

Striking a Chord with Spectral Sirens: Multiple Features in the Compact Binary Population Correlate with H_0

Utkarsh Mali^{1,2}  and Reed Essick^{1,2,3} 

¹ Canadian Institute for Theoretical Astrophysics, 60 St. George Street, Toronto, ON M5S 3H8, Canada; utkarsh.mali@utoronto.ca

² Department of Physics, University of Toronto, 60 St. George Street, Toronto, ON M5S 1A7, Canada

³ David A. Dunlap Department of Astronomy, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada

Received 2024 October 11; revised 2024 December 7; accepted 2024 December 10; published 2025 February 5

Abstract

Spectral siren measurements of the Hubble constant (H_0) rely on correlations between observed detector-frame masses and luminosity distances. Features in the source-frame mass distribution can induce these correlations. It is crucial, then, to understand (i) which features in the source-frame mass distribution are robust against model (re) parameterization, (ii) which features carry the most information about H_0 , and (iii) whether distinct features independently correlate with cosmological parameters. We study these questions using real gravitational-wave observations from the LIGO-Virgo-KAGRA Collaborations' third observing run. Although constraints on H_0 are weak, we find that current data reveals several prominent features in the mass distribution, including peaks in the binary black hole source-frame mass distribution near $\sim 9M_\odot$ and $\sim 32M_\odot$ and a roll-off at masses above $\sim 46M_\odot$. For the first time using real data, we show that all of these features carry cosmological information and that the peak near $\sim 32M_\odot$ consistently correlates with H_0 most strongly. Introducing model-independent summary statistics, we show that these statistics independently correlate with H_0 , exactly what is required to limit systematics within future spectral siren measurements from the (expected) astrophysical evolution of the mass distribution.

Unified Astronomy Thesaurus concepts: [Compact objects \(288\)](#); [Gravitational wave astronomy \(675\)](#); [Gravitational waves \(678\)](#); [Black holes \(162\)](#); [Neutron stars \(1108\)](#)

1. Introduction

The observation of gravitational waves (GWs) from compact binary coalescences (CBCs) observed with the advanced LIGO, Virgo, and KAGRA (LVK) detectors (F. Acernese et al. 2014; J. Aasi et al. 2015; T. Akutsu et al. 2020) provide a new window onto many astrophysical and cosmological phenomena (B. Abbott et al. 2016a, 2016b, 2019; R. Abbott et al. 2021b, 2023a, 2024). Many authors have proposed various methods to use CBCs as tracers of the Hubble relation. While the specific approaches vary (B. P. Abbott et al. 2017; M. Soares-Santos et al. 2019; J. M. Ezquiaga & D. E. Holz 2022), they all involve obtaining simultaneous estimates of both the luminosity distance (D_L) and redshift (z) for a set of sources. These estimates are then used to fit the Hubble relation $H(z)$, allowing us to extract key cosmological parameters.

Several electromagnetic (EM) approaches operate in this way; these include Hubble's original observations (E. Hubble 1929), recent catalogs of cepheids (A. G. Riess et al. 2020), the tip of the red giant branch (M. G. Lee et al. 1993; W. L. Freedman et al. 2020), and the J-region asymptotic giant branch (W. L. Freedman & B. F. Madore 2023). Often, with EM measurements, it is possible to precisely measure z , but more difficult to estimate D_L . Indeed, much of the discussion in the literature focuses on different ways to calibrate the local distance ladder and thereby improve estimates of D_L (W. L. Freedman et al. 2024).

Conversely, GW observations of CBCs can directly constrain D_L independently of the distance ladder and without relying on other observations (B. F. Schutz 1986). However, it

is much more difficult to reliably estimate z for GW sources. In general relativity, vacuum solutions to the Einstein field equations do not automatically contain a fixed scale (like the known rest-frame frequencies of atomic and molecular lines in EM observations) that can be used to measure z . As such, it is relatively easy to measure D_L with GW observations, and many approaches in GW cosmology focus on different ways of obtaining estimates of z .

Several of the most common approaches rely on EM data to obtain z . *Bright sirens* use observations of EM counterparts for individual events to measure z through an association with a host galaxy (D. E. Holz & S. A. Hughes 2005; N. Dalal et al. 2006; S. Nissanka et al. 2010, 2013; B. P. Abbott et al. 2017). *Dark sirens*, instead, do not rely on the identification of EM counterparts for individual events but rather probabilistically associate CBCs with catalogs of potential host galaxies (W. Del Pozzo 2012; M. Fishbach et al. 2019; M. Soares-Santos et al. 2019; R. Gray et al. 2020, 2022; B. Abbott et al. 2021a; S. Mukherjee et al. 2021, 2024; J. R. Gair et al. 2023; S. Mastrogiovanni et al. 2023; A. G. Hanselman et al. 2025). In both cases, EM observations effectively serve as a separate source of information about z for individual events.

However, by considering a catalog of CBCs, information about z can be inferred from GW data alone. That is, GW cosmology does not need to rely on EM data. Several approaches have been proposed (see, e.g., D. Chatterjee et al. 2021; J. M. Ezquiaga & D. E. Holz 2022; Y.-J. Li et al. 2024). In general, they rely on identifying an observable feature of individual CBC systems that correlates with the source-frame mass (m_s). Measurement of that parameter then provides information about m_s , and the GW data directly constrains the detector-frame (or redshifted) mass: $m_d = (1 + z)m_s$. Joint constraints on m_d and m_s thereby provide an estimate for z .



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

We focus on *spectral sirens*, in which features in the distribution of source-frame masses inferred from a catalog of CBCs provide additional information about m_s for individual events. This approach has been studied in several contexts, including for binary neutron star (BNS) and binary black hole (BBH) systems with current and proposed GW detectors (B. F. Schutz 1986; D. F. Chernoff & L. S. Finn 1993; C. Messenger & J. Read 2012; S. R. Taylor et al. 2012; H.-Y. Chen et al. 2018; W. M. Farr et al. 2019; R. Gray et al. 2020; D. Chatterjee et al. 2021; S. Mastrogiovanni et al. 2021, 2022; Z.-Q. You et al. 2021; J. M. Ezquiaga & D. E. Holz 2022; K. Leyde et al. 2022; M. Mancarella et al. 2022; C. Karathanasis et al. 2023; S. Mastrogiovanni et al. 2023; I. M. Hernandez & A. Ray 2024; G. Pierra et al. 2024; A. M. Farah et al. 2025).

Typically, authors assume the presence of a feature in the source-frame mass distribution, like a narrow peak (S. R. Taylor et al. 2012) or a rapid falloff (W. M. Farr et al. 2019; M. Mapelli 2020), which then induces a telltale pattern in the joint distribution of detector-frame masses and luminosity distance. See, e.g., J. M. Ezquiaga & D. E. Holz (2022) for a review.

The efficacy of spectral sirens has been simulated by many authors (S. R. Taylor et al. 2012; J. M. Ezquiaga & D. E. Holz 2022; A. Ray et al. 2024; A. M. Farah et al. 2025), but several questions remain. While the method may work in simulated universes, we wish to study its prospects with real observations. As such, we study

1. whether there are useful features in the source-frame mass distribution consistently inferred with different mass models,
2. which of these features (if any) carry the most information about $H(z)$, and
3. whether multiple features correlate equally and independently with $H(z)$ or if most of the information is associated with a single feature.

Since a potential weakness of spectral sirens is the unknown astrophysical evolution of the source-frame mass with time, the last point is of particular interest. That is, if the distribution of m_s is not independent of z , then correlations between m_s and z may not be due to the Universe's expansion, but instead simply associated with stellar astrophysics (e.g., changes in binary evolution). However, if the source-frame mass distribution contains multiple features across mass scales, it is improbable that the same astrophysical processes would affect all of them identically, whereas cosmological effects would. Put differently, stellar evolution operates differently at different mass scales but cosmology redshifts all GWs in the same way. Therefore, if multiple features in the source-frame mass distribution separately correlate with H_0 , then it is likely we will be able to break the degeneracy between astrophysical evolution of the source-frame mass distribution and $H(z)$. Our work shows that the source-frame mass distribution inferred from real (observed) GW data has multiple such features.

We begin by outlining our methodology and the observations used in Section 2. We examine the behavior of several different models of the source-frame mass distribution, each of which is described in Section 2.1 (see also Appendix B). Section 3 briefly discusses conclusions from the joint posteriors obtained by simultaneously inferring a Λ CDM cosmology and our mass distribution before Section 4 examines how different

features within the mass distribution encode information about H_0 . We identify which parameters correlate most strongly with H_0 , explore which types of behaviors lead to stronger correlations, and introduce several model-independent summary statistics. Interestingly, the summary statistics can encode cosmological information even better than individual model parameters. We show that current data supports the presence of local overdensities (“peaks”) in the source-frame mass distribution near $\sim 9M_\odot$ and $\sim 32M_\odot$ along with a “roll-off” near $\sim 46M_\odot$ and that all of these features correlate with H_0 . Section 5 then investigates whether these features are independent. We show that their correlations are predominantly driven by their independent correlations with H_0 , which is essential for a robust spectral siren constraint. We conclude in Section 6.

2. Methodology

We construct a hierarchical model and use 63 confidently detected CBCs from the LVK's third observing run (O3) to simultaneously infer the component mass distribution and cosmological parameters. We assume CBCs follow an inhomogeneous Poisson process, marginalizing over the overall rate of mergers. The likelihood of observed data (D_i) for each event (i), given parameters that describe the CBC merger density and the Hubble expansion $H(z)$ (and other population parameters, Λ) is then

$$p(\{D_i\}|\Lambda) \propto \frac{1}{\mathcal{E}^N} \prod_{i=1}^N \mathcal{Z}_i, \quad (1)$$

where we have defined the single-event evidence \mathcal{Z}_i

$$\mathcal{Z}_i(\Lambda) = p(D_i|\Lambda) = \int p(D_i|\theta) p(\theta|\Lambda) d\theta \quad (2)$$

and the detection probability \mathcal{E}

$$\mathcal{E}(\Lambda) = P(\text{det}|\Lambda) = \int P(\text{det}|\theta) p(\theta|\Lambda) d\theta. \quad (3)$$

Within these integrals, θ represents the single-event parameters, like component masses and spins, $p(D_i|\theta)$ is the likelihood of obtaining D_i given a signal described by θ , and $P(\text{det}|\theta)$ is the probability that a signal described by θ would be detected (marginalized over noise realizations; see ; R. Essick 2023; R. Essick & M. Fishbach 2024). See J. Skilling (2004), I. Mandel (2010), E. Thrane & C. Talbot (2019), I. Mandel et al. (2019) and S. Vitale et al. (2021) for reviews.

Our set of 63 confident events from O3 was obtained by selecting those events for which at least one search (either cWB (S. Klimenko et al. 2021), GstLAL (K. Cannon et al. 2021), MBTA (F. Aubin et al. 2021), and one of two PyCBC searches (A. Nitz et al. 2024)) reported a false alarm rate ≤ 1 yr. We only use events from O3 because real search sensitivity estimates (R. Abbott et al. 2021c) are not publicly available for other observing runs, and O3 contains the vast majority of the publicly available surveyed volume-time (and detected events).

Additionally, we consider events across the entire mass spectrum, including BNS, neutron star-black hole (NSBH), and BBH coalescences. Previous analyses have focused only on BBHs, applying ad hoc cuts based on the secondary mass inferred assuming a reference cosmology. De facto, they assume neutron stars (NSs) are uninformative a priori. Although Section 3 finds that much of the information about

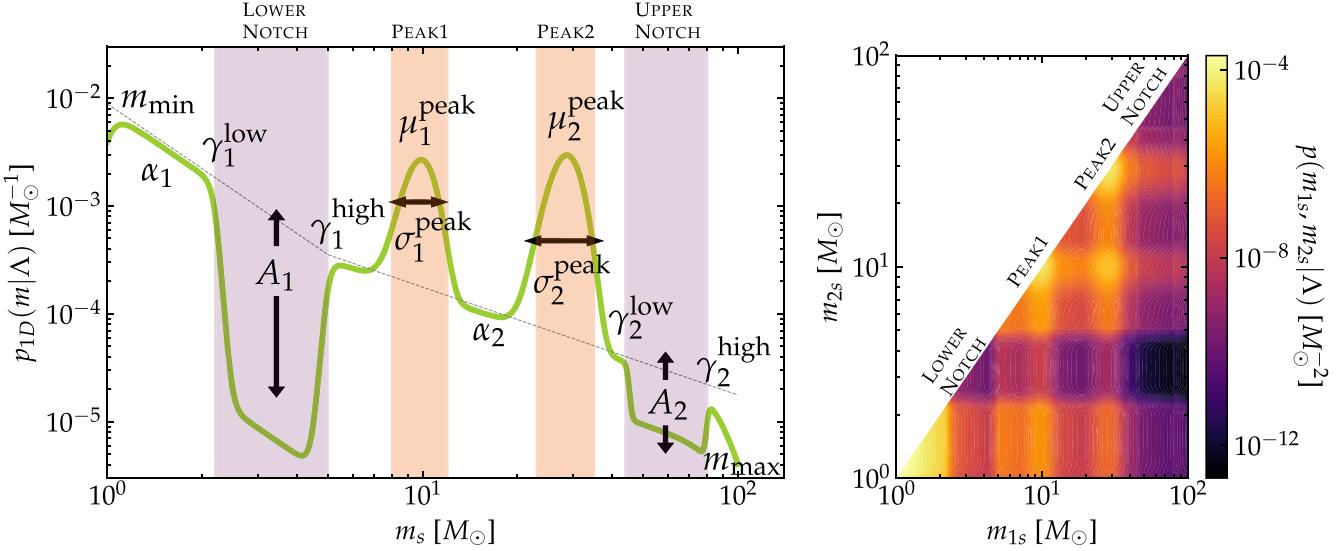


Figure 1. (Left) Schematic representation of $p_{1D}(m|\Lambda)$ (Equation (B11)) and associated hyperparameters. The model (green solid) is based on a broken power law (black dotted) with roll-offs at both high (m_{\max}) and low masses (m_{\min}). Additional Butterworth notch filters (purple) and Gaussian peaks (orange) are included as multiplicative factors. See Appendix B for details. (right) The associated joint distribution $p(m_{1s}, m_{2s}|\Lambda)$ (Equation (5)). Note that features in p_{1D} appear in both m_{1s} and m_{2s} (A. M. Farah et al. 2024).

cosmology from the current catalog is indeed associated with features at high masses, we also consider low-mass systems.

For our cosmological model, we assume flat Λ CDM. We fit for H_0 and the present-day matter density (Ω_m). The present-day radiation density (Ω_r) is fixed at 0.001. We assume the closure condition so that the present-day dark-energy density is $\Omega_\Lambda = 1 - \Omega_m - \Omega_r$. Appendix C lists all priors assumed within our inference.

2.1. Population Model and Mass Distributions

Within our analysis, the distribution of compact binaries is of primary importance. We focus on the source-frame mass distribution, making simple assumptions for the spin and redshift (z) distributions. We assume fixed distributions over component spins (uniform in magnitude and isotropic in orientation) and that binaries' redshift follows the local star formation rate ($\Phi(z)$) (see Appendix B).

$$p(z|\Lambda) \propto \frac{dV_c}{dz} \frac{1}{1+z} \Phi(z). \quad (4)$$

Here, V_c is the comoving volume and the factor, and $(1+z)^{-1}$ accounts for cosmological time-dilation between the source and detector frames. Note that $p(z|\Lambda)$ implicitly depends on $H(z)$ through V_c .

We extend a phenomenological source-frame mass distribution first introduced in M. Fishbach et al. (2020) and later used in A. Farah et al. (2022) and R. Abbott et al. (2023b). Specifically, we model the joint distribution of source-frame component masses (m_{1s} and m_{2s}) as

$$p(m_{1s}, m_{2s}|\Lambda) \propto p_{1D}(m_{1s}|\Lambda) p_{1D}(m_{2s}|\Lambda) \times f(m_{1s}, m_{2s}; \Lambda) \Theta(m_{1s} \geq m_{2s}), \quad (5)$$

where p_{1D} is a one-dimensional distribution over mass, ⁴ f is a pairing function that influences how often different component masses form binaries, and Θ is an indicator function that

⁴ Importantly, $p_{1D}(m|\Lambda)$ should not be confused with the marginal distributions over the primary or secondary masses $p(m_1|\Lambda)$, $p(m_2|\Lambda)$.

enforces our labeling convention: $m_{1s} \geq m_{2s}$. We truncate p_{1D} below $1 M_\odot$, but we do consider high-mass binaries ($m_{1s} > 100 M_\odot$) to which the LVK is currently sensitive. Figure 1 shows an example of $p_{1D}(m|\Lambda)$ and the corresponding $p(m_{1s}, m_{2s}|\Lambda)$.

Details of the exact functional forms and priors assumed within the inference are provided in Appendix B, but briefly, we model p_{1D} as a broken power law with roll-offs at both high and low masses, which is further augmented by a set of multiplicative filters that either remove notches or add peaks to the mass distribution. Specifically, we consider up to two Gaussian peaks and two Butterworth notch filters. These features give the mass model considerable flexibility while still providing convenient ways to downselect or “turn off” specific features. Sections 2.1.1–2.1.3 describe our models in more detail.

Before performing the integrals in Equations (2) and (3), we transform our model from source-frame mass and redshift to detector-frame mass ($m_d = m_s(1+z)$) and luminosity distance ($D_L(z)$). This allows us to conveniently incorporate the cosmological dependence of these transformations within our analysis and approximate our integrals (Equations (2) and (3)) as Monte Carlo samples over parameters that are directly measured (m_{1d} , m_{2d} , and D_L)

$$\begin{aligned} p(\theta|\Lambda) &= p(m_{1d}, m_{2d}, D_L, s_1, s_2|\Lambda) \\ &= p(m_{1s}) = \frac{m_{1d}}{1+z}, m_{2s} = \frac{m_{2d}}{1+z} \Big| \Lambda(1+z)^{-2} \\ &\quad \times p(z = z(D_L)|\Lambda) \Big| \frac{dD_L}{dz} \Big|^{-1} \\ &\quad \times p(s_1, s_2|\Lambda). \end{aligned} \quad (6)$$

Because we focus on the source-frame mass distribution, we assumed fixed distributions for the redshift and spins. That is, we only fit the hyperparameters associated with the mass distribution and the cosmological model. See Appendix B for more discussion.

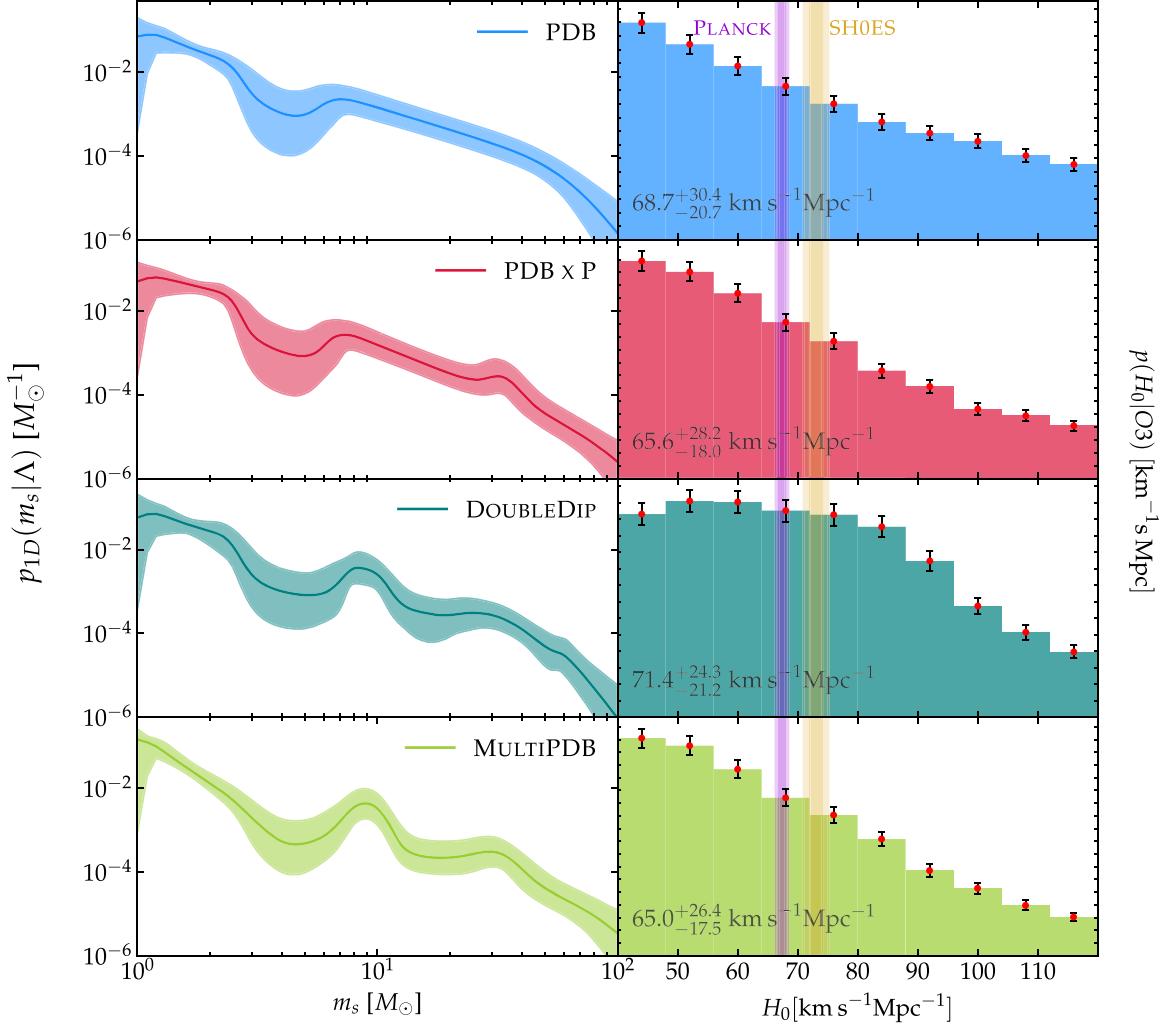


Figure 2. (Left) Posterior medians and 90% symmetric credible regions for $p_{1D}(m_s|\Lambda)$ as a function of mass conditioned on the O3 catalog for (top to bottom) PDB (Section 2.1.1), PDB \times P (Section 2.1.2), DOUBLEDIP and MULTIPDB (Section 2.1.3). In general, p_{1D} decreases as a function of mass, but models that allow for additional flexibility often find evidence for peaks at $\sim 9M_\odot$ and $\sim 32M_\odot$. (right) Corresponding posterior distributions (linear scale) for the Hubble parameter (H_0) along with 2σ error bars from the finite number of samples used. Median values and 90% symmetric credible regions are shown in each panel. We also show the 1σ and 2σ constraints from the Planck and SH0ES collaborations as vertical bars (N. Aghanim et al. 2020; A. G. Riess et al. 2022). Although mostly uninformative, the data primarily disfavors large values of H_0 . See Table 3 for priors.

We consider several variations of $p_{1D}(m|\Lambda)$, each with different subsets of features “switched on.” In order of increasing complexity, the following sections explain the motivation for each. Figure 2 shows their corresponding posterior distributions conditioned on the LVK’s O3 catalog.

2.1.1. Base Distribution (Power Law + Dip + Break)

First introduced in M. Fishbach et al. (2020) and later used in A. Farah et al. (2022) and R. Abbott et al. (2023b), the Power law + Dip + Break (PDB; top row of Figure 2) mass distribution is the simplest p_{1D} model we consider. This model includes a broken power law (indexes α_1 and α_2 in Figure 1) with high- and low-mass roll-offs (m_{\max} and m_{\min}) along with a single Butterworth notch ($\gamma_1^{\text{low}}, \gamma_1^{\text{high}}$, and A_1) to account for the dearth of compact objects observed between 3 and $5M_\odot$. We include this model to make comparisons with previous results and to provide a baseline for how much cosmological information is contained in a source-frame mass distribution that does not have any pronounced peaks.

2.1.2. Distributions with a Single Peak (PDB \times P)

We also consider a model with a single additional peak (PDB \times P: μ_2^{peak} , and σ_2^{peak} in Figure 1; second row of Figure 2). There is strong evidence that there is a local overdensity in the mass distribution around $30M_\odot$ compared to a single power law (R. Abbott et al. 2023b; A. M. Farah et al. 2023; T. A. Callister & W. M. Farr 2024), and PDB \times P models this with a Gaussian peak (see Appendix A). We generally find consistent behavior between PDB \times P and other models that include a single peak (R. Abbott et al. 2021d, 2023b).

Additionally, while our implementation includes additional features as multiplicative factors, it is common to include such features as additional components in a mixture model (B. Abbott et al. 2019; R. Abbott et al. 2021b; M. Zevin et al. 2021; J. Godfrey et al. 2024). While these models are not exactly equivalent, both approaches produce similar behavior. We demonstrate this in Appendix A by comparing PDB \times P to a mixture model of PDB with a separate Gaussian (PDB+P).

Table 1Pearson Correlation Coefficients (r) between Features in the Mass Distribution and H_0 from Joint Posterior Distributions Conditioned on CBCs from O3

	γ_1^{high}	μ_1^{peak}	$\hat{\mu}_1^{7:11}$	$\hat{\sigma}_1^{7:11}$	μ_2^{peak}	$(\mu_2^{\text{peak}} \sigma_2^{\text{peak}} < 8M_\odot)$	$\hat{\mu}_2^{25:40}$	$\hat{\sigma}_2^{25:40}$	m_{max}	m_{99}
PDB	-0.19	...	-0.16	+0.17	-0.38	-0.46	-0.38	-0.33
PDB \times P	-0.2	...	-0.24	+0.14	-0.48	-0.73	-0.68	-0.46	-0.04	-0.29
DOUBLEDIP	-0.06	...	-0.48	-0.37	-0.54	-0.09	-0.15	-0.26
MULTIPDB	+0.08	-0.49	-0.65	-0.17	-0.4	-0.66	-0.72	-0.37	-0.03	-0.23

Note. Individual model parameters are described in Section 2.1 and Figure 1. Model-independent summary statistics are described in Section 4.2.

2.1.3. Distributions with Multiple Peaks (DOUBLEDIP and MULTIPDB)

In addition to an overdensity at $\sim 30M_\odot$, there is growing evidence in favor of another feature near $\sim 10M_\odot$ (V. Tiwari & S. Fairhurst 2021; B. Edelman et al. 2023; A. M. Farah et al. 2023; T. A. Callister & W. M. Farr 2024; J. Godfrey et al. 2024; A. Ray et al. 2024). We introduce two models that try to capture this behavior.

First, DOUBLEDIP (third row in Figure 2) introduces a second Butterworth notch filter ($\gamma_2^{\text{low}}, \gamma_2^{\text{high}}$, and A_2 in Figure 1) because two peaks (local maxima) can be achieved by introducing a single notch (local minimum). Generally, the additional notch produces local maxima in p_{1D} at the expected mass scales a posteriori, although the posterior distribution for γ_2^{high} is bimodal and also supports a small overdensity at higher masses (i.e., it snaps to the component masses of GW190521; R. Abbott et al. 2020).

MULTIPDB (bottom row of Figure 2) builds upon PDB by considering a model that explicitly adds two additional Gaussian peaks: $(\mu_1^{\text{peak}}, \sigma_1^{\text{peak}})$ and $(\mu_2^{\text{peak}}, \sigma_2^{\text{peak}})$ in Figure 1. MULTIPDB consistently finds overdensities in p_{1D} at both $\sim 9M_\odot$ and $\sim 32M_\odot$.

3. Joint Posteriors Conditioned on O3

For each mass model, we sample from the joint posterior distribution for both the parameters describing the mass distribution and the Hubble relation. Figure 2 shows the median and 90% symmetric credible regions for p_{1D} as a function of mass along with the marginal posterior distributions for H_0 .

Generally, we find posteriors with the expected behavior that are consistent with previous results (R. Abbott et al. 2021d). All models prefer a dip between $3.0M_\odot$ and $6.4M_\odot$. All models also find a local maximum in p_{1D} around $\sim 9M_\odot$, although this manifests as a more pronounced “peak” in models which support multiple peaks (DOUBLEDIP and MULTIPDB). All models also consistently find that the power law steepens at higher masses through a roll-off above $46M_\odot$. Finally, all models that can support a local maximum near $\sim 32M_\odot$ find one a posteriori. Models with Gaussian peaks call this feature μ_2^{peak} , and it corresponds to γ_2^{high} in DOUBLEDIP (the upper edge of the second notch).

Hyperpriors for each mass model are shown in Tables 3 and 4, symmetric posterior credible regions for all other values included in our fits can be found in Table 5. We report their median values along with their 90% symmetric credible regions.

Additionally, there are qualitative similarities in the marginal posteriors for H_0 . All models primarily disfavor large values of H_0 a posteriori, which would correspond to large detector-frame masses given a fixed source-frame mass distribution. The H_0 posteriors obtained with PDB, PDB \times P, and MULTIPDB closely resemble the results from R. Abbott et al. (2021d), which assumed a source-frame primary-mass distribution consisting of a mixture of a single (unbroken) power law and Gaussian peak. Interestingly, the posterior obtained with

DOUBLEDIP has a local maximum near the values reported by the Planck and SH0ES collaborations (N. Aghanim et al. 2020; A. G. Riess et al. 2021). However, the posterior is very wide and remains consistent with the other mass models.

It is clear, then, that current catalogs of CBCs are not competitive with other estimates of H_0 . However, it is also clear that some cosmological information is encoded within the inference. We now examine exactly how information about H_0 manifests through correlations with different features in the source-frame mass distribution.

4. Which Features Carry Cosmological Information?

Our mass models support a wide range of features, including peaks, notches, and roll-offs. Upon first inspection, it may not be immediately clear which of these features would be most useful when attempting to infer the Hubble relation. We attempt to identify those features by examining the Pearson correlation coefficients (r) between H_0 and features in p_{1D} a posteriori (i.e., correlations induced by conditioning on the observed data). Table 1 reports the correlation coefficients between H_0 and several statistics representing these features.

We currently expect high-mass (BBH) features to carry more cosmological information than BNS masses. This is because BBH mergers are detectable to much larger redshifts than BNS.⁵ As such, more BBHs are detected within current catalogs, and cosmological effects are more apparent within BBHs. We begin by examining the correlations between H_0 and a few parameters of each mass distribution in Section 4.1.

As we will see, it quickly becomes apparent that relying on specific parameters of individual models may be difficult to scale, as models with different parametrizations can nevertheless capture the same behavior. As such, we additionally consider several model-independent summary statistics derived from p_{1D} and study how they correlate with H_0 in Section 4.2. These statistics can be extended to any mass model, even those without concise functional forms (B. Edelman et al. 2023; T. A. Callister & W. M. Farr 2024; A. Ray et al. 2024; A. M. Farah et al. 2025).

4.1. Parametric Descriptions of the Mass Distribution

Examining the behavior of the high-mass end of the distributions in Figure 2, we might expect either the shelf created by the roll-off around $46M_\odot$ or local maximum at $\sim 32M_\odot$ to carry cosmological information. In general, H_0 may correlate with many parameters, such as m_{max} (roll-off at high masses), μ_1^{peak} (overdensity near $9M_\odot$), and μ_2^{peak} (overdensity near $32M_\odot$). Some of these are listed in Table 1.

⁵ J. M. Ezquiaga & D. E. Holz (2022) claim that the lower mass gap will eventually dominate the constraint with next-generation GW detectors given the expected increase in the BNS detection rate.

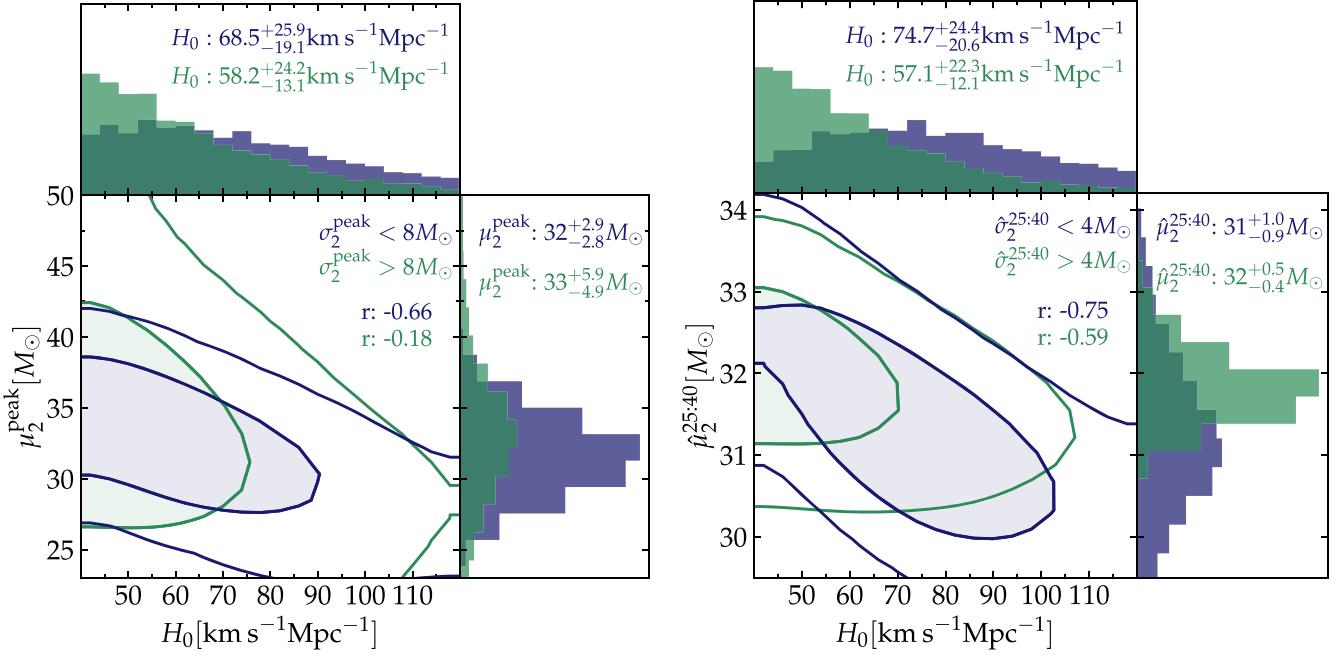


Figure 3. (Left) Joint and marginal posterior distributions for μ_2^{peak} and H_0 conditioned on the presence of wide (green, $\sigma_2^{\text{peak}} \geq 8 M_\odot$) and narrow (blue, $\sigma_2^{\text{peak}} \leq 8 M_\odot$) peaks with MULTIPDB. Contours in the joint distributions show the 50% and 90% highest-probability-density credible regions. Median and 90% symmetric credible regions are shown in the marginal distributions, and the Pearson correlation coefficients (r) are shown in the joint distribution. (right) Analogous model-independent summary statistics derived from p_{1D} ($\hat{\mu}_2^{25:40}$ and $\hat{\sigma}_2^{25:40}$) for MULTIPDB. We apply different thresholds on σ_2^{peak} and $\hat{\sigma}_2^{25:40}$ based on their one-dimensional posterior medians.

To begin, we examine the roll-off at high masses, which was previously identified as a useful feature in W. M. Farr et al. (2019). Even if the true mass distribution does not have any peaks, like PDB, the roll-off at high masses still carries cosmological information. That is, we do not need a peak to constrain the Hubble relation.

However, Table 1 also shows that when a model supports a peak near $32 M_\odot$, the location of that peak always correlates more strongly with H_0 than m_{max} . This is also apparent in Figures 5 and 13 of R. Abbott et al. (2021d). As such, while the presence of a peak is not necessary for spectral siren cosmology, it is helpful. Additionally, we consider which aspects of the peak help it carry cosmological information in Figure 3 using MULTIPDB. In line with predictions from S. R. Taylor et al. (2012), we find that wide peaks ($\sigma_2^{\text{peak}} \geq 8 M_\odot$) yield small correlation coefficients a posteriori, even smaller than what is observed between H_0 and m_{max} . However, relatively sharp peaks ($\leq 8 M_\odot$) yield strong correlations between H_0 and μ_2^{peak} . Interestingly, conditioning on a sharp peak significantly shifts the maximum of the marginal posterior for H_0 much closer to $70 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

Finally, Table 1 does not report a correlation between μ_2^{peak} and H_0 for DOUBLEDIP because that model does not contain that parameter. Interestingly, DOUBLEDIP does have a few parameters that moderately correlate with H_0 . The lower and upper edges of the high-mass dip (γ_2^{low} and γ_2^{high}) both correlate with H_0 with $r = -0.13$ and 0.18 , respectively.⁶ Additionally, m_{max} and η_{max} , which trace the location and steepness of the

falloff at masses above $\sim 40 M_\odot$, also correlate with H_0 ($r = -0.15$ and 0.27 , respectively). However, none of these parameters correlate as strongly with H_0 as the parameters describing peaks in other models (e.g., μ_2^{peak} in MULTIPDB). That is, even though DOUBLEDIP can reproduce qualitatively similar features as the other models in Figure 2, it lacks a single parameter that concisely captures the overall behavior of p_{1D} near $\sim 30 M_\odot$. This makes it clear that relying on correlations between individual parameters and H_0 may not be a robust way to describe the information contained within p_{1D} and suggests the need for model-independent summary statistics, which we explore in Section 4.2.

4.2. Model-independent Summary Statistics Derived from the Mass Distribution

So far, our analysis has focused on specific parameters in individual models. However, this may not generalize well, particularly for models that do not have concise parametric representations (T. A. Callister & W. M. Farr 2024; A. Ray et al. 2024; A. M. Farah et al. 2025). Instead, we now characterize general features in the mass distribution with model-independent summary statistics and then consider their correlations with H_0 .

We introduce several such statistics. Each is derived from p_{1D} , and therefore not directly applicable to the joint distribution $p(m_{1s}, m_{2s}|\Lambda)$ that may be constrained by other analyses. However, we believe they nevertheless capture important behavior and can be generalized to the joint distribution straightforwardly, particularly as GW catalogs seem to suggest a relatively strong preference for equal-mass binaries within the astrophysical distribution (M. Fishbach & D. E. Holz 2020; R. Abbott et al. 2023b).

First, we compute moments concerning p_{1D} over restricted ranges of mass. Specifically, we consider the mean and

⁶ Simpson's paradox (E. H. Simpson 1951) indicates that correlations between variates observed in the whole population can change when subsets of samples are examined separately. The γ_2^{high} correlations with H_0 demonstrate this behavior (since γ_2^{high} is bimodal). Both subpopulations independently anticorrelate with H_0 (we have confirmed this). However, when combined, their correlations appear positively correlated.

variance within a range of masses:

$$\hat{\mu}(X, Y) = \int_X^Y mq(m; X, Y)dm, \quad (7)$$

$$\hat{\sigma}^2(X; Y) = \int_X^Y [m - \hat{\mu}(X, Y)]^2 q(m; X, Y)dm, \quad (8)$$

where

$$q(m; X, Y) \equiv \frac{p_{1D}(m|\Lambda)}{\int_X^Y p_{1D}(m|\Lambda)dm} \Theta(X \leq m \leq Y). \quad (9)$$

Table 1 reports the correlations between H_0 and

$$\hat{\mu}_1^{7:11} \equiv \hat{\mu}(7 M_\odot, 11 M_\odot), \quad (10)$$

$$\hat{\sigma}_1^{7:11} \equiv \hat{\sigma}(7 M_\odot, 11 M_\odot), \quad (11)$$

as well as

$$\hat{\mu}_2^{25:40} \equiv \hat{\mu}(25 M_\odot, 40 M_\odot), \quad (12)$$

$$\hat{\sigma}_2^{25:40} \equiv \hat{\sigma}(25 M_\odot, 40 M_\odot), \quad (13)$$

which are intended to model the peaks observed with DOUBLEDIP and MULTIPDB near $\sim 9 M_\odot$ and $\sim 32 M_\odot$, respectively.

The precise behavior of these statistics depends somewhat on the integration bounds, which can complicate their interpretation. Put simply, $\hat{\mu}_2^{25:40}$ does not trace μ_2^{peak} exactly. Instead, it acts as an accumulation of information within the range over which μ_2^{peak} operates. Similarly, $\hat{\sigma}_2^{25:40}$ will not perfectly trace σ_2^{peak} . We do not expect our correlations between these statistics and H_0 to greatly change with different integral bounds and have confirmed that changing them does not alter our conclusions.

We also consider the 99th percentile (m_{99}) of p_{1D} as a model-independent proxy for the location of the roll-off at high masses, analogous to m_{max} in Section 4.1. Again, we find the same general trends for different percentiles above $\sim 95\%$.

Interestingly, Table 1 shows that the model-independent statistics (almost) always correlate with H_0 more strongly than their parametric analogs. At first glance, this may be surprising, as one might expect individual model parameters to trace features in the mass distribution better than ad hoc summary statistics. However, as we see in Figure 3, the correlation of some parameters with H_0 can depend on the values taken by other parameters. When σ_2^{peak} is large, μ_2^{peak} correlates poorly with H_0 ($r = 0.18$). In this scenario, Table 1 shows that m_{max} is a better correlator. Conversely, the summary statistics can capture the relevant behavior within the mass distribution regardless of the behavior of individual parameters (i.e., when σ_2^{peak} is large and μ_2^{peak} no longer correlates strongly with H_0 , $\hat{\mu}_2^{25:40}$ can still pick up on the roll-off associated with m_{max}).

Fundamentally, then, it is likely that our summary statistics correlate more strongly with H_0 because they are sensitive to the overall shape of p_{1D} . Individual model parameters may be degenerate, meaning that the same approximate shape of p_{1D} may be obtained with several different parameter combinations. A degeneracy between parameters may weaken their individual correlations with H_0 . Our summary statistics, which may not be closely tied to an individual parameter, allow us to study the relationship between H_0 and the overall shape of p_{1D} .

Figure 3 further demonstrates that our model-independent statistics capture the expected behavior. The correlation

between $\hat{\mu}_2^{25:40}$ and H_0 also improves when we condition on small $\hat{\sigma}_2^{25:40}$ (i.e., a narrow feature). This observation, combined with the fact that $\hat{\mu}_2^{25:40}$ consistently has one of the largest correlation coefficients in Table 1 across all models, strongly suggests that the bulk of the cosmological information in current GW catalogs is carried by binaries with source-frame masses covered by $\hat{\mu}_2^{25:40}$ (i.e., $25\text{--}40 M_\odot$).

This makes sense. High-mass binaries are detected at larger distances (and redshifts), meaning that they can have a larger cosmological imprint on their detector-frame masses. We also tend to detect more of them. However, the properties of the features themselves also play a role. Figure 3 shows that sharper peaks correlate with H_0 more strongly than broad peaks. Furthermore, even though systems from the low-mass feature have a smaller detection horizon, the peak near $\sim 9 M_\odot$ may correlate with H_0 nearly as well as the peak near $\sim 32 M_\odot$ because it is sharper.

We can also understand the signs of the correlations between the summary statistics in Table 1 and H_0 as follows. The correlations observed a posteriori are driven by the different combinations of source-frame mass and H_0 that can predict the same detector-frame mass, which for low z is approximately

$$m_d \approx \left(1 + \frac{H_0 D_L}{c}\right) m_s. \quad (14)$$

The locations of peaks (e.g., $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$) are anticorrelated with H_0 because increasing either peak will tend to increase m_s , which in turn must be compensated by a decrease in H_0 . Table 1 shows that γ_1^{high} (upper edge of the lower notch) is also usually anticorrelated with H_0 , for the same reason.

The behavior of the summary statistics for the width of the peaks ($\hat{\sigma}_1^{7:11}$ and $\hat{\sigma}_2^{25:40}$) can be more complicated, though. In the presence of a peak, we observe that $\hat{\sigma}$ is anticorrelated with H_0 . This is because larger $\hat{\sigma}$ produce a wider astrophysical source-frame mass distribution and because more massive binaries are easier to detect than less massive binaries, this shifts the mean source-frame mass in the detected distribution to larger values. As such, increasing $\hat{\sigma}$ produces a similar effect to increasing $\hat{\mu}$, which requires H_0 to decrease to maintain the same m_d in Equation (14).

Interestingly though, in the absence of a peak, $\hat{\sigma}_1^{7:11}$ is positively correlated with H_0 . This is the opposite of the behavior $\hat{\sigma}_2^{25:40}$ demonstrates. We understand this as the effect of a strong anticorrelation between γ_1^{high} and $\hat{\sigma}_1^{7:11}$ in the absence of a peak near $\sim 9 M_\odot$. That is, when γ_1^{high} increases, it tends to cut out part of the mass distribution within the range spanned by $\hat{\sigma}_1^{7:11}$, which produces a narrower distribution and a smaller value of $\hat{\sigma}_1^{7:11}$ (see posterior credible regions for γ_1^{high} in Table 5). Therefore, we observe a positive correlation between $\hat{\sigma}_1^{7:11}$ and H_0 because $\hat{\sigma}_1^{7:11}$ is anticorrelated with γ_1^{high} , which in turn is anticorrelated with H_0 .

Finally, we again note that, while peaks are helpful for cosmological constraints, they are not necessary. Figure 4 compares the joint posteriors between H_0 and m_{max} , m_{99} for PDB, which, unlike PDB \times P, DOUBLEDIP, and MULTIPDB, does not support a peak near $\sim 32 M_\odot$. Both m_{max} and m_{99} correlate with H_0 comparably. Broadly similar behavior is seen across all models, and similar behavior manifests in $\hat{\mu}_1^{7:11}$ and $\hat{\sigma}_1^{7:11}$ in PDB and PDB \times P. Again, this is likely because m_{99} is sensitive to the overall shape of p_{1D} , whereas m_{max} may not be important

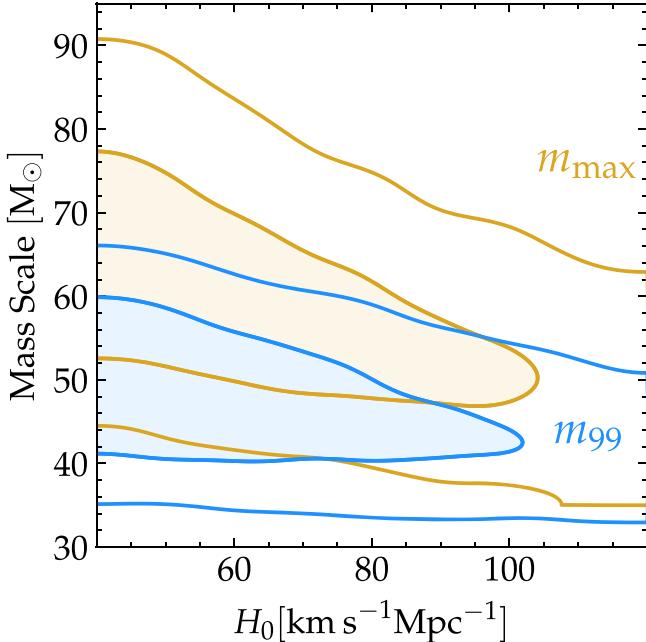


Figure 4. Joint posterior distributions for H_0 and m_{\max} , m_{99} obtained with PDB. Contours denote the 50% and 90% highest-probability-density credible regions. Both mass scales trace the location of the roll-off at high masses. Even for models without prominent peaks, cosmological information is still encoded in the mass distribution.

if there is a large peak in the same mass range. This often seems to be the case in Figure 2.

5. Independent Correlations between Multiple Features and H_0

As seen in Table 1, $\hat{\mu}_1^{7:11}$ can correlate with H_0 nearly as strongly as $\hat{\mu}_2^{25:40}$, particularly for models that support multiple peaks (DOUBLEDIP and MULTIPDB). Figure 5 shows the joint posteriors for MULTIPDB. This would appear to be great news for spectral siren cosmology, as one of the main advantages of the method is that multiple features in the mass distribution can be used to constrain the Hubble relation at the same time.

However, it is not immediately clear whether these features independently correlate with H_0 . That is, it could be the case that $\hat{\mu}_1^{7:11}$ appears to correlate with H_0 only because it correlates with $\hat{\mu}_2^{25:40}$ as they are both related to the shape of the overall mass distribution. Conditioning the posterior on the observed catalog can easily introduce correlations like this. Put another way, $\hat{\mu}_1^{7:11}$ could correlate with $\hat{\mu}_2^{25:40}$ even if we fixed H_0 . In this way, information about H_0 could pass from one feature to another. Therefore, it could be the case that the correlation between $\hat{\mu}_1^{7:11}$ and H_0 observed in Table 1 and Figure 5 are only present because $\hat{\mu}_1^{7:11}$ separately correlates with $\hat{\mu}_2^{25:40}$ which in turn correlates with H_0 . For $\hat{\mu}_1^{7:11}$ to be useful in spectral siren measurements, it must directly correlate with H_0 independently of $\hat{\mu}_2^{25:40}$.

Another way to phrase this is whether $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ are conditionally independent given H_0 ($\hat{\mu}_1^{7:11} \perp \hat{\mu}_2^{25:40} | H_0$). Figure 6 directly addresses this by plotting the joint posterior distributions for $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ conditioned on several (small) ranges of H_0 . We see that $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ are correlated in their joint posterior, but essentially all of the correlation is due

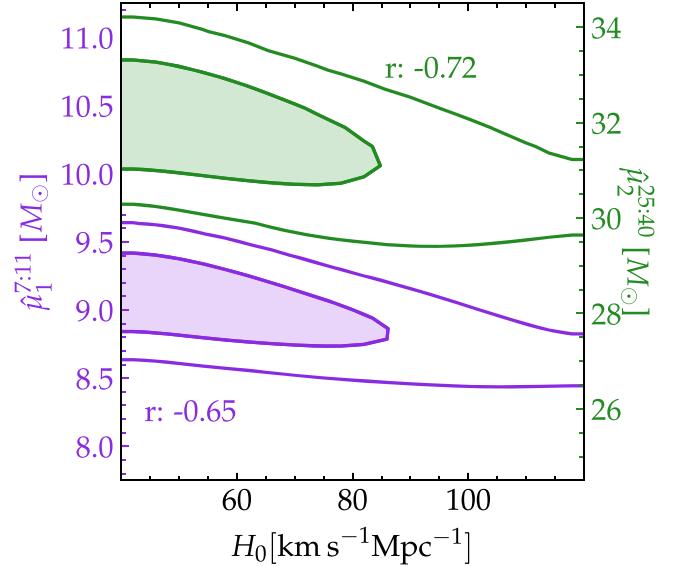


Figure 5. Joint posteriors between H_0 and $\hat{\mu}_1^{7:11}$ (purple), $\hat{\mu}_2^{25:40}$ (green). Contours denote the 50% and 90% highest-probability-density credible regions. The estimator functions defined in Equation (9) are applied to MULTIPDB. They correlate similarly for both local overdense regions.

to their separate correlations with H_0 . In fact, $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ are almost completely uncorrelated when we condition on H_0 , but their joint posterior shifts to larger values as we increase H_0 .⁷

To quantify how much of the correlation between $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ is due to their separate correlations with H_0 , we consider a simple model in which $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ separately depend linearly on H_0 .⁸

$$\hat{\mu}_1^{7:11} = (aH_0 + b) + \delta, \quad (15)$$

$$\hat{\mu}_2^{25:40} = (\alpha H_0 + \beta) + \epsilon, \quad (16)$$

where δ and ϵ are drawn from a joint distribution that does not depend on H_0 . If we additionally assume $H_0 \sim p(H_0)$, it is straightforward to show that

$$C[\hat{\mu}_1^{7:11}, \hat{\mu}_2^{25:40}] = \frac{C[\hat{\mu}_1^{7:11}, H_0]C[\hat{\mu}_2^{25:40}, H_0]}{V[H_0]} + C[\delta, \epsilon], \quad (17)$$

where $C[x, y]$ denotes the covariance between x and y and $V[x] = C[x, x]$ is the variance of x . The first term represents the covariance induced by the fact that both $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ depend on the same variable (H_0), and the second term corresponds to the “extra” covariance between the two that would exist even at fixed H_0 : $C[\hat{\mu}_1^{7:11}, \hat{\mu}_2^{25:40}|H_0]$. It is natural to compare the size of these terms, and we define the ratio

$$\mathcal{R} = \left| \frac{C[\hat{\mu}_1^{7:11}, H_0]C[\hat{\mu}_2^{25:40}, H_0]}{C[\hat{\mu}_1^{7:11}, \hat{\mu}_2^{25:40}]V[H_0] - C[\hat{\mu}_1^{7:11}, H_0]C[\hat{\mu}_2^{25:40}, H_0]} \right|. \quad (18)$$

⁷ This is another manifestation of Simpson’s paradox.

⁸ There is no compelling reason to believe the dependence is this simple; we only use this model to motivate more general expressions.

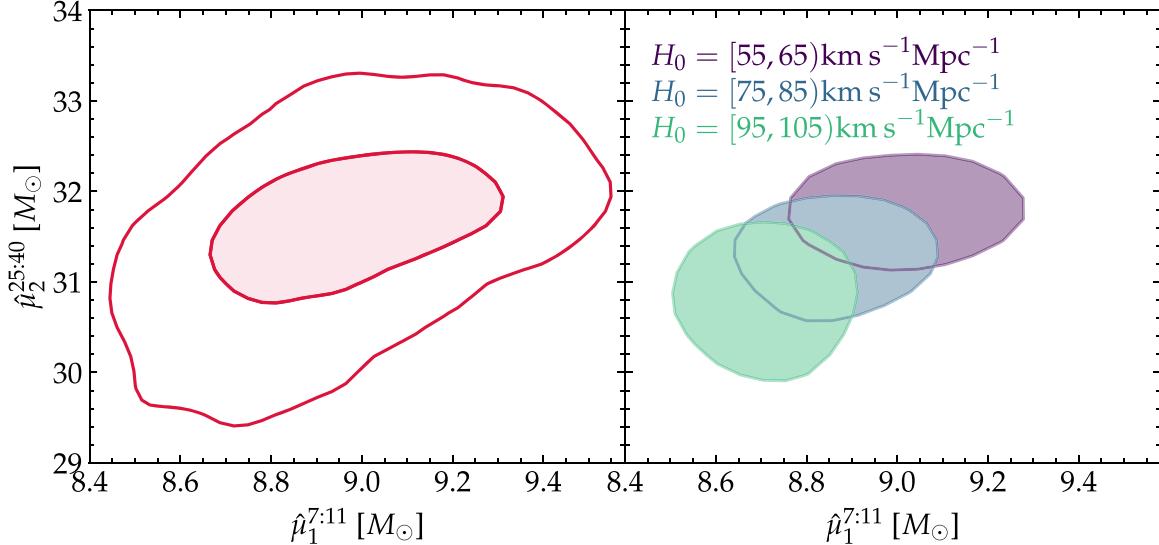


Figure 6. (Left) Joint posterior distribution over $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ in MULTIPDB. Contours denote the 50% and 90% highest-probability-density credible regions. (Right) 50% posterior credible regions from the joint posterior conditioned on different ranges of H_0 : 55–65 (purple), 75–85 (blue), and 95–105 (teal) $\text{km s}^{-1} \text{Mpc}^{-1}$. We omit the 90% credible regions for clarity, as they significantly overlap.

Table 2

Pearson Correlation Coefficients (r) for $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ Along with the Ratio of Covariances from Their Joint Dependence on H_0 and All Other Factors (\mathcal{R} , Equation (18)), Ideally $R \gg 1$

	r	\mathcal{R}
PDB	-0.009	0.87
PDB \times P	0.146	7.92
DOUBLEDIP	0.242	19.17
MULTIPDB	0.467	126.28

Note. All models with significant correlations between $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ also have large \mathcal{R} .

If $\mathcal{R} \gg 1$, then the vast majority of the correlation between $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$ is due to their separate dependence on H_0 , which is the ideal situation for a spectral sirens constraint. Table 2 summarizes these results, and we find that the model-independent estimates’ covariance is indeed almost entirely driven by their separate correlations with H_0 . The lone exception to this is PDB, which can be explained by the poor correlation between $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$. Even then, the correlation from the separate dependence on H_0 is almost as large as all other factors combined.

Although we restrict our analysis to $\hat{\mu}_1^{7:11}$ and $\hat{\mu}_2^{25:40}$, a similar analysis can be applied to other estimators. We leave such studies for future work.

6. Discussion

We have shown that the observed mass distribution, as inferred with CBCs at all mass scales from O3, appears to contain multiple features that separately correlate strongly with H_0 . Future analyses can use this work to further improve spectral siren constraints. However, we make several simplifying assumptions in our analysis that may require further study.

First, although we considered several variations of the one-dimensional latent mass distribution p_{1D} , we only considered a single functional form for a mass-dependent pairing function in

Equation (5). We also assume that m_{1s} and m_{2s} are drawn from the same p_{1D} . Alternatives have been proposed in the literature. A. M. Farah et al. (2024) looked at whether there is evidence that m_{1s} and m_{2s} are drawn from the same underlying 1D mass distributions. They found that current data slightly favor our model choice. M. Fishbach & D. E. Holz (2020) and A. Farah et al. (2022) examined different pairing functions with earlier GW catalogs; both found that our pairing function is a reasonable description of the data. It is also similar to the common model assumption that $p(m_{2s}|m_{1s}) \propto m_{2s}^\beta$ (A. M. Farah et al. 2024).

In general, model misspecification of this kind is a persistent concern in any inference that assumes a specific functional form for the mass distribution. G. Pierra et al. (2024) show that mass model choices may bias H_0 measurements. We believe our parametric models provide a reasonable fit to the current GW catalog, particularly for BBH masses where a strong preference for equal-mass binaries has been observed (M. Fishbach & D. E. Holz 2020). Recent nonparametric mass models also show that the source-frame mass distribution we use is realistic and matches the observed data (S. Rinaldi & W. Del Pozzo 2021; V. Tiwari & S. Fairhurst 2021; T. A. Callister & W. M. Farr 2024; A. Ray et al. 2024; A. M. Farah et al. 2025). What’s more, just as models with peaks may be better able to constrain H_0 than models without peaks, models with more complicated pairing functions (that depend on scales in the source frame) should only improve the ability of spectral sirens to constrain H_0 . In this respect, our assumption may in fact be conservative.

Misspecification of the spin distribution may also affect our inference of H_0 , primarily through its impact on the inferred masses. Specifically, it is known that spins are correlated with masses within single-event uncertainties, the effective spin and mass ratio appear to be correlated at the population level (T. A. Callister et al. 2021), and the spin population is unlikely to be uniform in magnitude and isotropic in orientation (R. Abbott et al. 2023b). However, we expect our assumption of a fixed, broad spin distribution to have a small impact on the overall finding that there are multiple relatively prominent

peaks in the mass distribution. It is the presence of those peaks which drives our main conclusions.

In addition to our mass and spin distribution, our analysis makes several assumptions about the evolution of the distribution of merging binaries with redshift. In particular, one may be concerned that we do not know the true redshift distribution $p(z|\Lambda)$ and that our assumption of a fixed redshift distribution (which tracks the star formation rate) is overly optimistic. While additional uncertainty in $p(z|\Lambda)$ is likely to weaken posterior constraints on H_0 , we do not expect it to completely spoil our ability to constrain H_0 . This may change the typical redshift at which we observe events, but it is unlikely to change the relationship between m_d and z , which is what connects back to specific features in $p(m_{1s}, m_{2s}|\Lambda)$ upon which we base our measurement. However, additional work is warranted to show that this is indeed the case.⁹

The other significant assumption about the redshift dependence within our analysis is that the source-frame mass distribution $p(m_{1s}, m_{2s}|\Lambda)$ is independent of z . In reality, this almost certainly is not the case.¹⁰ It is known that metallicity correlates with redshift and, therefore, stars that form earlier in the Universe form in more metal-poor environments (P. Madau & M. Dickinson 2014). We expect the metallicity of the formation environment to impact the masses of the stellar remnants left behind (P. Madau & T. Fragos 2017). Additionally, multiple separate formation channels may have been active during different epochs of cosmic history. If these formation channels preferentially produce different types of binaries, then, again, the source-frame mass distribution may depend on redshift (K. K. Y. Ng et al. 2021; L. A. C. van Son et al. 2022; C. S. Ye & M. Fishbach 2024). Given that the location and shapes of features in the source-frame mass distribution may shift with redshift, some authors have raised the reasonable concern that this astrophysical evolution could be confused with cosmological redshift. Both could manifest as a dependence of the detector-frame mass distribution on the luminosity distance.

While our assumption that the source-frame mass distribution is independent of redshift does not allow us to directly address this concern, we do provide evidence that it may not be a showstopper. Specifically, while we may expect the source-frame mass distribution to evolve, we do not expect it to evolve in the same way at all mass scales. However, a cosmological redshift will affect all mass scales in the same way.

This is similar to how redshifts observed for both atomic and molecular transition lines are evidence for cosmological redshift rather than a conspiracy of changes in the underlying physics for each set of lines. Changes in the electron mass could affect atomic transitions in, e.g., hydrogen, but molecular transitions in H_2 are controlled by the mass of the proton. It is more parsimonious to infer cosmological redshift than to contrive a model that changes the mass of both the electron and the proton simultaneously to mimic the observed behavior.

In the same way, if multiple features in the source-frame mass distribution correlate separately with H_0 (as we have shown is the case in Section 5), we expect to be able to break

⁹ Single-event measurement uncertainties on D_L can be broad, and different population priors may be able to significantly alter our posterior beliefs about individual events.

¹⁰ Current observations do not rule out redshift evolution of the source-frame mass distribution, but it is not required to explain the data either (M. Fishbach et al. 2021).

the degeneracy between astrophysical and cosmological effects. See A. M. Farah et al. (2025) for an explicit demonstration of how this could work with simulated data.

So far, the features we identified are within the BBH portion of the mass distribution. However, it is expected that the larger detection rate of BNS with next-generation GW detectors will eventually drive spectral siren constraints (J. M. Ezquiaga & D. E. Holz 2022). BNS and NSBH will also provide other tracers of the Hubble relation, as tidal effects in the GW waveform will not redshift and can be used as an alternate measurement of the source-frame masses. In general, any binary property that does not redshift but is correlated with the source-frame masses, like NS maximum mass or component spins, could be used in this way (C. Messenger & J. Read 2012; D. Chatterjee et al. 2021; T. Ghosh et al. 2022, 2024; S. Mukherjee 2022; H.-Y. Chen et al. 2024). However, it is difficult to measure tides and spins in the current catalog of CBCs, so it is unclear how much of a near-term improvement these additional features could provide. We note that additional information of this type should only improve our ability to measure H_0 from GW catalogs without EM counterparts: the more independent features, the easier it is to disentangle astrophysical evolution from cosmology.

Our results show that it is, in fact, quite difficult to completely remove cosmological information from the observed distribution of CBC properties. When we condition our hierarchical model on observed data, any features in the source-frame mass distribution, such as dips, gaps, peaks, and roll-offs, correlate with H_0 (and the full Hubble relation). We identify several robust, model-independent features and show that they correlate strongly with H_0 across a range of models. We also show that these features correlate with H_0 independently, which is the best-case scenario for spectral siren measurements. Looking ahead, even in the presence of redshift evolution in the source-frame mass distribution, current data suggests we live in a universe in which spectral sirens can provide an accurate measurement of H_0 .

Acknowledgments

We sincerely thank Aditya Vijaykumar, José Ezquiaga, Maya Fishbach, Amanda Farah, Alex Hanselman, and Phillip Landry for the useful comments and discussions throughout the project. U.M. and R.E. are supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) through a Discovery grant (RGPIN-2023-03346). This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation.

Software: NumPy (C.R. Harris et al. 2020), pandas (pandas development team 2020), NumPyro (D. Phan et al. 2019; E. Bingham et al. 2019), jax (J. Bradbury et al. 2018), matplotlib (J. D. Hunter 2007), ScipY (P. Virtanen et al. 2020).

Appendix A Adding versus Multiplying to Create Additional Features

As mentioned in Section 2.1, there are several nearly equivalent methods for adding additional peaks to the mass distribution. We focus on multiplicative filters that can either add a peak or remove a notch based on the sign of their amplitude parameter (see Appendix B). However, it is also

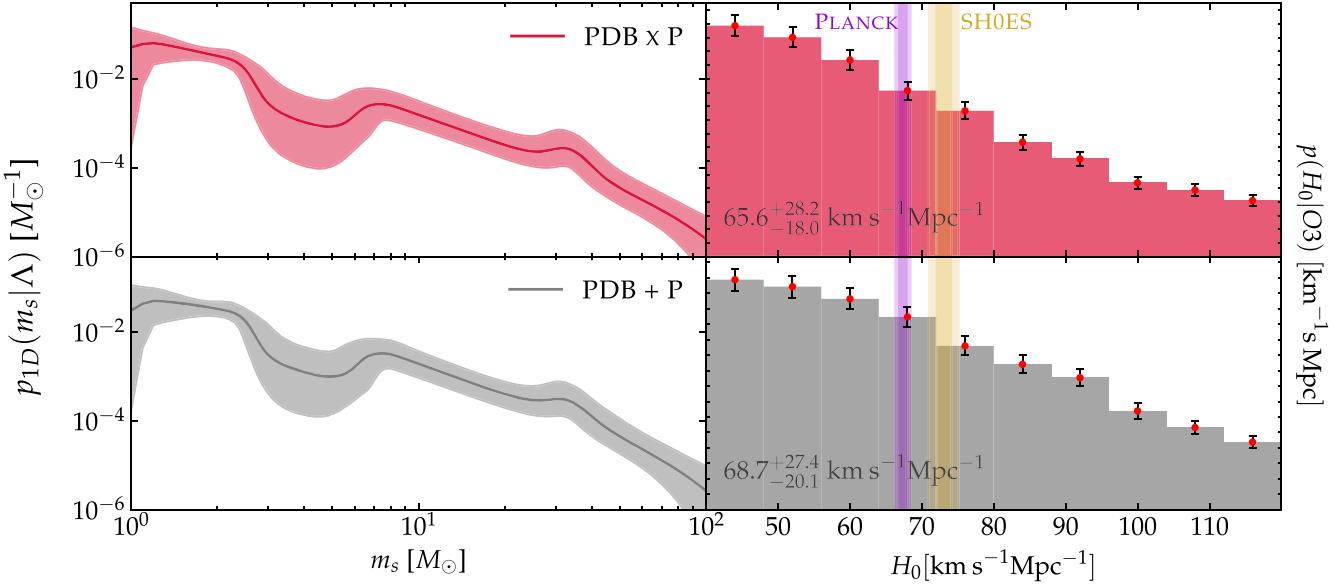


Figure 7. Similar to Figure 2, we compare two different extensions of PDB, each with a single peak: PDB \times P (Equation (B7)) and PDB+P (Equation (B9)).

common to add a peak as an additional component within a mixture model. We consider such a mixture for p_{1D} by summing PDB and a Gaussian peak (PDB + P; Equation (B9)).

We find quantitatively similar results with PDB \times P and PDB + P (Figure 7). Both marginal posteriors for H_0 yield similar constraints, the credible regions for p_{1D} show similar features, and features in each mass model show similar correlations with H_0 .

This suggests that, as expected, a model’s behavior does not strongly depend on the precise implementation of additional peaks, and analysts should choose whichever implementation is

function introduced in A. Farah et al. (2022).

$$f(m_{1s}, m_{2s}; \beta_1, \beta_2) = \begin{cases} (m_{2s}/m_{1s})^{\beta_1} & \text{if } m_{2s} \leq 5 M_\odot \\ (m_{2s}/m_{1s})^{\beta_2} & \text{if } m_{2s} > 5 M_\odot \end{cases} \quad (\text{B1})$$

This pairing function introduces a discontinuity in $p(m_{1s}, m_{2s} | \Lambda)$ at the fixed switch-point $m_{2s} = 5 M_\odot$. It is not clear whether such a discontinuity is appropriate, but it is not thought to affect our conclusions.

Within this framework, we also construct more complicated p_{1D} by starting with a base model (PDB)

$$p_{\text{PDB}}(m | \Lambda) \propto p_{\text{BPL}}(m | \alpha_1, \alpha_2, m_{\text{brk}}) h(m | m_{\text{min}}, \eta_{\text{min}}) n(m | \gamma_1^{\text{low}}, \gamma_1^{\text{high}}, \eta_1^{\text{low}}, \eta_1^{\text{high}}, A_1) \ell(m | m_{\text{max}}, \eta_{\text{max}}). \quad (\text{B2})$$

easiest to control. For example, one does not need to normalize p_{1D} within each likelihood call if additional features are included as multiplicative factors (the normalization cancels term-by-term between \mathcal{Z}_i and \mathcal{E} in Equation (1)). However, it may be necessary (and expensive!) to numerically normalize each mixture model component within each likelihood evaluation to obtain interpretable mixing fractions.

Appendix B Details of the Population Model

We define the population model as a product of distributions over the binary’s source-frame component masses (Section B.1), component spins (Section B.2), and redshift (Section B.3).

B.1. Source-frame Mass Models

As shown in Equation (5), we separate the source-mass distribution into two one-dimensional mass distributions and a pairing function. While we change p_{1D} to add additional features, we always consider a simple model for the pairing

This combines a broken power law

$$p_{\text{BPL}}(m | \alpha_1, \alpha_2, m_{\text{brk}}) = \begin{cases} (m/m_{\text{brk}})^{\alpha_1} & \text{if } m < m_{\text{brk}} \\ (m/m_{\text{brk}})^{\alpha_2} & \text{if } m \geq m_{\text{brk}} \end{cases} \quad (\text{B3})$$

with Butterworth low-pass and high-pass filters

$$\ell(m | k, \eta) = (1 + (m/k)^\eta)^{-1} \quad (\text{B4})$$

$$h(m | k, \eta) = 1 - \ell(m | k, \eta) \quad (\text{B5})$$

and a notch filter

$$\begin{aligned} n(m | \gamma_1^{\text{low}}, \gamma_1^{\text{high}}, \eta_1^{\text{low}}, \eta_1^{\text{high}}, A_1) \\ = 1 - A_1 h(m | \gamma_1^{\text{low}}, \eta_1^{\text{low}}) \ell(m | \gamma_1^{\text{high}}, \eta_1^{\text{high}}). \end{aligned} \quad (\text{B6})$$

Depending on the sign of A_1 , the notch filter can either remove a notch ($A_1 > 0$) or add a peak ($A_1 < 0$).

We then extend this model to include multiple peaks and dips. PDB \times P adds a single additional peak through a multiplicative filter.

$$\begin{aligned} p_{\text{PDB} \times \text{P}}(m | \Lambda) = p_{\text{PDB}}(m | \Lambda) \\ \times (1 + c_2 \mathcal{G}(m | \mu_2^{\text{peak}}, \sigma_2^{\text{peak}})), \end{aligned} \quad (\text{B7})$$

where $\mathcal{G}(m|\mu, \sigma)$ is a Gaussian distribution with mean μ and standard deviation σ .

$$\mathcal{G}(m|\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(m-\mu)^2}{2\sigma^2}\right). \quad (\text{B8})$$

Alternatively, PDB+P (Appendix A) constructs a mixture model.

$$p_{\text{PDB+P}}(m|\Lambda) = (1 - \lambda_2)p_{\text{PDB}}(m|\Lambda) + \lambda_2 \mathcal{G}(m|\mu_2^{\text{peak}}, \sigma_2^{\text{peak}}). \quad (\text{B9})$$

DOUBLEDIP extends PDB by adding another notch filter at high masses.

$$p_{\text{DOUBLEDIP}}(m|\Lambda) = p_{\text{PDB}}(m|\Lambda) \times n(m|\gamma_2^{\text{low}}, \gamma_2^{\text{high}}, \eta_2^{\text{low}}, \eta_2^{\text{high}}, A_2), \quad (\text{B10})$$

while MULTIPDB instead extends PDB by adding a multiplicative filter containing two Gaussian peaks.

$$p_{\text{MULTIPDB}}(m|\Lambda) = p_{\text{PDB}}(m|\Lambda) \times (1 + c_1 \mathcal{G}(m|\mu_1^{\text{peak}}, \sigma_1^{\text{peak}}) + c_2 \mathcal{G}(m|\mu_2^{\text{peak}}, \sigma_2^{\text{peak}})). \quad (\text{B11})$$

B.2. Spin Model

We assume both components' spins are independently and identically distributed uniformly in magnitude and isotropically in orientation. The distribution over Cartesian spin components

is therefore

$$p(s_1, s_2|\Lambda) \propto \frac{1}{|s_1|^2 |s_2|^2}. \quad (\text{B12})$$

B.3. Redshift Model

As discussed in Section 2.1, we assume the merger rate follows the star formation rate (Equation (4)). We follow M. Fishbach et al. (2018) and define

$$\Phi(z) = 0.015 \left(\frac{(1+z)^{2.7}}{1 + [(1+z)/2.9]^{5.6}} \right) \frac{M_\odot}{\text{yr Mpc}^3}. \quad (\text{B13})$$

The rest of our redshift model implicitly depends on the flat Λ CDM cosmology assumed through the comoving volume (V_c). See Table 3 for the precise values of our cosmological model.

Appendix C Priors and Posteriors

Tables 3 and 4 list the priors assumed within our analysis. All priors are uniform over a restricted range, and we denote the uniform distribution between X and Y as $U(X, Y)$.

Table 5 lists posterior medians and 90% symmetric credible regions for all model parameters for all models. It additionally lists posterior credible regions for the model-independent summary statistics introduced in Section 4.2.

Table 3
Hyperpriors for Parameters That are Common to Each Mass Model

	Parameter		Description	Prior
Cosmology	H_0	[km s ⁻¹ Mpc ⁻¹]	Present expansion rate	$U(40, 120)$
	Ω_m		Matter density	$U(0, 1)$
	Ω_r		Radiation density	0.001
	Ω_k		Curvature	0
Pairing function	Ω_Λ		Cosmological constant	$1 - \Omega_m - \Omega_r - \Omega_k$
	β_1		Spectral index below $5M_\odot$	$U(0, 10)$
	β_2		Spectral index above $5M_\odot$	$U(0, 10)$
Broken power law	α_1		Spectral index below m_{brk}	$U(-5, 5)$
	α_2		Spectral index above m_{brk}	$U(-5, 5)$
	m_{brk}	[M_\odot]	Dividing point for α_1 and α_2	$U(2, 5)$
High-pass filter	m_{min}	[M_\odot]	Roll-off scale for low masses	$U(0.5, 1.2)$
	η_{min}		Sharpness of the roll-off at m_{min}	$U(25, 50)$
Low-pass filter	m_{max}	[M_\odot]	Roll-off scale for high masses	$U(35, 100)$
	η_{max}		Sharpness of the roll-off at m_{max}	$U(0, 10)$
Low-mass notch	γ_1^{low}	[M_\odot]	Lower edge of low-mass notch	$U(2.3, 4)$
	η_1^{low}		Sharpness of the roll-off at γ_1^{low}	$U(0, 50)$
	γ_1^{high}	[M_\odot]	Upper edge of low-mass notch	$U(4, 8)$
	η_1^{high}		Sharpness of the roll-off at γ_1^{high}	$U(0, 50)$
	A_1		Depth of low-mass notch	$U(0, 1)$

Note. We denote a uniform distribution between X and Y as $U(X, Y)$. For fixed parameters, we simply report the value assumed.

Table 4
Additional Hyperpriors for Each Mass Model

	Parameter	Description	PDB	PDB+P	PDB \times P	DOUBLEDIP	MULTIPDB
High-mass notch	γ_2^{low}	[M_\odot]	Lower edge of high-mass notch	$U(6, 60)$
	η_2^{low}		Sharpness of the roll-off at γ_2^{low}	$U(0, 50)$
	γ_2^{high}	[M_\odot]	Upper edge of high-mass notch	$U(6, 60)$
	η_2^{high}		Sharpness of the roll-off at γ_2^{high}	$U(0, 50)$
Low-mass peak	A_2		Depth of high-mass notch	0	0	0	$U(0, 1)$
	μ_1^{peak}	[M_\odot]	Location of low-mass peak
	σ_1^{peak}	[M_\odot]	Width of low-mass peak	$U(1, 40)$
High-mass peak	c_1		Height of the low-mass peak	0	0	0	$U(0, 100)$
	μ_2^{peak}	[M_\odot]	Location of high-mass peak	...	$U(20, 60)$	$U(20, 60)$...
	σ_2^{peak}	[M_\odot]	Width of high-mass peak	...	$U(1, 40)$	$U(1, 40)$...
	c_2		Height of high-mass peak	0	0	$U(0, 100)$	0
	λ_2		Mixing frac. of high-mass peak	0	$U(0, 1)$	0	0

Table 5
Posterior Medians and 90% Symmetric Credible Regions of All Hyperparameters

	Parameter	PDB	PDB \times P	DOUBLEDIP	MULTIPDB
Cosmology	H_0	[$\text{km s}^{-1} \text{Mpc}^{-1}$]	$68.71^{+43.93}_{-26.20}$	$65.60^{+44.69}_{-23.12}$	$71.44^{+38.58}_{-28.21}$
	Ω_m		$0.46^{+0.48}_{-0.42}$	$0.43^{+0.50}_{-0.39}$	$0.41^{+0.52}_{-0.38}$
Pairing function	β_1		$0.96^{+1.01}_{-0.68}$	$0.96^{+0.99}_{-0.70}$	$0.89^{+0.94}_{-0.67}$
	β_2		$2.73^{+2.01}_{-1.27}$	$2.51^{+1.76}_{-1.16}$	$2.69^{+1.85}_{-1.16}$
Broken power law	α_1		$-2.11^{+2.03}_{-1.93}$	$-1.21^{+2.14}_{-1.61}$	$-1.75^{+2.68}_{-2.61}$
	α_2		$-1.75^{+0.35}_{-0.30}$	$-2.29^{+0.45}_{-0.47}$	$-1.55^{+1.85}_{-0.87}$
	m_{brk}	[M_\odot]	$3.05^{+1.71}_{-0.97}$	$3.60^{+1.27}_{-1.47}$	$3.39^{+1.42}_{-1.23}$
High-pass filter	m_{min}	[M_\odot]	$0.87^{+0.30}_{-0.33}$	$0.86^{+0.31}_{-0.32}$	$0.86^{+0.31}_{-0.33}$
	η_{min}		$37.46^{+11.33}_{-11.28}$	$37.57^{+11.24}_{-11.39}$	$37.48^{+11.09}_{-11.13}$
Low-pass filter	m_{max}	[M_\odot]	$58.42^{+25.05}_{-15.58}$	$79.39^{+18.74}_{-36.00}$	$40.68^{+29.91}_{-5.31}$
	η_{max}		$5.53^{+3.71}_{-3.01}$	$3.50^{+5.84}_{-3.22}$	$5.90^{+3.88}_{-3.32}$
	γ_1^{low}	[M_\odot]	$2.66^{+0.88}_{-0.33}$	$2.56^{+0.60}_{-0.24}$	$2.61^{+0.88}_{-0.28}$
	η_1^{low}		$27.30^{+20.25}_{-21.06}$	$27.77^{+19.90}_{-19.85}$	$28.15^{+19.77}_{-20.95}$
Low-mass notch	γ_1^{high}	[M_\odot]	$6.05^{+1.38}_{-1.37}$	$6.38^{+1.24}_{-1.23}$	$7.24^{+0.70}_{-1.94}$
	η_1^{high}		$30.06^{+17.82}_{-21.01}$	$30.59^{+17.42}_{-20.23}$	$30.80^{+17.21}_{-21.42}$
	A_1		$0.88^{+0.11}_{-0.45}$	$0.92^{+0.07}_{-0.25}$	$0.92^{+0.08}_{-0.37}$
High-mass notch	γ_2^{low}	[M_\odot]	$10.98^{+2.67}_{-1.44}$
	η_2^{low}		$32.83^{+15.57}_{-21.53}$
	γ_2^{high}	[M_\odot]	$28.13^{+31.32}_{-8.41}$
	η_2^{high}		$28.78^{+19.26}_{-23.83}$
Low-mass peak	A_2		$0.86^{+0.11}_{-0.29}$
	μ_1^{peak}	[M_\odot]
	σ_1^{peak}	[M_\odot]
	c_1	
High-mass peak	$\frac{c_1}{\sigma_1^{\text{peak}} \sqrt{2\pi}}$	
	μ_2^{peak}	[M_\odot]	...	$34.74^{+8.03}_{-5.44}$...
	σ_2^{peak}	[M_\odot]	...	$4.25^{+16.18}_{-2.82}$...
	c_2		...	$32.25^{+47.62}_{-19.12}$...
Summary statistics	$\frac{c_2}{\sigma_2^{\text{peak}} \sqrt{2\pi}}$...	$3.10^{+4.83}_{-2.25}$...
	$\hat{\mu}_1^{7:11}$	[M_\odot]	$8.74^{+0.16}_{-0.05}$	$8.67^{+0.26}_{-0.08}$	$8.83^{+0.29}_{-0.29}$
	$\hat{\sigma}_1^{7:11}$	[M_\odot]	$1.13^{+0.01}_{-0.05}$	$1.12^{+0.01}_{-0.08}$	$1.04^{+0.10}_{-0.20}$
	$\hat{\mu}_2^{25:40}$	[M_\odot]	$31.35^{+0.17}_{-0.23}$	$31.90^{+1.69}_{-1.67}$	$31.33^{+1.06}_{-0.49}$
	$\hat{\sigma}_2^{25:40}$	[M_\odot]	$4.24^{+0.03}_{-0.06}$	$3.98^{+0.34}_{-1.01}$	$4.13^{+0.11}_{-0.19}$
	m_{99}	[M_\odot]	$46.88^{+13.58}_{-11.10}$	$42.52^{+12.98}_{-8.23}$	$45.59^{+13.08}_{-9.51}$
					$44.90^{+15.76}_{-10.31}$

Note. We have added the term $\frac{c}{\sigma^{\text{peak}} \sqrt{2\pi}}$. It should be interpreted as an amplitude representing the height of the peak compared to the surrounding mass distribution.

ORCID iDs

Utkarsh Mali <https://orcid.org/0009-0003-1285-2788>
 Reed Essick <https://orcid.org/0000-0001-8196-9267>

References

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, *CQGra*, **32**, 074001
 Abbott, B., Abbott, R., Abbott, T., et al. 2016a, *PhRvL*, **116**, 061102
 Abbott, B., Abbott, R., Abbott, T., et al. 2016b, *PhRvX*, **6**, 041015
 Abbott, B., Abbott, R., Abbott, T., et al. 2019, *PhRvX*, **9**, 031040
 Abbott, B., Abbott, R., Abbott, T., et al. 2021a, *ApJ*, **909**, 218
 Abbott, B. P., Abbott, R., Abbott, T., et al. 2017, *PhRvL*, **119**, 161101
 Abbott, R., Abbott, T., Abraham, S., et al. 2020, *PhRvL*, **125**, 101102
 Abbott, R., Abbott, T. D., Acernese, F., et al. 2023a, *PhRvX*, **13**, 041039
 Abbott, R., Abbott, T. D., Acernese, F., et al. 2024, *PhRvD*, **109**, 022001
 Abbott, R., Abbott, T., Abraham, S., et al. 2021b, *PhRvX*, **11**, 021053
 Abbott, R., Abe, H., Acernese, F., et al. 2021c, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run O3 Search Sensitivity Estimates, Zenodo, doi:[10.5281/zenodo.5546676](https://doi.org/10.5281/zenodo.5546676)
 Abbott, R., Abe, H., Acernese, F., et al. 2021d, *ApJ*, **949**, 76
 Abbott, R., Abbott, T. D., Acernese, F., et al. 2023b, *PhRvX*, **13**, 011048
 Acernese, F., Agathos, M., Agatsuma, K., et al. 2014, *CQGra*, **32**, 024001
 Aghanim, N., Akrami, Y., Ashdown, M., et al. 2020, *A&A*, **641**, A6
 Akutsu, T., Ando, M., Arai, K., et al. 2020, *PTEP*, **2021**, 05A101
 Aubin, F., Brightenti, F., Chierici, R., et al. 2021, *CQGra*, **38**, 095004
 Bingham, E., Chen, J. P., Jankowiak, M., et al. 2019, *J. Mach. Learn. Res.*, **20**, 1, <https://jmlr.org/papers/v20/18-403.html>
 Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, JAX: Composable Transformations of Python+NumPy Programs, 0.3.13., <https://github.com/google/jax>
 Callister, T. A., & Farr, W. M. 2024, *PhRvX*, **14**, 021005
 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
 Chatterjee, D., K. R., A.H., Holder, G., et al. 2021, *PhRvD*, **104**, 083528
 Chen, H.-Y., Ezquiaga, J. M., & Gupta, I. 2024, *CQGra*, **41**, 125004
 Chen, H.-Y., Fishbach, M., & Holz, D. E. 2018, *Natur*, **562**, 545
 Chernoff, D. F., & Finn, L. S. 1993, *ApJ*, **411**, L5
 Dalal, N., Holz, D. E., Hughes, S. A., & Jain, B. 2006, *PhRvD*, **74**, 063006
 Del Pozzo, W. 2012, *PhRvD*, **86**, 043011
 Edelman, B., Farr, B., & Doctor, Z. 2023, *ApJ*, **946**, 16
 Essick, R. 2023, *PhRvD*, **108**, 043011
 Essick, R., & Fishbach, M. 2024, *ApJ*, **962**, 169
 Ezquiaga, J. M., & Holz, D. E. 2022, *PhRvL*, **129**, 061102
 Farah, A., Fishbach, M., Essick, R., Holz, D. E., & Galaudage, S. 2022, *ApJ*, **931**, 108
 Farah, A. M., Callister, T. A., Ezquiaga, J. M., Zevin, M., & Holz, D. E. 2025, *ApJ*, **978**, 153
 Farah, A. M., Edelman, B., Zevin, M., et al. 2023, *ApJ*, **955**, 107
 Farah, A. M., Fishbach, M., & Holz, D. E. 2024, *ApJ*, **962**, 69
 Farr, W. M., Fishbach, M., Ye, J., & Holz, D. E. 2019, *ApJL*, **883**, L42
 Fishbach, M., Doctor, Z., Callister, T., et al. 2021, *ApJ*, **912**, 98
 Fishbach, M., Essick, R., & Holz, D. E. 2020, *ApJ*, **899**, L8
 Fishbach, M., Gray, R., Hernandez, I. M., et al. 2019, *ApJL*, **871**, L13
 Fishbach, M., & Holz, D. E. 2020, *ApJL*, **891**, L27
 Fishbach, M., Holz, D. E., & Farr, W. M. 2018, *ApJL*, **863**, L41
 Freedman, W. L., & Madore, B. F. 2023, *JCAP*, **2023**, 050
 Freedman, W. L., Madore, B. F., Hoyt, T., et al. 2020, *ApJ*, **891**, 57
 Freedman, W. L., Madore, B. F., Jang, I. S., et al. 2024, arXiv:[2408.06153](https://arxiv.org/abs/2408.06153)
 Gair, J. R., Ghosh, A., Gray, R., et al. 2023, *AJ*, **166**, 22
 Ghosh, T., Biswas, B., & Bose, S. 2022, *PhRvD*, **106**, 123529
 Ghosh, T., Biswas, B., Bose, S., & Kapadia, S. J. 2024, arXiv:[2407.16669](https://arxiv.org/abs/2407.16669)
 Godfrey, J., Edelman, B., & Farr, B. 2024, arXiv:[2304.01288](https://arxiv.org/abs/2304.01288)
 Gray, R., Hernandez, I. M., Qi, H., et al. 2020, *PhRv*, **101**, 122001
 Gray, R., Messenger, C., & Veitch, J. 2022, *MNRAS*, **512**, 1127
 Hanselman, A. G., Vijaykumar, A., Fishbach, M., & Holz, D. E. 2025, *ApJ*, **979**, 9
 Harris, C.R., Millman, K.J., van der Walt, S.J., et al. 2020, *Natur*, **585**, 357
 Hernandez, I. M., & Ray, A. 2024, arXiv:[2404.02522](https://arxiv.org/abs/2404.02522)
 Holz, D. E., & Hughes, S. A. 2005, *ApJ*, **629**, 15
 Hubble, E. 1929, *PNAS*, **15**, 168
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Karathanasis, C., Mukherjee, S., & Mastrogiovanni, S. 2023, *MNRAS*, **523**, 4539
 Klimenko, S., Vedovato, G., Necula, V., et al. 2021, cWB Pipeline Library: v6.4.0, cWB-6.4.0, Zenodo, doi:[10.5281/zenodo.4419902](https://doi.org/10.5281/zenodo.4419902)
 Lee, M. G., Freedman, W. L., & Madore, B. F. 1993, *ApJ*, **417**, 553
 Leyde, K., Mastrogiovanni, S., Steer, D., Chassande-Mottin, E., & Karathanasis, C. 2022, *JCAP*, **2022**, 012
 Li, Y.-J., Tang, S.-P., Wang, Y.-Z., & Fan, Y.-Z. 2024, *ApJ*, **976**, 153
 Madau, P., & Dickinson, M. 2014, *ARA&A*, **52**, 415
 Madau, P., & Fragos, T. 2017, *ApJ*, **840**, 39
 Mancarella, M., Genoud-Prache, E., & Maggiore, M. 2022, *PhRvD*, **105**, 064030
 Mandel, I. 2010, *PhRvD*, **81**, 084029
 Mandel, I., Farr, W. M., & Gair, J. R. 2019, *MNRAS*, **486**, 1086
 Mapelli, M. 2020, *FrASS*, **7**, 38
 Mastrogiovanni, S., Laghi, D., Gray, R., et al. 2023, *PhRvD*, **108**, 042002
 Mastrogiovanni, S., Leyde, K., Karathanasis, C., et al. 2021, *PhRvD*, **104**, 062009
 Mastrogiovanni, S., Leyde, K., Karathanasis, C., et al. 2022, in Proc. of the European Physical Society Conf. on High Energy Physics (Trieste: SISSA), **098**
 Messenger, C., & Read, J. 2012, *PhRvL*, **108**, 091101
 Mukherjee, S. 2022, *MNRAS*, **515**, 5495
 Mukherjee, S., Krolewski, A., Wandelt, B.D., & Silk, J. 2024, *ApJ*, **975**, 189
 Mukherjee, S., Wandelt, B. D., Nissanke, S. M., & Silvestri, A. 2021, *PhRvD*, **103**, 043520
 Ng, K. K. Y., Vitale, S., Farr, W. M., & Rodriguez, C. L. 2021, *ApJL*, **913**, L5
 Nissanke, S., Holz, D. E., Dalal, N., et al. 2013, arXiv:[1307.2638](https://arxiv.org/abs/1307.2638)
 Nissanke, S., Holz, D. E., Hughes, S. A., Dalal, N., & Sievers, J. L. 2010, *ApJ*, **725**, 496
 Nitz, A., Harry, I., Brown, D., et al. 2024, gwastro/pycbc: v2.3.3 release of PyCBC, Zenodo, doi:[10.5281/zenodo.10473621](https://doi.org/10.5281/zenodo.10473621)
 pandas development team 2020, pandas-dev/pandas: Pandas, latest, Zenodo, doi:[10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134)
 Phan, D., Pradhan, N., & Jankowiak, M. 2019, arXiv:[1912.11554](https://arxiv.org/abs/1912.11554)
 Pierra, G., Mastrogiovanni, S., Perriès, S., & Mapelli, M. 2024, *PhRvD*, **109**, 083504
 Ray, A., Hernandez, I. M., Breivik, K., & Creighton, J. 2024, arXiv:[2404.03166](https://arxiv.org/abs/2404.03166)
 Riess, A. G., Casertano, S., Yuan, W., et al. 2021, *ApJL*, **908**, L6
 Riess, A. G., Yuan, W., Casertano, S., Macri, L. M., & Scolnic, D. 2020, *ApJL*, **896**, L43
 Riess, A. G., Yuan, W., Macri, L. M., et al. 2022, *ApJL*, **934**, L7
 Rinaldi, S., & DelPozzo, W. 2021, *MNRAS*, **509**, 5454
 Schutz, B. F. 1986, *Nature*, **323**, 310
 Simpson, E. H. 1951, *Journal of the Royal Statistical Society B*, **13**, 238
 Skilling, J. 2004, in AIP Conf. Proc. 735, Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. R. Fischer (Melville, NY: AIP), **395**
 Soares-Santos, M., Palmese, A., Hartley, W., et al. 2019, *ApJL*, **876**, L7
 Taylor, S. R., Gair, J. R., & Mandel, I. 2012, *PhRvD*, **85**, 023535
 Thrane, E., & Talbot, C. 2019, *PASA*, **36**, e010
 Tiwari, V., & Fairhurst, S. 2021, *ApJL*, **913**, L19
 van Son, L.A.C., de Mink, S.E., Callister, T., et al. 2022, *ApJ*, **931**, 17
 Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
 Vitale, S., Gerosa, D., Farr, W. M., & Taylor, S. R. 2021, in Handbook of Gravitational Wave Astronomy, ed. C. Bambi et al. (Singapore: Springer),
 Ye, C. S., & Fishbach, M. 2024, *ApJ*, **967**, 62
 You, Z.-Q., Zhu, X.-J., Ashton, G., Thrane, E., & Zhu, Z.-H. 2021, *ApJ*, **908**, 215
 Zevin, M., Bavera, S. S., Berry, C. P. L., et al. 2021, *ApJ*, **910**, 152