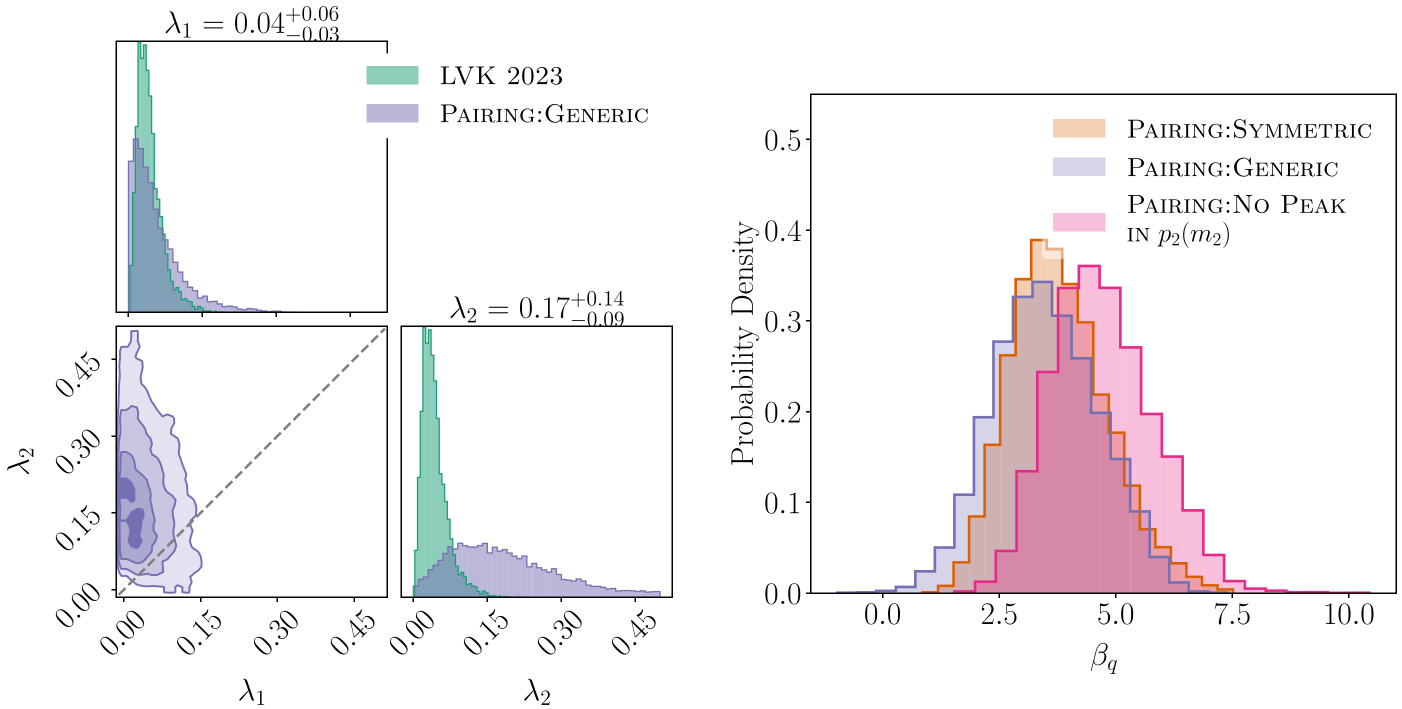

Figure 6. Hyperposteriors under pairing function models. Left: corner plot of hyperparameters governing the height of the Gaussian peak for the primary (λ_1) and secondary (λ_2) mass distributions under the PAIRING:GENERIC model. The medians and 90% credible intervals are shown above their respective marginal distributions. The medians indicate that roughly 17% of BBHs have secondary masses in the Gaussian peak while 4% have primary masses in the peak, though this preference for a larger peak in $p_2(m_2)$ may be due to the relatively poor constraint on λ_2 rather than being a true preference in the data. The lower left portion of the two-dimensional posterior is excluded, indicating that there must be a peak in either $p_1(m_1)$ or $p_2(m_2)$, and a slight preference exists for the peak to be in $p_2(m_2)$ (upper left portion of plot) or in both distributions (central portion of plot), though the peak being in $p_1(m_1)$ only is not completely ruled out. For reference, the one-dimensional posterior on λ under the LVK 2023 model is shown in green, with $\lambda = 0.04^{+0.03}_{-0.02}$ (Abbott et al. 2023a). The dashed gray line indicates where $\lambda_1 = \lambda_2$. Right: marginal posteriors on β_q, the hyperparameter controlling the steepness of the pairing function under the PAIRING:SYMMETRIC, PAIRING:GENERIC, and PAIRING:NO PEAK IN $p_2(m_2)$ models. When we set $\lambda_2 = 0$, β_q increases, indicating a preference for equal-mass components. This behavior is likely caused to accommodate an excess of events with $m_2 \sim 35\,M_\odot$, which can be caused either by a peak in $p_2(m_2)$ or by a peak in $p_1(m_1)$ and a strong preference for equal-mass binaries (see discussion in Appendix A.1).

be large. This indicates that either one of the two distributions has a large peak while the other has no peak, or that both distributions have moderate peaks.

In fact, the posterior on λ_2 peaks at a higher value than that on λ_1. Under the LVK 2023 model, the one-dimensional posterior on λ peaks between the marginal λ_1 and λ_2 posteriors of the PAIRING:GENERIC model. Additionally, in the PAIRING:GENERIC model, λ_1 has more support at zero relative to λ, while λ_2 has reduced support at zero. This hints at the intriguing possibility that the secondary masses may be driving the nonzero value of λ found by Abbott et al. (2021a, 2023a). In other words, it is possible that the $\sim 30\,M_\odot$ peak exists in the secondary mass distribution rather than in the primary mass distribution. However, the data are still consistent with $\lambda_1 = \lambda_2$: the dashed gray line in Figure 6 intersects the 1.5σ contour of the hyperposterior.

As shown in Appendix A.3, the LVK 2023 formalism behaves similarly to PAIRING:NO PEAK IN $p_2(m_2)$, which is nested within PAIRING:GENERIC when $\lambda_2 = 0$. Therefore, ruling out $\lambda_2 = 0$ would indicate that the pairing function formalism is strongly preferred. We find that $\lambda_2 = 0$ is disfavored but not ruled out: there is 4.2 times higher posterior density at $\lambda_2 = 0.17$ (its median a posteriori value) than at $\lambda_2 = 0$. It is therefore difficult to tell with the current number of detections whether the data prefer for the peak to exist in $p_1(m_1)$, $p_2(m_2)$, or both.

These three possibilities are somewhat degenerate because of the preference for equal-mass BBHs. To illustrate this, the right panel of Figure 6 compares the power-law spectral index, β_q, of the pairing function under PAIRING:GENERIC and PAIRING:NO PEAK IN $p_2(m_2)$. When there is no peak in $p_2(m_2)$, the β_q posterior shifts to higher values, which correspond to a larger preference for $q \sim 1$. This is because the PAIRING:NO PEAK IN $p_2(m_2)$ model has a peak in $p_1(m_1)$, so it allows for a high fraction of secondary masses to lie within the peak by making more binaries equal-mass. For comparison, we also plot the posterior on β_q under PAIRING:SYMMETRIC to show how a different set of assumptions about λ_2 changes β_q. The shift in β_q when $\lambda_2 = 0$ is larger than when $\lambda_2 = \lambda_1$, suggesting that it is driven by an excess of events with $m_2 \sim 35\,M_\odot$ rather than by other model assumptions or a different realization of the inference.

In summary, we find modest evidence for a peak in $p_2(m_2)$, suggesting that pairing function models may be preferred over the LVK 2023 formalism, though current observations do not allow us to definitively rule out that the peak exists only in $p_1(m_1)$. The data is consistent with $\Lambda_1 = \Lambda_2$, so there is no strong evidence against the possibility that m_1 and m_2 are drawn from the same underlying distribution, up to a pairing function.

3.4. BBHs Are Picky

Pairing functions provide an intuitive way to quantify whether the properties of one BH in a binary influence those of its companion. Binaries that pair up independently of each component's mass are described by a pairing function with $\beta_q = 0$, whereas a preference for equal-mass binaries is described by $\beta_q > 0$. We find $\beta_q > 0$ to be >99.99% for PAIRING:SYMMETRIC and 99.95% for PAIRING:GENERIC. This is consistent with an earlier study by Fishbach & Holz (2020), who used the pairing function formalism on GWTC-1 (Abbott et al. 2019a). However, BBHs are not maximally picky: the posterior on β_q does not rail against the high end of the prior bounds.

Our inferred two-dimensional mass distribution is generally consistent with that shown by Sadiq et al. (2023), who used a nonparametric approach and assumed that the population is symmetric under $m_1 \leftrightarrow m_2$. They showed evidence for peaks in the mass ratio distribution at $q \sim 0.5$ when $m_1 \sim 15\,M_\odot$ and $m_1 \sim 70\,M_\odot$, which they interpreted as a lack of preference for similar-mass pairings. However, we note that these peaks in the mass ratio distribution translate to peaks in the secondary mass distribution at $\sim 7\,M_\odot$ and $\sim 35\,M_\odot$, the same locations at which the primary mass distribution exhibits peaks. Parameterizations that assume symmetry under $m_1 \leftrightarrow m_2$ but do not use a pairing function are unable to disentangle the effects of preference for similar-mass pairings from features in one or both component mass distributions. We therefore conclude that both our results and those presented in Sadiq et al. (2023) are consistent with a preference for similar-mass pairings as well as with structure in the secondary mass distribution.

In this work, we only examine power-law forms for the pairing function, but more complex models are under investigation and will be presented in future work. If the pairing function is a more complex function of mass ratio, or if it depends on multiple parameters, such as the primary mass or spin (e.g., Farah et al. 2022), it may become easier to distinguish between pairing function models and models parameterized similarly to LVK 2023, as complex pairing functions may find more support away from the $m_1 = m_2$ diagonal. The same is true if the marginal mass ratio distribution has multiple modes, or if mass ratio is allowed to correlate with other parameters (e.g., the effective spin; Callister et al. 2021).

4. Looking to the Future

The LVK 2023 and pairing function models differ most in their predictions for the number of events in the region $m_1 \in [35, 40] \cap m_2 \in [m_{\min}, 30]$, because they disagree on whether to place m_2 in the Gaussian peak or somewhat evenly throughout the available parameter space. We can therefore determine which model will be preferred in the future by counting the number of detected events in that region.

It is expected that 260^{+330}_{-150} BBHs will be detected by the end of the LVK's fourth observing run (O4), and 870^{+1100}_{-480} BBHs will be detected by the end of the fifth observing run (O5; Kiendrebeogo et al. 2023). With 260 (870) total BBH detections, we expect 23.6 ± 4.6 (79.0 ± 8.5) BBHs to fall in the region $m_1 \in [35, 40] \cap m_2 \in [m_{\min}, 30]$ under the LVK 2023 model and 12.6 ± 3.5 (42.3 ± 6.3) BBHs under the PAIRING:SYMMETRIC model. This means that by the end of O4, we will be able to distinguish the LVK 2023 and PAIRING:SYMMETRIC models to $>2\sigma$, and we will be able to distinguish them to $>4\sigma$ by the end of O5.

Of course, it will be necessary to perform a full hierarchical analysis, as measurement uncertainties of detected systems will cause events to scatter into and out of this region. Additionally, in a hierarchical Bayesian context, posterior predictive checks performed on the event-level parameters (such as the true masses of individual events) are less sensitive than those performed on the observed data, such as the strain or the maximum-likelihood event parameters (Sinharay & Stern 2003; Gelman 2006; Bayarri & Castellanos 2007; Loredo 2013). This is why, for example, Fishbach et al. (2020) performed posterior predictive checks on the maximum-likelihood values of the observed and predicted events rather than on the posterior distributions of those events. Therefore, more sensitive posterior predictive checks may be able to distinguish between the two frameworks with fewer events than we project here.

5. Summary and Implications for Astrophysics of Merging BBHs

We present the first analysis of the secondary mass distribution of merging BBHs. This allows us to explore whether the primary and secondary component masses in merging BBHs are drawn from the same distribution, or whether the BBH mass distribution is asymmetric in $m_1 \leftrightarrow m_2$ as is commonly assumed in BBH population studies. We find the data to be consistent with two possibilities: either the primary and secondary mass distributions are similar, or a larger peak exists in the secondary mass distribution than in the primary mass distribution. In either scenario, a peak likely exists in the secondary mass distribution. This possibility is not considered in many previous analyses of the BBH population, which fix the secondary mass distribution to being a power law (e.g., Abbott et al. 2019b, 2021a, 2023a) or assume that m_1 and m_2 are interchangeable (e.g., Fishbach & Holz 2020; Sadiq et al. 2023). The existence of this secondary mass peak has implications for the formation channels of merging BBHs and the origin of the $\sim 35\,M_\odot$ peak.

If the mass distribution is indeed symmetric under $m_1 \leftrightarrow m_2$ as in our preferred model PAIRING:SYMMETRIC, this may imply that a large fraction of merging BBHs are formed through dynamical assembly, in which the two component BHs are born through a similar process and then find each other in a dense stellar environment. In this case, the pairing function and its dependence on mass ratio and/or total mass may encode valuable information about dynamical processes like mass segregation and binary exchanges.

The appearance of a peak in both primary and secondary mass distributions is consistent with an origin in the BH remnant mass function. However, the peak location at $31.4^{+2.3}_{-2.6}\,M_\odot$ is in tension with predictions for pair-instability supernovae, so this feature is likely caused by another astrophysical process. As another application of our work, we suggest that future spectral siren measurements consider the PAIRING:SYMMETRIC model when inferring the Hubble constant, because it provides an improved measurement of the peak location relative to the LVK 2023 model and the peak location is strongly correlated with the inferred value of H_0.

On the other hand, the data remain consistent with an asymmetric BBH mass distribution $p(m_1, m_2)$, in which $p_1(m_1) \neq p_2(m_2)$, as in the PAIRING:GENERIC model. This would imply that a significant fraction of merging BBHs originate from field binaries, in which "primary" and "secondary" are physically meaningful labels if they tend to correspond to the firstborn and secondborn BHs in a binary. The component mass distributions $p_1(m_1)$ and $p_2(m_2)$ would then encode the binary IMF and the highly uncertain physics of binary stellar evolution. Specifically, the location and prominence of features in each component mass distribution may provide insight into how mass loss versus mass accretion affects stellar evolution and collapse, including their effect on supernovae and the BH remnant mass function. For example, when mass transfer is unstable, a peak is expected at $\sim 15\,M_\odot$ in the primary mass distribution, but near the minimum BH mass in the secondary mass distribution, though the relative locations of these peaks depend on the common envelope and supernova kick physics (van Son et al. 2022b).

Notably, if the BBH mass distribution is asymmetric, we find that it is possible that no peak exists in the primary mass distribution, and the previously identified peak in primary mass is actually driven by an overdensity in the secondary mass distribution. A peak that is more prominent in $p_2(m_2)$ than in $p_1(m_1)$ may imply that pulsational pair-instability supernovae are more frequent or are shifted to lower masses among secondborn BHs.

The degree of asymmetry in the BBH mass distribution provides insight into the frequency of mass inversion in binary stars, providing a complementary probe to BH spins (Mould et al. 2022). If mass inversion never occurs, m_1 will typically be the remnant of the donor star. However, if mass inversion is the norm, m_1 will be the remnant of the accretor star. A perfectly symmetric BBH distribution under isolated binary evolution would imply that mass inversion happens exactly 50% of the time, causing the primary and secondary mass distributions to be indistinguishable even though the binary physics imparts different distributions on the firstborn versus the secondborn BH, but this is statistically unlikely. Determining whether primary and secondary BH masses follow the same underlying distribution is therefore a novel and promising probe of formation channels. However, this test should be interpreted within a full BBH population inference, as theoretical models should make consistent predictions for the mass, spin, and redshift distributions.

We only explore parametric models in this work, and can therefore only comment on the features in the mass distribution explicitly parameterized by our model. Future investigations may find it beneficial to use a nonparametric approach that separately fits the primary and secondary mass distributions by employing a pairing function.

We additionally find that BHs in merging binaries have a strong preference to pair with similar-mass BHs for all forms of the secondary mass distribution considered in this work. This is consistent with the results presented in Fishbach & Holz (2020), though we more confidently exclude the scenario in which BBHs pair independently of the mass ratio since we now have many more detected BBHs.

By the end of the LVK's fifth observing run, we expect to confidently distinguish between the different scenarios presented here for the BBH mass distribution. In this work, we find that secondary masses in merging BBH systems likely display a peak at $\sim 35\,M_\odot$, whereas previous results identified this peak exclusively among primary masses. With a few hundred additional BBH observations, we expect to determine whether both component mass distributions have a peak at the same location and if the peak is more prominent among secondary masses.

Directly incorporating the distribution of secondary masses can serve as an important tool to constrain the formation mechanisms of BBHs.

Acknowledgments

The authors thank Sharan Banagiri, Thomas Dent, Zoheyr Doctor, Will Farr, Davide Gerosa, Tom Loredo, and Jam Sadiq for helpful discussions. A.M.F. is supported by the National Science Foundation Graduate Research Fellowship Program under grant No. DGE-1746045. D.E.H. is supported by NSF grants AST-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. This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based on work supported by NSF's LIGO Laboratory, which is a major facility fully funded by the National Science Foundation. The authors are grateful for the computational resources provided by the LIGO Laboratory and supported by National Science Foundation grants PHY-0757058 and PHY-0823459.

Facilities: LIGO, EGO:Virgo.

Software: `gwpopulation` (Talbot et al. 2019), `bilby` (Ashton et al. 2019), `scipy` (Virtanen et al. 2020), `corner.py` (Foreman-mackey 2016).

Appendix A
Comparison of Common Mass Distribution Parameterizations

A.1. Differences between Mass Ratio Distributions and Pairing Functions

Given a form for the primary mass distribution, there are several ways to construct a two-dimensional mass distribution. We discuss two possibilities here that are common in the literature, showing the different effects they have on the resulting two-dimensional mass distribution.

The main qualitative differences between the two parameterizations are illustrated in Figure 7. The top row has three examples of mass distributions that can be described by a model of the form

$$p(m_1, m_2|\Lambda_m, \beta) = p(m_1|\Lambda_m)p(q|m_1, m_{\min}, \beta) \tag{A1}$$

for different values of β. Here, m_1 is the mass of the heavier component in the binary, m_2 is the mass of the lighter component, and $q = m_2/m_1 \leqslant 1$ is the mass ratio. This can be equivalently written as

$$p(m_1, m_2|\Lambda_m, \beta) = p(m_1|\Lambda_m)p(m_2|m_1, \beta), \tag{A2}$$

since $m_2 = qm_1$. The parameterization described in Equation (A1) is used by all parametric models presented in Abbott et al. (2021a, 2023a) that were used to model the primary mass distribution of BBHs, such as BROKEN POWER LAW and POWER LAW + PEAK, and is also commonly used in other analyses (e.g., Fishbach & Holz 2017; Kovetz et al. 2017; Talbot & Thrane 2018; Tiwari 2021; Edelman et al. 2022a; Callister & Farr 2023; Godfrey et al. 2023). For the remainder of this appendix, we will refer to models parameterized by Equation (A1)

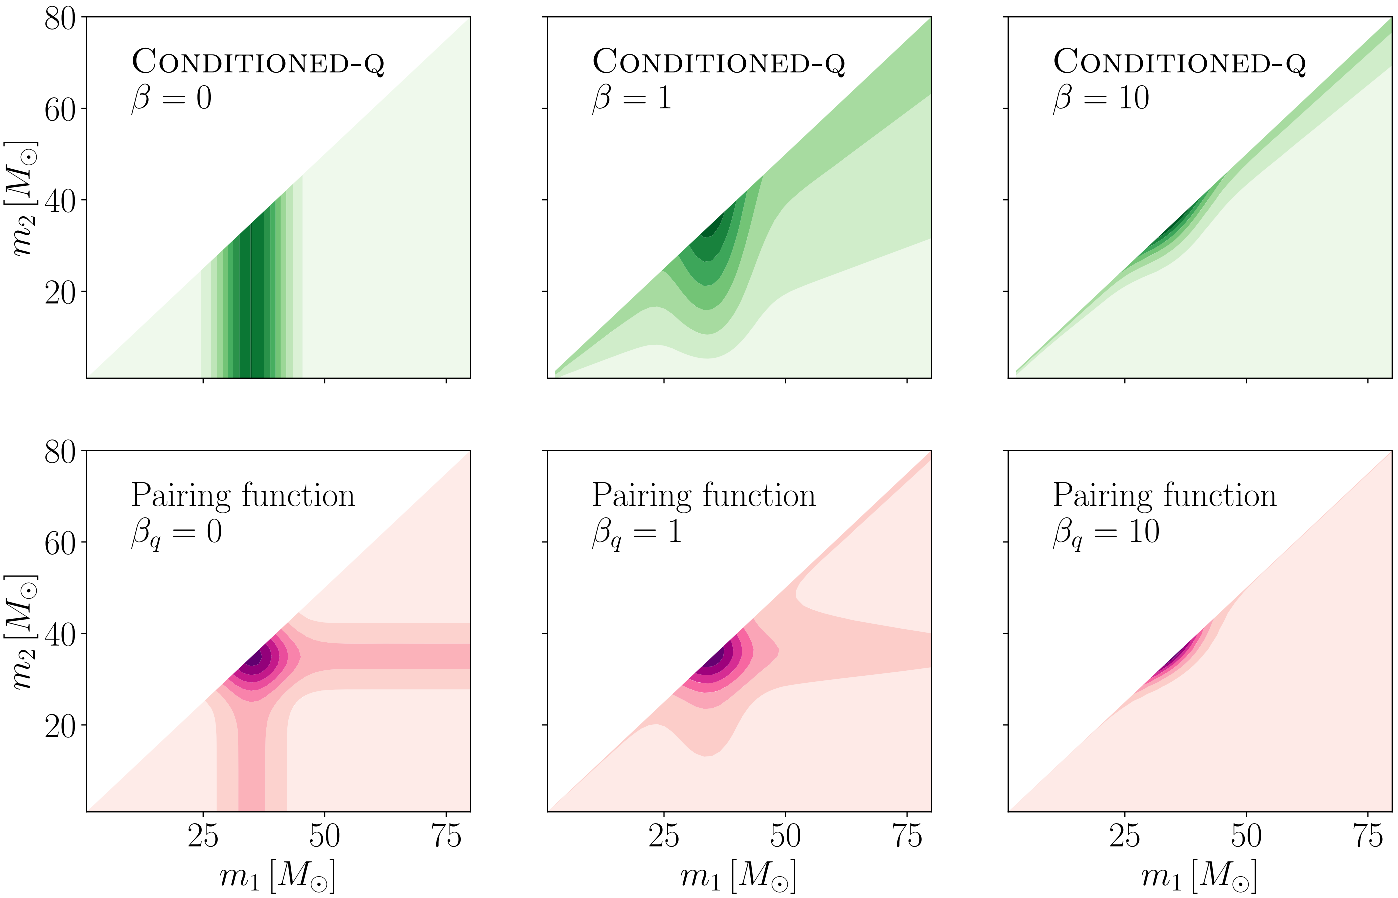

Figure 7. Illustration of some possible two-dimensional mass distributions under the commonly used "LVK 2023" formalism described by Equation (A1) (top row) and the pairing function formalism described by Equation (A3) (bottom row). Overdensities/peaks in the mass distribution appear as darker filled contours in these figures. The LVK 2023 formalism is only able to produce models with features in the m_1 distribution, as shown by the vertical bands in the top row, whereas the pairing function formalism can model features in either m_1 or m_2, or both. The different columns correspond to different power-law spectral indices for the mass ratio distribution (top row, β) and the mass-ratio-dependent pairing function (bottom row, β_q). In the case of a steeply rising mass ratio distribution, or if components strongly prefer to pair with nearly equal-mass partners, the LVK 2023 model and the pairing function model become difficult to tell apart and likely explain the data equally well, as shown by the two panels in the leftmost column. The diagonal contours in the middle and right columns are caused by a preference for equal-mass binaries and follow lines of constant mass ratio.

as "CONDITIONED-Q," since they require the mass ratio distribution to be explicitly conditioned on the primary mass.

The bottom row of Figure 7 has three examples of mass distributions that can be described using a "pairing function," f (Fishbach & Holz 2020). Models with pairing functions have the form

$$p(m_1, m_2|\Lambda) = p_1(m_1|\Lambda_1)p_2(m_2|\Lambda_2)f(q;\, \beta_q) \tag{A3}$$

where $f(q)$ is a pairing function that depends on the mass ratio of the system[10], and $\Lambda = \{\Lambda_1, \Lambda_2, \beta_q\}$ is the set of all model hyperparameters. In this work, we use a pairing function of the form $f(q;\, \beta_q) = q^{\beta_q}\Theta(q \leqslant 1)$, though other forms may provide a better fit to the data (e.g., Farah et al. 2022). In the examples illustrated in Figure 7, the primary and secondary mass distributions are equivalent, so $p_1(m) = p_2(m) \equiv p(m)$. Alternatively, $\Lambda_1 = \Lambda_2$. This describes a situation in which there is a single underlying mass distribution from which both components are drawn. The pairing function then describes how likely the two components are to be combined in a merging binary based on their mass ratio. A pairing function that prefers equal-mass binaries will cause a marginal mass ratio distribution that has more support near $q = 1$, but the inverse is not necessarily true.

The parameterization in Equation (A3) factorizes the possibility that components in binaries prefer to be of near-equal mass and the possibility that the primary and secondary masses have distinct probability distributions. In other words, pairing functions allow us to model the secondary mass separately from the primary mass, while also allowing for the possibility that component BHs prefer to pair with similar-mass BHs.

In Figure 7, a peak at $35\,M_\odot$ is placed in both models to show the effects of such features in both cases. One consequence of LVK 2023 models is that features such as Gaussian peaks can only appear in the the primary mass distribution. This is shown by the vertical bands in the top row of Figure 7 and the lack of horizontal bands, since a band in the vertical (horizontal) direction is caused by a peak or dip in the primary (secondary) mass distribution for a range of secondary (primary) masses. For pairing function models, features can appear in $p_1(m_1)$, $p_2(m_2)$, or both. We have illustrated the case in which the same feature appears in both component mass distributions, and this appears as both vertical and horizontal bands in the bottom row of Figure 7. If Λ_1 and Λ_2 are allowed

---

[10] In principle, the pairing function can be parameterized in terms of any observable parameter (e.g., total mass).

to differ, features could appear in only one of these distributions. This would cause there to be only horizontal bands if features only existed in $p_2(m_2)$, and only vertical bands if features only existed in $p_1(m_1)$. Features are also able to appear in different locations in $p_1(m_1)$ versus $p_2(m_2)$ under the pairing function formalism. However, the LVK 2023 formalism only allows for bands in the vertical direction, meaning that it is not flexible enough to capture the true underlying distributions with features in $p_2(m_2)$. The behavior of the LVK 2023 formalism can in general be approximated by the pairing function formalism, while the opposite is not true.

The different columns in Figure 7 correspond to different power-law spectral indices for the mass ratio distribution (top row, β) and the mass-ratio-dependent pairing function (bottom row, β_q). The top row's leftmost panel has a uniform mass ratio distribution, the middle panel has a mass ratio distribution that mildly favors equal-mass binaries, and the right panel's mass ratio distribution strongly favors equal-mass binaries. Analogously, the bottom row's leftmost panel shows a model where components in the binary are allowed to pair up independently of the mass ratio, the middle panel shows a model where components have a slight preference to pair up with partners that are of equal mass, and the rightmost panel shows the case where components are "picky": they almost always pair up with equal-mass partners (Fishbach & Holz 2020). When the mass ratio distribution is broad, or when components pair up nearly independently of the mass ratio, the LVK 2023 models produce noticeably different distributions from those produced by the pairing function models. However, in the case of very picky binaries or a steeply rising mass ratio distribution, the LVK 2023 and pairing function models become difficult to tell apart, and likely explain the data equally well. There is therefore a degeneracy between the steepness of the pairing function and the existence of distinct features in the two mass distributions (see Tiwari 2023 for a discussion of this phenomenon in terms of Jacobian transformations).

Fortunately, as we show in Section 3.4, we measure $\beta \sim 3.5$ and $\beta_q \sim 1$, so the data lie somewhere between the middle and rightmost columns.[11] This means that differentiating between the two scenarios will be difficult, but possible given enough data.

Mass distributions of compact objects are often visualized through a plot of the marginal component mass distributions. The marginal m_2 distribution is defined as

$$p(m_2|\Lambda) = \int dm_1 p(m_1, m_2|\Lambda), \tag{A4}$$

where $p(m_1, m_2|\Lambda)$ can be parameterized in the way described in Equation (A1) or (A3). However, in all but the top left panel in Figure 7, the marginal secondary mass distribution exhibits a peak between $30\,M_\odot$ and $40\,M_\odot$, even though the secondary mass distribution is not able to have any features on its own. This is because features in the primary mass distribution induce features in the marginal secondary mass distribution if equal-mass binaries are preferred: when a binary's primary mass is within the peak, its secondary mass is also likely to be in that peak simply because $m_2 \simeq m_1$ is preferred. It is therefore difficult to distinguish between these different scenarios by looking at the marginal distributions alone. Instead, we analyze two-dimensional distributions such as the ones illustrated in Figure 7, as well as secondary mass distributions conditioned at various values of m_1. The latter can be thought of as one-dimensional slices of the former.

It is our goal to determine whether the data prefer models described by pairing functions or those described by LVK 2023 functions. We also aim to determine if the primary and secondary masses are drawn from the same distribution, and whether we can draw physical insights from features that appear in the primary mass distribution, the secondary mass distribution, or both.

A.2. "Pickiness"

Under the LVK 2023 formalism, we cannot determine how BHs in binaries choose their companions, though we can gain some insight from their marginal mass ratio distribution. LVK 2023 parameterizes the mass ratio distribution as a power law with spectral index β, where $\beta > 0$ corresponds to mass ratio distributions with more support for similar-mass binaries. Abbott et al. (2023a) found $\beta = 1.1^{+1.7}_{-1.3}$, which similarly indicates a preference for equal-mass binaries. Note that the value for β under LVK 2023 is noticeably different from the value found for β_q under the pairing function models because the two parameters cause different behaviors in the two-dimensional mass distribution within their respective models. A low value for β does not imply that BBHs are not picky, just that the marginal mass ratio distribution rises slowly. As shown in Appendix B.3, the pairing function models and LVK 2023 model all produce near-identical marginal mass ratio distributions, despite different values for β_q and β. Low values for β are partially due to the fact that the low-q end of the marginal mass ratio distribution will always be suppressed because of the existence of a minimum BH mass. This minimum mass makes it impossible to get extreme mass ratios unless m_1 is very large, and the m_1 distribution has very little support for $m_1 \gtrapprox 60\,M_\odot$. Therefore, β does not need to be large in order for the marginal mass ratio distribution to strongly disfavor unequal-mass binaries. In fact, the existence of a minimum mass plus a tapering at high m_1 means that even when $\beta < 0$, the marginal mass ratio distribution rises toward $q = 1$.

A.3. Mimicking LVK 2023 with a Pairing Function Model

We show that the PAIRING:NO PEAK IN $p_2(m_2)$ model approximates the morphology of the LVK 2023 model, which uses the LVK 2023 formalism.

The two rightmost panels of Figure 8 show the two-dimensional PPDs for the LVK 2023 and PAIRING:NO PEAK IN $p_2(m_2)$ models. Both exhibit vertical bands and no horizontal bands, and have a similar-magnitude drop in merger rate moving away from the diagonal.

The leftmost panel of Figure 8 shows the conditional m_2 distribution for both models. While the slopes of the two models differ slightly, a peak in m_2 appears in neither. We therefore find PAIRING:NO PEAK IN $p_2(m_2)$ to be an appropriate proxy for the behavior of the LVK 2023 model for the purposes of this work. A future measurement ruling out $\lambda_2 = 0$ with high confidence would therefore serve as evidence for the preference for pairing function models over LVK 2023–like models.

---

[11] The value for β under LVK 2023 is different from the value found for β_q under the pairing function models because the two parameters cause different behaviors in the two-dimensional mass distribution within their respective models. A low value for β does not imply that BBHs are not picky, just that the marginal mass ratio distribution rises slowly. We discuss this in more detail in Appendix A.2.

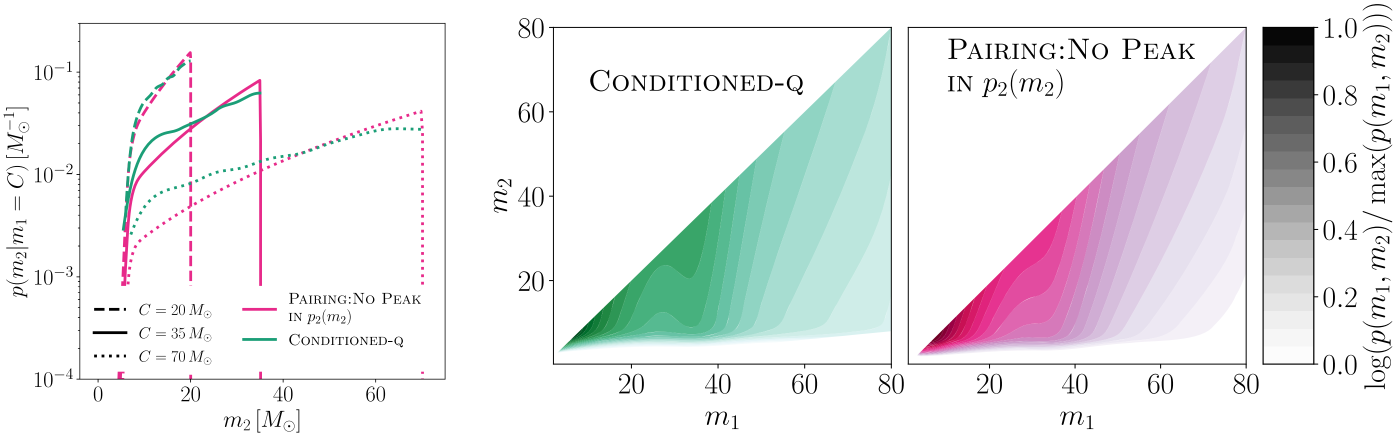

Figure 8. Two-dimensional (right) and conditional m_2 (left) PPDs for the LVK 2023 and PAIRING:NO PEAK IN $p_2(m_2)$ models.

Appendix B
General Population Model

B.1. Base Model

We model the two-dimensional mass distribution by constructing separate one-dimensional distributions for the primary and secondary masses and combining draws from these two distributions according to a pairing function, as in Equation (A3). Fishbach & Holz (2020) found the pairing function is most informative when parameterized by the mass ratio of the binary, so we adopt a pairing function described by a power law in mass ratio.

Because our aim is to learn whether the primary and secondary masses are consistent with being drawn from the same distribution (up to a pairing function and subject to the constraint that $m_1 \geqslant m_2$), we describe the primary and secondary mass distributions separately, but according to the same functional form. We model each of the one-dimensional mass distributions as a mixture model between a smoothed power-law component and a Gaussian component G in order to make direct comparisons to the POWER LAW + PEAK model used by the LVK population analysis to describe the distribution of primary masses (Talbot & Thrane 2018; Abbott et al. 2023a). Explicitly,

$$\begin{aligned} p(m_1|\Lambda_1) \propto &\left[(1-\lambda_1)\Theta(m_1 > m_{\min,1})\Theta(m_1 < m_{\max,1})\right.\\ &\left.\times\left(\frac{\alpha_1+1}{m_{\max,1}^{\alpha_1+1} - m_{\min,1}^{\alpha_1+1}}\right)m_1^{-\alpha_1} + \lambda_1 G(m_1|\mu_1, \sigma_1)\right]\\ &\times S(m_1|m_{\min,1}, \delta_1)\\ p(m_2|\Lambda_2) \propto &\left[(1-\lambda_2)\Theta(m_2 > m_{\min,2})\Theta(m_2 < m_{\max,2})\right.\\ &\left.\times\left(\frac{\alpha_2+1}{m_{\max,2}^{\alpha_2+1} - m_{\min,2}^{\alpha_2+1}}\right)m_2^{-\alpha_2} + \lambda_2 G(m_2|\mu_2, \sigma_2)\right]\\ &\times S(m_2|m_{\min,2}, \delta_2). \end{aligned} \tag{B1}$$

Here, $G(m_{\{1,2\}}|\mu_{\{1,2\}}, \sigma_{\{1,2\}})$ is a normalized Gaussian distribution with mean $\mu_{\{1,2\}}$ and width $\sigma_{\{1,2\}}$. The parameter $\lambda_{\{1,2\}}$ is a mixing fraction determining the relative prevalence of mergers in the power-law and Gaussian components, and $S(m_{\{1,2\}}|m_{\min,\{1,2\}}, \delta_{\{1,2\}})$ is a smoothing function that rises from 0 to 1 over the interval $(m_{\min,\{1,2\}}, m_{\min,\{1,2\}} + \delta_{\{1,2\}})$. Λ_1 is then the set of hyperparameters $\{m_{\min,1}, m_{\max,1}, \alpha_1, \lambda_1, \mu_1, \sigma_1, \delta_1\}$ and $\Lambda_2 = \{m_{\min,2}, m_{\max,2}, \alpha_2, \lambda_2, \mu_2, \sigma_2, \delta_2\}$.

In all models considered in this work, the redshift distribution is modeled as a power law with spectral index κ (Fishbach et al. 2018) such that

$$p\left(z,\right), \propto , \frac{dV_c}{dz}, \left(\frac{1}{1+z}\right),\left(, 1, + ,z\right)^{\kappa}. \tag{B2}$$

We use the DEFAULT spin model from Abbott et al. (2021a, 2023a) to describe the spin magnitudes and tilts of each component. These are the same redshift and spin distributions used with the POWER LAW + PEAK mass model in the analysis presented in Abbott et al. (2023a).

A fit to the model described above is provided in Appendix B.2, though we focus on specific variations nested within this more general model for the remainder of this work.

B.2. Fit to General Base Model

We present results of a fit to the most general form of the pairing function model described in Appendix B.1 (Equation (B1)). Here, we do not fix any parameters to be equal between Λ_1 and Λ_2, and instead infer them separately.

Figure 9 shows the posterior of all mass-related hyperparameters within this model in the form of a corner plot. Because of the large number of free parameters, several are not well constrained. Nonetheless, all parameters describing $p_1(m_1)$ are consistent with those describing $p_2(m_2)$. Additionally, strong correlations exist between α_1, α_2, and β_q, making it difficult to meaningfully constrain all three parameters. We therefore set $\alpha_1 = \alpha_2$ in the PAIRING:GENERIC model.

Figure 10 shows the underlying distributions, $p_1(m_1)$ and $p_2(m_2)$, inferred by the most general version of the base model. As expected, the constraints on these distributions weaken relative to those of the PAIRING:GENERIC model. However it is still clear that $p_1(m_1)$ and $p_2(m_2)$ are consistent with each other and that it is possibile that $p_2(m_2)$ has a larger peak than $p_1(m_1)$ does. Notably, there now seems to be very little evidence for a peak in $p_1(m_1)$, whereas a clear peak still exists in $p_2(m_2)$.

These results are consistent with what has been presented using the PAIRING:GENERIC model in the main body of this work.

B.3. Model Comparison

In this section, we compare the various models considered in this work, including the LVK 2023 model from LVK 2023.

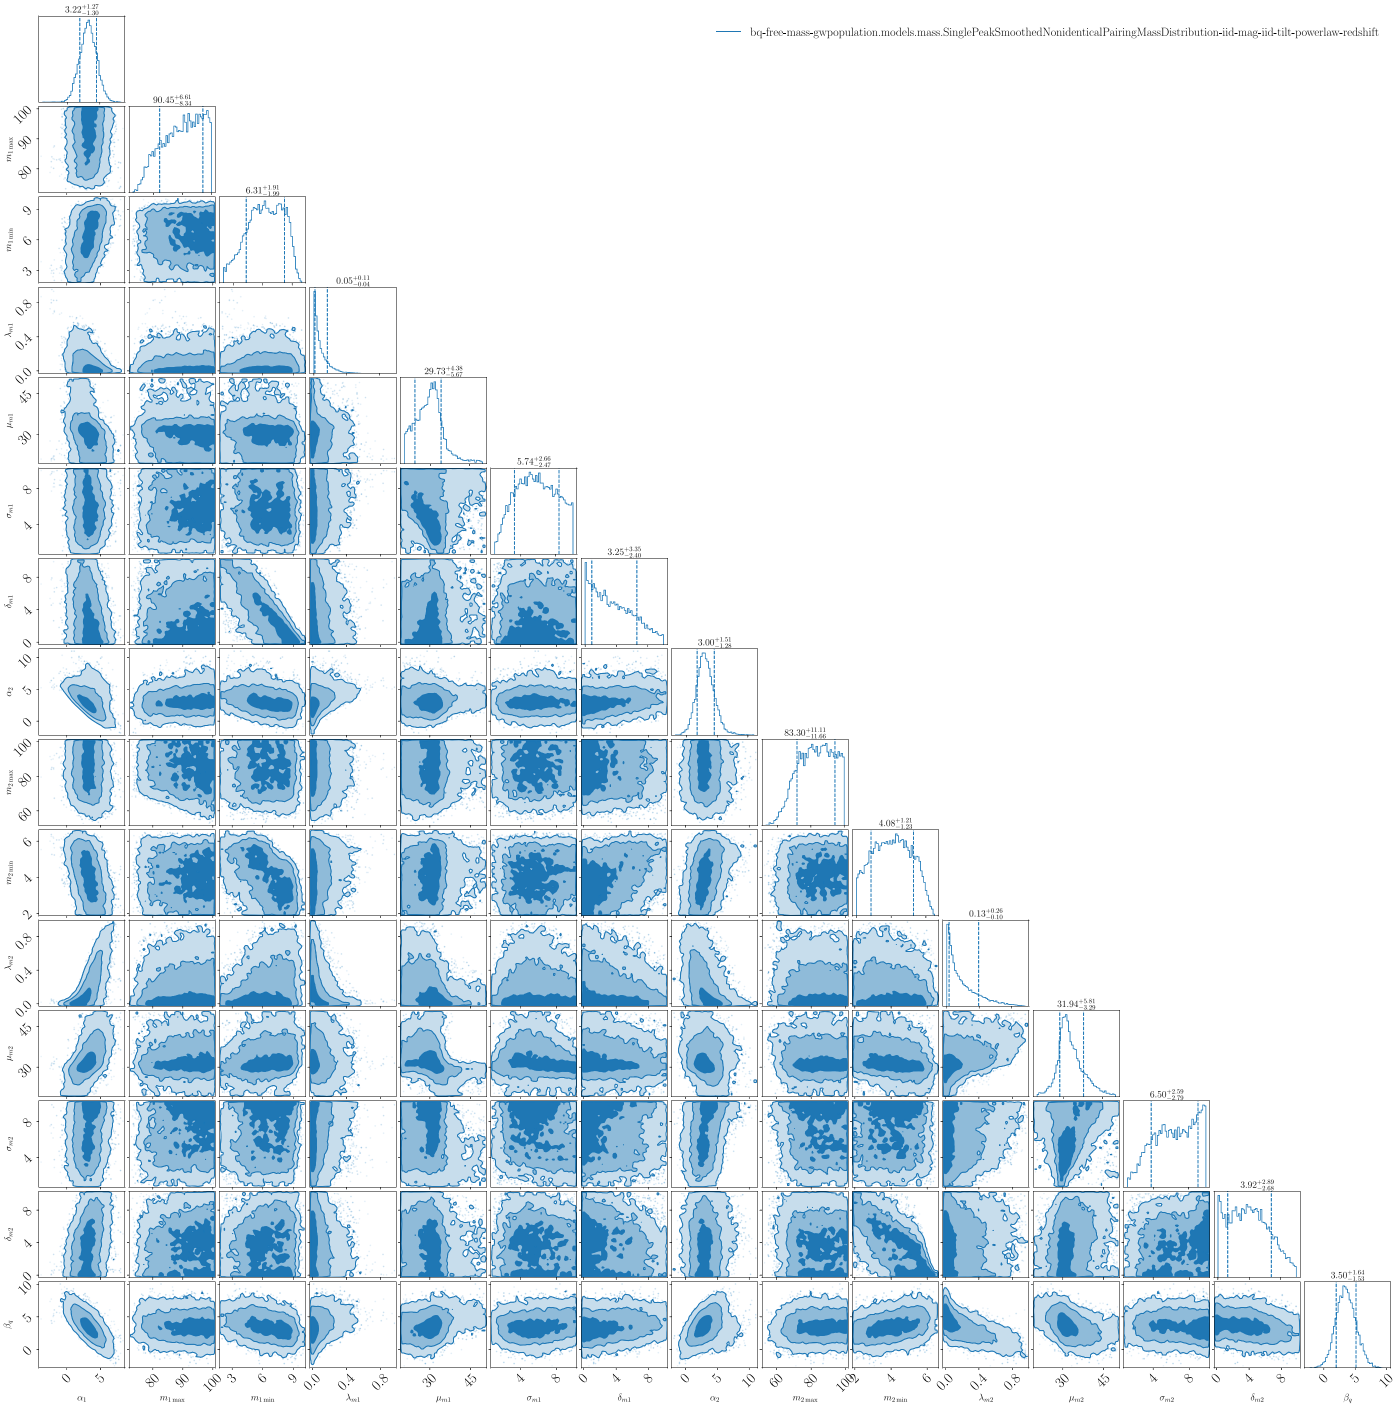

Figure 9. Corner plot of all mass-related hyperposteriors in the most general form of the base model. Power-law spectral indices governing the slope of the primary and secondary mass distributions are degenerate with each other and with the pairing function power law, so only two of the three parameters can be meaningfully constrained at a time. All other parameters that perform the same function in $p_1(m_1)$ and $p_2(m_2)$ are consistent between the two underlying distributions.

Table 1 lists the hyperprior choices made to construct each model, along with descriptions of each hyperparameter.

Figure 11 shows the m_1, m_2, and q PPDs, marginalized over all other parameters. These show a high level of agreement between the different models, despite their two-dimensional PPDs differing in Figure 2. Notably, the marginal secondary mass distribution exhibits a peak even for the LVK 2023 model. As discussed in Appendix A.1, this is caused by the preference for equal masses: if m_1 is in the $\sim 35\,M_\odot$ peak, a preference for $m_2 \approx m_1$ will cause m_2 to also preferentially be in that region, causing higher probability density for $m_2 \sim 35\,M_\odot$ than elsewhere. Marginal PPDs are therefore not very sensitive to whether m_2 prefers its own features independently of the existence of a preference for equal-mass binaries. Instead, we examine the underlying component mass distributions explicitly modeled by PAIRING:SYMMETRIC and PAIRING:GENERIC, as well as the values of informative hyperparameters in these models, to determine whether $p_2(m_2)$ differs from $p_1(m_1)$. For comparisons to LVK 2023, we plot secondary mass distributions conditioned on various values of m_1.

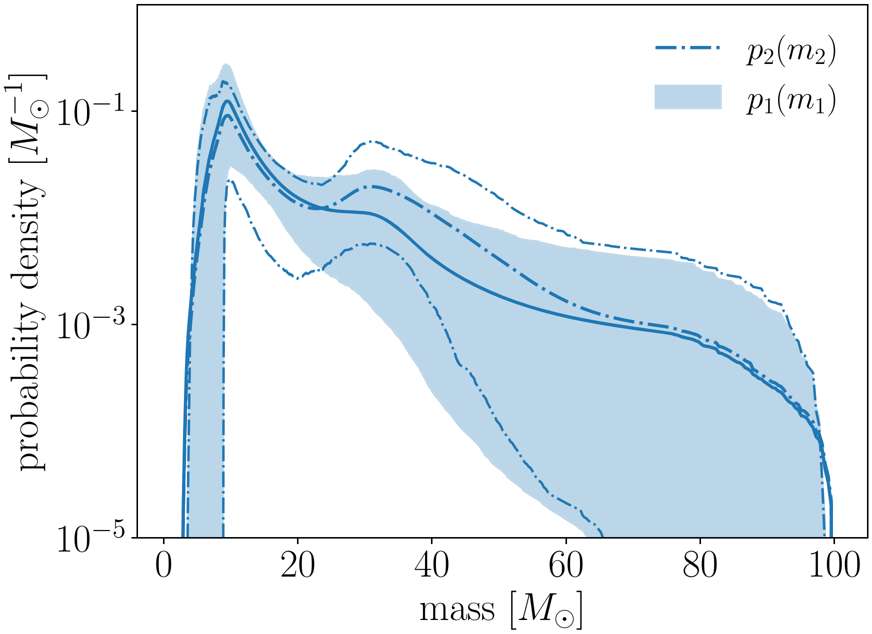

Figure 10. Underlying distributions (i.e., before a pairing function is applied) of the primary (shaded band) and secondary (dotted–dashed lines) masses under the most general form of the base model. Despite fitting all parameters separately between the two distributions, we find that they appear consistent with each other. However, both are relatively poorly constrained. Interestingly, the support for a peak in $p_1(m_1)$ lessens in this more generalized model, while the support for a peak in $p_2(m_2)$ remains the same as in PAIRING:GENERIC.

Table 1
Hyperparameters of Our Mass Model and Hyperpriors Corresponding to Specific Model Variations

Parameter	Description	Prior			
		Base Model	PAIRING: GENERIC	PAIRING: SYMMETRIC	LVK 2023
β_q	Spectral index for the power law of the pairing function	U(−4, 12)			...
β	Spectral index for the power law of the mass ratio distribution	...			U (−4, 12)
α_1	Spectral index for the power law of the primary mass distribution	U(−4, 12)			
α_2	Spectral index for the power law of the secondary mass distribution	U(−4, 12)	$\alpha_2 = \alpha_1$		...
$m_{\min,1}$	Minimum mass of the primary mass distribution	U(2 $M_\odot$, 10 $M_\odot$)			
$m_{\min,2}$	Minimum mass of the secondary mass distribution	U(2 $M_\odot$, 10 $M_\odot$)	$m_{\min,2} = m_{\min,1}$		...
$m_{\max,1}$	Maximum mass of the primary mass distribution	U(30 $M_\odot$, 100 $M_\odot$)			
$m_{\max,2}$	Maximum mass of the secondary mass distribution	U(30 $M_\odot$, 100 $M_\odot$)	$m_{\max,2} = m_{\max,1}$		...
δ_1	Range of tapering at the low end of the primary mass distribution	U(0 $M_\odot$, 10 $M_\odot$)			
δ_2	Range of tapering at the low end of the secondary mass distribution	U(0 $M_\odot$, 10 $M_\odot$)	$\delta_2 = \delta_1$		...
λ_1	Fraction of systems with primary mass in the Gaussian component	U(0, 1)			
λ_2	Fraction of systems with secondary mass in the Gaussian component	U(0, 1)		$\lambda_2 = \lambda_1$	...
μ_1	Mean of the Gaussian component in the primary mass distribution	U(20 $M_\odot$, 50 $M_\odot$)			
μ_2	Mean of the Gaussian component in the secondary mass distribution	U(20 $M_\odot$, 50 $M_\odot$)		$\mu_2 = \mu_1$	...
σ_1	Width of the Gaussian component in the primary mass distribution	U(1 $M_\odot$, 10 $M_\odot$)			
σ_2	Width of the Gaussian component in the secondary mass distribution	U(1 $M_\odot$, 10 $M_\odot$)		$\sigma_2 = \sigma_1$	...

Notes. For comparison, we also show the hyperpriors for the LVK 2023 model, which uses the LVK 2023 framework. We denote a 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 ellipses.

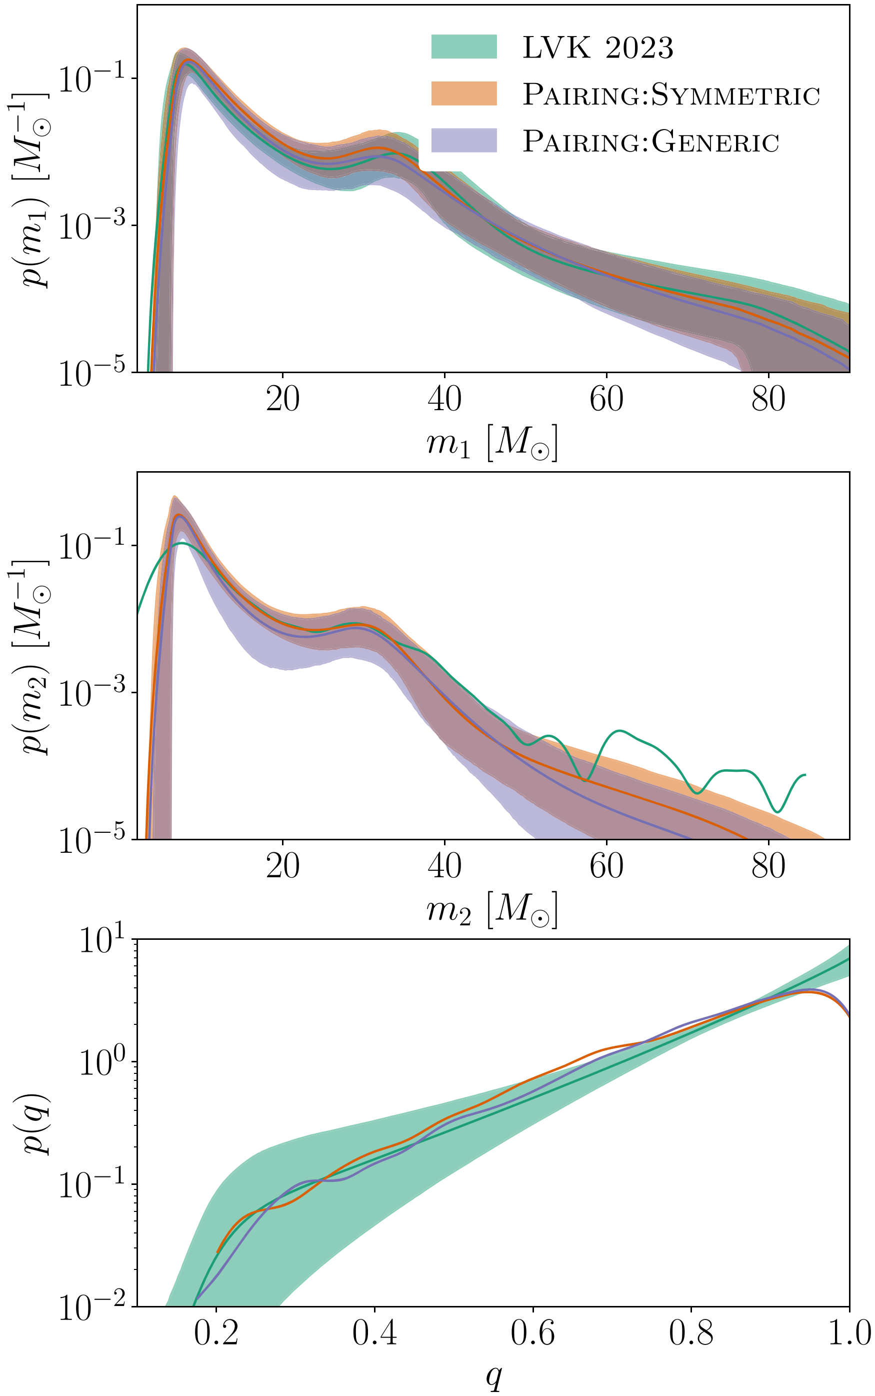

Figure 11. Marginal PPDs of primary mass (top), secondary mass (middle), and mass ratio (bottom) for LVK 2023 (green), PAIRING:SYMMETRIC (orange), and PAIRING:GENERIC (purple). Solid lines are the mean value and shaded regions represent the 90% credible interval. These are obtained by integrating the two-dimensional PPDs shown in Figure 2 along each dimension in turn. Since the LVK 2023 model is parameterized in terms of m_1 and q, we reconstruct its marginal m_2 distribution by sampling from the joint distribution and creating a kernel density estimate, causing some artificial "wiggles" in these plots. The same is true for the marginal q distributions of the pairing function models. Despite the two-dimensional PPDs appearing different between LVK 2023 and the other two models, the marginal PPDs all appear similar. This is because features in the primary mass distribution induce features in the marginal secondary mass distribution if equal-mass binaries are preferred. Therefore, marginal PPDs are not sensitive to the features in the two-dimensional mass distribution in which we are interested.

Appendix C
Statistical Framework

In this appendix, we describe the hierarchical Bayesian analysis used to fit the population models described in Appendix B and Section 2 to BBHs in GWTC-3 (Abbott et al. 2023c, 2023d).

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

$$p(\Lambda|\{D_j\}) = p(\Lambda) \prod_j^N \frac{p(D_j|\Lambda)}{\mathcal{E}(\Lambda)}, \tag{C1}$$

where

$$p(D_j|\Lambda) = \int dm_1 dm_2 ds_1 ds_2 dz\, p(z)p(s_1, s_2)p(m_1, m_2|\Lambda) \times p(D_j|m_1, m_2, s_1, s_2, z) \tag{C2}$$

is the marginal likelihood for the *j*th event,

$$\mathcal{E}(\Lambda) = \int dm_1 dm_2 ds_1 ds_2 dz\, p(z)p(s_1, s_2) \times p(m_1, m_2|\Lambda)P(\mathrm{det}|m_1, m_2, s_1, s_2, z) \tag{C3}$$

is the fraction of detectable events in the population described by Λ, and $P(\mathrm{det}|m_1, m_2, s_1, s_2, z)$ 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.

In practice, the high-dimensional integrals in Equations (C2) and (C3) are approximated via importance sampling (see Essick & Fishbach 2021; Essick & Farr 2022 for a detailed explanation of this process). Given a set of N_j event-level samples drawn from the posterior for the *j*th event that used a reference prior $p_{\mathrm{ref}}(m_1, m_2, s_1, s_2, z)$, we approximate

$$\frac{p(D_j|\Lambda)}{p_{\mathrm{ref}}(D_j)} \approx \frac{1}{N_j}\sum_k^{N_j} \frac{p(m_1^{(k)}, m_2^{(k)}, s_1^{(k)}, s_2^{(k)}, z^{(k)}|\Lambda)}{p_{\mathrm{ref}}(m_1^{(k)}, m_2^{(k)}, s_1^{(k)}, s_2^{(k)}, z^{(k)})} \tag{C4}$$

where $p_{\mathrm{ref}}(D_j)$, 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_{draw}, we can approximate Equation (C3) with a sum over the subset of m detected signals:

$$\mathcal{E}(\Lambda) \approx \frac{1}{M}\sum_k^{m} \frac{p(m_1^{(k)}, m_2^{(k)}, s_1^{(k)}, s_2^{(k)}, z^{(k)}|\Lambda)}{p_{\mathrm{draw}}(m_1^{(k)}, m_2^{(k)}, s_1^{(k)}, s_2^{(k)}, z^{(k)})}. \tag{C5}$$

We obtain the detectable set of m signals by injecting waveforms of these signals into the measured detector strain and running the search pipelines (Usman et al. 2016; Aubin et al. 2021; Cannon et al. 2021; Drago et al. 2021) used by the LVK to detect events to obtain a false-alarm rate (FAR) for each injected signal. This process was performed by the LVK for its population analysis (Abbott et al. 2023a) and we use the resulting publicly released data product in this work (LIGO Scientific Collaboration et al. 2021). We then consider an event detected if it has a FAR of less than 1 per year in at least one pipeline, and apply this threshold both to injected signals and to the real GW events in GWTC-3.

We sample from the posterior distribution in Equation (C1) using the approximations in Equations (C4) and (C5) 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 hyperposteriors and the hyperpriors described in Appendix B.1.

ORCID iDs

Amanda M. Farah https://orcid.org/0000-0002-6121-0285
Maya Fishbach https://orcid.org/0000-0002-1980-5293
Daniel E. Holz https://orcid.org/0000-0002-0175-5064

References

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, CQGra, 32, 074001
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, ApJL, 882, L24
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021a, ApJL, 913, L7
Abbott, R., Abbott, T. D., Acernese, F., et al. 2023c, PhRvX, 13, 041039
Abbott, R., Abbott, T. D., Acernese, F., et al. 2023a, PhRvX, 13, 011048
Abbott, R., Abe, H., Acernese, F., et al. 2023b, ApJ, 949, 76
Abbott, R., Abe, H., Acernese, F., et al. 2023d, ApJS, 267, 29
Acernese, F., Agathos, M., Agatsuma, K., et al. 2014, CQGra, 32, 024001
Akutsu, T., Ando, M., Arai, K., et al. 2021, PTEP, 2021, 05A101
Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, ApJS, 241, 27
Aubin, F., Brighenti, F., Chierici, R., et al. 2021, CQGra, 38, 095004
Baxter, E. J., Croon, D., McDermott, S. D., & Sakstein, J. 2021, ApJL, 916, L16
Bayarri, M. J., & Castellanos, M. E. 2007, StaSc, 22, 322
Broekgaarden, F. S., Stevenson, S., & Thrane, E. 2022, ApJ, 938, 45
Callister, T. A., & Farr, W. M. 2023, arXiv:2302.07289
Callister, T. A., Haster, C. J., Ng, K. K. Y., Vitale, S., & Farr, W. M. 2021, ApJL, 922, L5
Cannon, K., Caudill, S., Chan, C., et al. 2021, SoftX, 14, 100680
Chernoff, D. F., & Finn, L. S. 1993, ApJL, 411, L5
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
Drago, M., Klimenko, S., Lazzaro, C., et al. 2021, SoftX, 14, 100678
Edelman, B., Doctor, Z., Godfrey, J., & Farr, B. 2022a, ApJ, 924, 101
Edelman, B., Farr, B., & Doctor, Z. 2023, ApJ, 946, 16
Essick, R., & Farr, W. 2022, arXiv:2204.00461
Essick, R., & Fishbach, M. 2021, On Estimating Rates from Monte-Carlo Integrals over Injection Sets, https://dcc.ligo.org/T2000100
Ezquiaga, J. M., & Holz, D. E. 2021, ApJL, 909, L23
Ezquiaga, J. M., & Holz, D. E. 2022, PhRvL, 129, 061102
Farag, E., Renzo, M., Farmer, R., Chidester, M. T., & Timmes, F. X. 2022, ApJ, 937, 112
Farah, A., Fishbach, M., Essick, R., Holz, D. E., & Galaudage, S. 2022, ApJ, 931, 108
Farah, A. M., Edelman, B., Zevin, M., et al. 2023, ApJ, 955, 107
Farmer, R., Renzo, M., de Mink, S., Fishbach, M., & Justham, S. 2020, ApJL, 902, L36
Farmer, R., Renzo, M., de Mink, S. E., Marchant, P., & Justham, S. 2019, ApJ, 887, 53
Farr, W. M., Fishbach, M., Ye, J., & Holz, D. E. 2019, ApJL, 883, L42
Fishbach, M., Farr, W. M., & Holz, D. E. 2020, ApJL, 891, L31
Fishbach, M., & Holz, D. E. 2017, ApJL, 851, L25
Fishbach, M., & Holz, D. E. 2020, ApJL, 891, L27
Fishbach, M., Holz, D. E., & Farr, W. M. 2018, ApJL, 863, L41
Foreman-Mackey, D. 2016, JOSS, 1, 24
Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, ApJ, 749, 91
Gallegos-Garcia, M., Berry, C. P. L., & Kalogera, V. 2022, ApJ, 955, 133
Gelman, A. 2006, BayAn, 1, 515
Godfrey, J., Edelman, B., & Farr, B. 2023, arXiv:2304.01288
Grudić, M. Y., Offner, S. S. R., Guszejnov, D., Faucher-Giguère, C. A., & Hopkins, P. F. 2023, OJAp, 6, 48
Heger, A., & Woosley, S. E. 2002, ApJ, 567, 532
Hu, R. C., Zhu, J. P., Qin, Y., et al. 2022, ApJ, 928, 163
Kalogera, V. 1996, ApJ, 471, 352
Kiendrebeogo, R. W., Farah, A. M., Foley, E. M., et al. 2023, ApJ, 958, 158
Kovetz, E. D., Cholis, I., Breysse, P. C., & Kamionkowski, M. 2017, PhRvD, 95, 103010
Laplace, E., Justham, S., Renzo, M., et al. 2021, A&A, 656, A58
LIGO Scientific CollaborationVirgo CollaborationKAGRA Collaboration 2021, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run ? O3 Search Sensitivity Estimates, v2, Zenodo, doi:10.5281/zenodo.5546676
Loredo, T. J. 2004, in AIP Conf. Proc. 735, 24th Int. Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering (Melville, NY: AIP), 195
Loredo, T. J. 2013, in Astrostatistical Challenges for the New Astronomy, ed. J. M. Hilbe (Berlin: Springer), 14
Mandel, I., & Farmer, A. 2022, PhR, 955, 1
Mandel, I., Farr, W. M., & Gair, J. R. 2019, MNRAS, 486, 1086
Mandel, I., Müller, B., Riley, J., et al. 2021, MNRAS, 500, 1380
Mapelli, M. 2020, FrASS, 7, 38

Messenger, C., & Read, J. 2012, PhRvL, 108, 091101

Mould, M., Gerosa, D., Broekgaarden, F. S., & Steinle, N. 2022, MNRAS, 517, 2738

Oh, M., Fishbach, M., Kimball, C., Kalogera, V., & Ye, C. 2023, ApJ, 953, 152

Olejak, A., & Belczynski, K. 2021, ApJL, 921, L2

Ray, A., Magaña Hernandez, I., Mohite, S., Creighton, J., & Kapadia, S. 2023, ApJ, 957, 37

Sadiq, J., Dent, T., & Gieles, M. 2023, arXiv:2307.12092

Sadiq, J., Dent, T., & Wysocki, D. 2022, PhRvD, 105, 123014

Schneider, F. R. N., Podsiadlowski, P., & Laplace, E. 2023, ApJL, 950, L9

Sinharay, S., & Stern, H. S. 2003, J. Stat. Plan. Inference, 111, 209

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

Stevenson, S., Ohme, F., & Fairhurst, S. 2015, ApJ, 810, 58

Stevenson, S., Sampson, M., Powell, J., et al. 2019, ApJ, 882, 121

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

Talbot, C., & Thrane, E. 2018, ApJ, 856, 173

Taylor, S. R., Gair, J. R., & Mandel, I. 2012, PhRvD, 85, 023535

Tiwari, V. 2021, CQGra, 38, 155007

Tiwari, V. 2023, MNRAS, 527, 298

Usman, S. A., Nitz, A. H., Harry, I. W., et al. 2016, CQGra, 33, 215004

van Son, L. A. C., de Mink, S. E., Renzo, M., et al. 2022a, ApJ, 940, 184

van Son, L. A. C., Mink, S. E. d., Callister, T., et al. 2022b, ApJ, 931, 17

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261

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

Zevin, M., & Bavera, S. S. 2022, ApJ, 933, 86

Zevin, M., Pankow, C., Rodriguez, C. L., et al. 2017, ApJ, 846, 82