

Measurement prospects for the pair-instability mass cutoff with gravitational waves

Matthew Mould^{1,2,3,4,*} Jack Heinzel^{2,3,4} Sofía Álvarez-López^{2,3,4}
 Cailin Plunkett^{2,3,4} Noah E. Wolfe^{2,3,4} and Salvatore Vitale^{2,3,4}

¹*Nottingham Centre of Gravity and School of Mathematical Sciences, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom*

²*LIGO Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

³*Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

⁴*Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*



(Received 11 February 2026; accepted 31 March 2026; published 11 May 2026)

Pair-instability supernovae leave behind no compact remnants, resulting in a predicted gap in the distribution of stellar black-hole masses. Gravitational waves from binary black-hole mergers probe the relevant mass range and analyses of the LIGO–Virgo–KAGRA catalog (GWTC-4) indicate a possible mass cutoff at $40 - 50 M_{\odot}$. However, the robustness of this result remains unclear. To this end, using full Bayesian parameter estimation, we simulate gravitational-wave catalogs with and without such a mass cutoff and then test whether its presence or absence is correctly inferred with parametric population models. For catalogs similar to GWTC-4, confident identification of a cutoff is not guaranteed, but the best constraints among our simulations are compatible with results from GWTC-4 when the model includes a cutoff. Conversely, spurious identification of a cutoff is unlikely. For catalogs expected by the end of the O4 observing run, uncertainty in the cutoff mass is reduced by $\gtrsim 20\%$, but a cutoff at $40 - 50 M_{\odot}$ yields only a lower bound on the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction rate, our most stringent constraints on the S factor at 300 keV being $S_{300} \gtrsim 125$ keV b at 90% credibility. Relative uncertainties on the Hubble parameter H_0 from gravitational-wave data alone can still be up to 100%. We also analyze GWTC-4 with the nonparametric PIXELPOP population model, finding that some mass features are more prominent than in parametric models but a sharp cutoff is not required. However, the parametric model passes a likelihood-based predictive test in GWTC-4 and the PIXELPOP results are consistent with those from our simulated catalogs with a cutoff. Such tests are necessary to make astrophysical claims from gravitational-wave catalogs.

DOI: [10.1103/sqy9-nnqq](https://doi.org/10.1103/sqy9-nnqq)

I. INTRODUCTION

The LIGO [1–3], Virgo [4], and KAGRA [5] (LVK) gravitational-wave (GW) detectors have made > 100 observations of the mergers of stellar-mass black holes (BHs) in compact binaries [6–10]. A fundamental open question is the origin of these systems [11–13], but astrophysical insights have been tentative [14–17]. One particular physical process that may occur in the late-stage evolution of a massive progenitor star is the pair instability [18–21]. Core contraction increases the temperature sufficiently for pair production, reducing internal pressure

support. Rebounding shocks powered by explosive nuclear burning can be so energetic as to completely disrupt the star in a pair-instability supernova (PISN) or eject material in a series of pulsational episodes with eventual collapse to a BH in a pulsational PISN [22]. The outcome is a “mass gap” from $\sim 50 - 120 M_{\odot}$ in which BHs do not form directly via stellar collapse [23–27]. Though there are candidate PISN observations, these are not confirmed (see, e.g., Refs. [28,29], and references therein). There are also significant theoretical uncertainties that impact the location of the purported mass gap [27,30–34]—most notably, nuclear reaction rates for $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ carbon–oxygen burning [25,26,35–37].

As GW observations constrain the source BH properties such as mass, they can be used to search for the imprints of (pulsational) PISNe. Following the first GW transient catalog (GWTC-1) [6], Ref. [14] found that $< 1\%$ of sources in the underlying population have masses $> 45 M_{\odot}$, potentially consistent with a PISN gap. After

*Contact author: matthew.mould@nottingham.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

GWTC-2 [7], Ref. [15] found that—while a sharp mass cutoff is not required to fit the data—only $\sim 3\%$ of systems have component masses $>45M_{\odot}$. They also found evidence of an overdensity of sources with mass $\sim 35M_{\odot}$, possibly due to pulsational PISNe [38]; later work suggested this interpretation is disfavored or offered alternative explanations [28,39–42]. Reference [16] analyzed the population following the additions of events through GWTC-3 [8,9], finding that a strongly suppressed merger rate in the PISN mass gap is not required by the GW data, reflecting the addition of higher-mass detections.

Independent analyses did find evidence for a PISN cutoff, however, starting from GWTC-2 [43,44]. Reference [45] then reanalyzed GWTC-3 with a simulation-based population model accounting for both stellar-origin BHs and repeated BH mergers [46], finding a lower edge for the PISN mass gap $\sim 40M_{\odot}$ and a positive population-level correlation between BH masses and spins. Subsequent work supported these findings with other physically motivated parametrizations [40,47] and more agnostic population models [48–53].

The recent release of GWTC-4 [10] more than doubles the number of confident binary BH detections compared to GWTC-3, allowing for sharper constraints on the population [17]. References [54,55] found evidence for a truncation or gap starting at $\sim 45M_{\odot}$ in the mass distribution of the lighter of the two BHs in merging binaries (the secondary), but not in that of the heavier (primary) BH, suggesting an astrophysical mechanism that populates the PISN gap for the preferentially heavier BH. The source of GW231123 sits above the gap as the heaviest BH merger in the LVK catalog [56], which—if its component BHs are the direct result of stellar collapse—implies the gap ends at $\sim 120M_{\odot}$. Whereas the lower edge of the gap is informed by the entire catalog, the upper edge depends on strong prior assumptions and ambiguous interpretation of GW231123 alone [57,58]: one can either constrain the upper edge of the PISN gap *assuming* the BHs of GW231123 trace stellar collapse or constrain the probability that this binary requires an alternative formation scenario *assuming* a gap location [56] but not both. Reference [59] analyzed GWTC-4 using a model with a transition in the BH spin distribution as a function of mass [51,52] rather than a separate secondary-mass truncation but similarly found a feature at $\sim 45M_{\odot}$. In these works, associating the inferred mass feature with PISNe placed constraints on the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction rate [25], albeit with uncertainties $\gtrsim 100\%$; cf. $\approx 20\%$ combined uncertainty from nuclear theory and laboratory experiments extrapolated to temperatures relevant for nuclear burning in massive stars [35,60,61].

The growing GW dataset and agreement between subsequent population analyses [62–66] lend credence to the (re)emerging interpretation that PISNe have a detectable and *detected* impact on the BH mass spectrum; however, other recent works disagree on this interpretation [67,68].

Overall, no analyses have tested the robustness of these conclusions, i.e., “*assuming such a feature really is present in the population, how well should we expect to measure it?*”, and conversely, “*is it possible to spuriously infer such a feature when it does not exist?*”. Unfortunately, there is a general lack of such validation studies for GW population results at present (though see, e.g., Refs. [69–73] for a few examples), without which any astrophysical conclusions are not solid. We tackle this issue here in the context of measuring the PISN mass cutoff from GW catalogs, serving (i) as an example of how population measurements should be tested going forward and (ii) as a projection of how well PISN physics (such as the edge of the mass gap and nuclear reaction rates) as well as cosmological parameters may be constrained by the end of the fourth LVK observing run (O4).

In Sec. II, we analyze the latest GW dataset and summarize population-level constraints, including the purported pair-instability mass cutoff. We perform a comprehensive suite of simulation studies to validate the results of these analyses in Sec. III and make projections for future catalogs in Sec. IV. Concluding discussion is presented in Sec. V.

II. OBSERVATIONAL SUMMARY

A. Data and models

We summarize results from our reanalyses of the cumulative GWTC-4 catalog [10,74] of GW detections confidently identified as binary BH mergers, following Ref. [17]. Unlike Ref. [54] who inferred both lower and upper mass-gap edges, we exclude GW231123 from our analyses owing to its uncertain interpretation and focus only on the lower edge. For comparison, we also analyze the 69 confident BH mergers through GWTC-3 [16].

We consider three phenomenological population models similar to but not the same as those in Refs. [17,54]. These are described in Appendix A and summarized in Table I, including Bayesian evidences from GWTC-4 and GWTC-3. We also use the more flexible PIXELPOP model [75–77] to test the impact of model assumptions. Though we infer the population over binary BH masses, spins, and merger redshift, here we focus on BH masses, as relevant for the PISN cutoff. The BH mass distributions inferred with each model are shown in Fig. 1.

B. Parametric results

Our phenomenological-model results confirm those of Ref. [54]. The model that allows for a truncation in secondary masses m_2 different from that in primary masses m_1 —SINGLE POWER LAW+2PEAKS2CUTOFF—is strongly preferred, according to Bayes factors $\log_{10} \mathcal{B} \approx 2.3$ over the same model but with equal m_1 and m_2 maxima—SINGLE POWER LAW+2PEAKS—and $\log_{10} \mathcal{B} \approx 2.8$ over our model most similar to the default from Ref. [17]—BROKEN POWER LAW+2PEAKS. This evidence becomes significant

TABLE I. Summary of the mass population models we consider, along with the Bayes factors $\mathcal{B}$ in favor of each phenomenological model (top three rows) relative to that which mostly closely resembles the default model of Ref. [17], BROKEN POWER LAW+2PEAKS (third row). We do not compute the Bayesian evidence for PIXELPOP as comparison between nonparametric and parametric models are not meaningful. Bayes factors are from analyses of the binary BH mergers through GWTC-4 but excluding GW231123, while those from GWTC-3 are in brackets.

Model	Description	$\log_{10} \mathcal{B}$
SINGLE POWER LAW +2PEAKS+CUTOFF	m_1 distribution consists of a power law and two Gaussians, With low-mass smoothing, free minimum and maximum. m_2 distribution is a power law with the same minimum and Low-mass smoothing as m_1 , but an independent maximum.	2.8 ± 0.2 (0.3 ± 0.1)
SINGLE POWER LAW +2PEAKS	As above, but m_1 and m_2 have the same maximum.	0.5 ± 0.2 (-0.1 ± 0.1)
BROKEN POWER LAW +2PEAKS	As above, but with a BROKEN POWER LAW in m_1 .	0.0 (0.0)
PIXELPOP	Binned $m_1 - m_2$ plane with nearest-neighbor coupling.	...

only after the addition of new observations in GWTC-4, as there is no model preference according to Bayes factors within numerical uncertainties based on GWTC-3 alone.

As seen in Fig. 1, the preferred model (top row) finds posterior median and 90% equal-tailed credible interval (CI) for the maximum m_2 of $45^{+7}_{-5} M_\odot$ from GWTC-4, much more constrained than $53^{+46}_{-13} M_\odot$ from GWTC-3. The maximum m_1 is not constrained as well but is $> 100 M_\odot$ with 99% credibility. This is in stark contrast to the two models (second and third rows of Fig. 1) that assume equal maxima for m_1 and m_2 . For example, for the SINGLE POWER LAW+2PEAKS model, the posterior for the maximum mass is peaked more strongly at $85^{+36}_{-9} M_\odot$ compared to the GWTC-3 measurement of $94^{+47}_{-20} M_\odot$, though still with a tail to higher masses in the posterior. The discrepant maxima between models highlight the impact and potential misspecification of prior model choices [57]. Though we excluded GW231123, this does not impact the inferred m_2 cutoff or the broader mass distribution (cf. Ref. [54]).

The BROKEN POWER LAW+2PEAKS model is only preferred (according to the Bayes factor) over the SINGLE POWER LAW+2PEAKS model when the m_1 and m_2 maxima are fixed as in Ref. [17] (not shown in Table I). When a BROKEN POWER LAW is included, the power law is steeper after the break at only 88% (78%) credibility for GWTC-4 (GWTC-3). Its location is not well constrained from the prior and the measurement uncertainty does not reduce from GWTC-3 to GWTC-4; cf. an $\sim 30\%$ reduction in uncertainty on the location of the peak of the mass distribution at $\sim 10 M_\odot$ between catalogs. Allowing for a BROKEN POWER LAW and independent m_1 and m_2 maxima yields a Bayes factor ($\log_{10} \mathcal{B} \approx 2.4 \pm 0.2$, also not shown in Table I) lower than the equivalent model with an unbroken power law (SINGLE POWER LAW+2PEAKS+CUTOFF).

Assuming the m_2 maximum in the SINGLE POWER LAW+2PEAKS+CUTOFF model reflects the lower edge of the PISN mass gap, we follow Refs. [25,78] to convert its posterior (but limit the BH mass range to $40 - 100 M_\odot$ to avoid errors from extrapolating their simulations, excluding $\approx 8\%$ of posterior

samples) to that for the astrophysical S factor S_{300} for the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ nuclear reaction at a temperature 300 keV—a cross section with significant influence on post main sequence stellar evolution and the pair-instability gap [35]. We find $S_{300} = 233^{+193}_{-108}$ keV b, in broad agreement with Refs. [54,59]. Because of the large uncertainty, this is also broadly consistent with independent constraints [35,60,61], though Refs. [28,67] suggest higher values for the maximum BH mass produced by pulsational PISNe may be more likely.

We also test whether or not the identification of an m_2 mass cutoff improves cosmological constraints based only on GW data using the so-called “spectral-siren” method [79,80]. Following Ref. [80], we assume a flat Λ CDM model with matter density fixed to $\Omega_m = 0.3075$ [81] and infer the Hubble parameter H_0 . For the parametric models in Table I, we find that the widths of the 90% posterior CIs are $\sim 60(80)$ kmMpc $^{-1}$ s $^{-1}$ for GWTC-4 (GWTC-3). Models allowing for the m_2 mass cutoff do have the smallest uncertainties but only by $\sim 10\%$.

C. Nonparametric results

Next, we turn to the nonparametric constraints in the bottom row of Fig. 1, using PIXELPOP to infer the BH merger rate in $m_1 - m_2$ bins [75] (see Appendix A for more details of the model and Appendix C for results including GW231123). The peaks in the m_2 distribution are more prominent compared to the parametric models [82], suggesting current default models [17] need to be revised. The distributions do not decline sharply at low and high masses, reflecting the agnostic PIXELPOP prior [75], unlike the assumption of strict minima and maxima in the parametric models. Instead, in Fig. 1 we show the 99%, 99.9%, and 99.99% percentile masses, computed from the marginal cumulative distributions for m_1 and m_2 from PIXELPOP, which from GWTC-4 (GWTC-3) are $33^{+3}_{-3} M_\odot$, $49^{+16}_{-7} M_\odot$, and $82^{+48}_{-23} M_\odot$ ($33^{+5}_{-3} M_\odot$, $52^{+45}_{-12} M_\odot$, and $91^{+50}_{-32} M_\odot$) for m_2 , respectively (see Table II later for a summary). Although the 99.9% m_2 is fairly well constrained and consistent




FIG. 1. Mass distributions of BHs in merging binaries inferred from GWTC-3 (blue) and GWTC-4 excluding GW231123 (orange) using three parametric models (top three rows) and a nonparametric model (bottom row). The masses for GW231123 (green) are included for reference. The left column shows central 90% posterior credible regions, with population constraints in terms of the probability density of logarithmic BH mass. The right column shows the posteriors of the maximum-mass parameters for the parametric models and three upper mass percentiles (99%, 99.9%, and 99.99%) for the nonparametric model. Primary (m_1) and secondary (m_2) BH masses are distinguished by filled and unfilled regions, respectively.

TABLE II. Population percentiles for secondary BH masses $m_2/M_\odot$ inferred using PIXELPOP from the LVK catalogs (top two rows) and simulated catalogs of 150 (middle three rows) or 300 (bottom three rows) O4-like events. Values are posterior medians and in brackets are widths of 90% CIs. For simulations, ranges cover ten catalog realizations.

Catalog	99%	99.9%	99.99%
GWTC-3	33 (8)	52 (57)	91 (82)
GWTC-4	33 (6)	49 (23)	82 (71)
SINGLE POWER LAW+2PEAKS+CUTOFF	31–37 (7–14)	45–50 (12–16)	58–64 (54–76)
SINGLE POWER LAW+2PEAKS	35–43 (8–20)	56–68 (13–25)	69–88 (52–67)
BROKEN POWER LAW+2PEAKS	37–40 (10–18)	55–65 (20–28)	76–87 (59–73)
SINGLE POWER LAW+2PEAKS+CUTOFF	33–37 (4–7)	43–47 (9–13)	56–61 (37–60)
SINGLE POWER LAW+2PEAKS	33–43 (8–12)	60–75 (12–19)	73–84 (31–52)
BROKEN POWER LAW+2PEAKS	36–41 (9–11)	58–68 (12–24)	70–84 (44–54)

TABLE III. Slopes (logarithmic density over logarithmic mass) of secondary-mass (m_2) distributions between 10 and $20M_\odot$ (second column), 30 and $60M_\odot$ (third column), and 30 and $150M_\odot$ (fourth column) inferred by PIXELPOP from the LVK catalogs (top two rows) and simulated catalogs of 150 (middle three rows) or 300 (bottom three rows) O4-like events. Values are posterior medians, in brackets are widths of 90% CIs, and percentages are posterior probabilities that slopes are steeper than over 10 – $20M_\odot$. For simulated catalogs, ranges cover ten catalog realizations.

Catalog	10 – $20M_\odot$	30 – $60M_\odot$	30 – $150M_\odot$
GWTC-3	3.3 (5.1)	6.6 (5.2): 94%	6.3 (7.7): 87%
GWTC-4	3.1 (4.2)	6.5 (4.1): 97%	6.3 (6.5): 93%
SINGLE POWER LAW+2PEAKS+CUTOFF	1.8–4.2 (4.9–6.4)	6.9–8.9 (7.2–8.7): 92%–97%	7.8–8.8 (8.5–11.2): 91%–98%
SINGLE POWER LAW+2PEAKS	1.4–3.9 (4.8–6.4)	3.1–5.7 (5.0–6.9): 58%–91%	6.8–7.8 (7.7–9.8): 89%–98%
BROKEN POWER LAW+2PEAKS	0.3–3.2 (4.5–6.6)	4.3–6.2 (5.2–8.0): 67%–98%	6.6–8.1 (7.6–9.9): 92%–99%
SINGLE POWER LAW+2PEAKS+CUTOFF	2.6–3.8 (4.1–5.3)	7.7–8.8 (6.3–7.4): 95%–98%	7.8–8.5 (8.2–8.7): 93%–98%
SINGLE POWER LAW+2PEAKS	2.1–3.9 (4.3–5.5)	3.4–5.2 (4.6–6.0): 60%–87%	7.0–7.9 (7.2–8.7): 89%–98%
BROKEN POWER LAW+2PEAKS	1.0–3.8 (4.3–5.5)	4.3–6.4 (4.8–5.8): 60%–94%	7.0–8.1 (7.8–8.7): 92%–99%

with the m_2 cutoff in the single power law + two peaks + cutoff model, the greater flexibility of PIXELPOP is balanced against larger uncertainties and the presence of mergers with $m_2 > 50M_\odot$ is not ruled out.

However, the PIXELPOP m_2 distribution does decline more steeply than that for m_1 and is steeper above $\sim 30M_\odot$ than over $\sim 10\text{--}20M_\odot$. The bin-to-bin gradients (logarithmic density over logarithmic mass) for $30M_\odot < m_2 < 60M_\odot$ (chosen as it has the same logarithmic width as the 10 – $20M_\odot$ interval) and $30M_\odot < m_2 < 150M_\odot$ (chosen as $150M_\odot$ is the highest PIXELPOP bin edge) are, respectively, $-6.5^{+2.0}_{-2.1}$ and $-6.3^{+3.0}_{-3.5}$ ($-6.6^{+2.5}_{-2.7}$ and $-6.3^{+3.6}_{-4.1}$) for GWTC-4 (GWTC-3). These are steeper than the slope $-3.1^{+2.0}_{-2.1}$ ($-3.3^{+2.4}_{-2.7}$) over $10M_\odot < m_2 < 20M_\odot$ at 97% and 93% (94% and 87%) posterior credibility (see Table III later for a summary). Therefore, while no strict mass cutoff is inferred with the PIXELPOP model, the number of high-mass BHs in binary mergers declines more steeply for secondary than primary components, similar to the single power law + two peaks + cutoff results.

D. Predictive checks

Given the differences between the SINGLE POWER LAW +2PEAKS+CUTOFF and PIXELPOP results, we test whether the

former is a reasonable model to begin with. We use a posterior predictive check (PPC) [83] but, unlike common approaches in GW population studies, we compare observed and predicted sources self-consistently (see Ref. [71] for a discussion). For each observed event, we select the parameter estimate with the highest likelihood¹ among its posterior samples. We compare the observed highest-likelihood secondary masses to predictions from the SINGLE POWER LAW +2PEAKS+CUTOFF model marginalized over posterior uncertainty from GWTC-4 by (i) selecting one population posterior sample, (ii) selecting a catalog (equal in size to the real one) of simulated detections from this population, (iii) selecting the highest-likelihood sample for each, and (iv) repeating many times; more details of the simulated catalogs using full Bayesian parameter estimation (PE) results from Refs. [73,84] are given in Appendix B.

The results of this check are shown in Fig. 2. We find that the observed events lie within the expected data-level cumulative distribution across (highest-likelihood) secondary BH masses, suggesting that the differences between the parametric and nonparametric results in Fig. 1 do not lead to

¹We use this term instead of “maximum likelihood” as the sample with the highest likelihood among a set of posterior samples has no guarantee of being the global maximum.




FIG. 2. Self-consistent data-level PPC showing the cumulative distribution of highest-likelihood secondary BH masses for: events in GWTC-4 (black); central 50%, 90%, and 99% posterior credible regions for simulated observations based on the SINGLE POWER LAW+2PEAKS+CUTOFF fit to GWTC-4 (darker to lighter blue); and the highest-likelihood population prediction (red). Central 50%, 90%, and 99% posterior credible intervals for the astrophysical maximum secondary mass (darker to lighter gray) are included for reference.

significantly discrepant observable properties. In fact, the mass range that sees the largest deviation is well away from a possible PISN gap, lying at $m_2 \sim 7\text{--}8M_\odot$, where the observed distribution lies above the 95th percentile of predictions; elsewhere, it lies within at most the central 90% credible region. The observed distribution also turns over around $\sim 40\text{--}50M_\odot$, where the m_2 cutoff is inferred to lie. We also perform a more stringent one-sample Kolmogorov–Smirnov test by comparing the distribution of observed highest-likelihood secondary masses to only the predicted distribution with the highest population likelihood among the set of population posterior samples. Under the null hypothesis, the observations should fall uniformly against the cumulative predicted distribution, for which the Kolmogorov–Smirnov p value is 0.97.

As a final check, we consider the event with the most massive highest-likelihood secondary BH—GW190521 [85,86] (recall GW231123 is excluded from our analyses). In Fig. 3, we compare the source masses from the original LVK PE results to our population-informed measurements [87–89] using the SINGLE POWER LAW+2PEAKS+CUTOFF and PIXELPOP models. The PIXELPOP posterior contains much more support for $m_2 > 50M_\odot$, consistent with the lack of a sharp truncation in Fig. 1. This implies that even if the secondary BH of GW190521 lies below a PISN mass gap, it is not apparent without stronger model assumptions.




FIG. 3. Posterior distributions for the source masses of GW190521 inferred from the original LVK PE results (blue) and population-informed measurements using the PIXELPOP (purple) and SINGLE POWER LAW+2PEAKS+CUTOFF (red) models. The lower left panel shows the 50%, 90%, and 99% credible regions for the joint posteriors of primary (m_1) and secondary (m_2) masses, with the gray region excluded by the definition $m_1 \geq m_2$, and the diagonal panels show the individual marginal posteriors.

III. VALIDATING GWTC-4 CONSTRAINTS

We now test whether or not: (a) the GWTC-4 constraints on the mass cutoff are as expected, including the strength of model preference, when such a cutoff is present in the underlying population; (b) a mass cutoff can be spuriously inferred from the catalog when there is no such feature; (c) a steepening decline of the primary BH mass (m_1) spectrum can be more easily confused for a secondary-mass (m_2) truncation; and (d) populations with a cutoff can be identified and distinguished from those without using PIXELPOP.

To do so, we simulate GW catalogs of 150 events with full Bayesian PE from Refs. [73,84], commensurate with the GWTC-4 catalog. We consider three different underlying populations according to the three parametric models in Table I, with parameters selected consistent with results from GWTC-4 when analyzed with each of those models (see Appendix B). We repeat each analysis on ten different catalog realizations to assess variations in results. Full details on an efficient bootstrap procedure we use to generate these catalogs by reusing previous PE results for simulated sources are given in Appendix B.

A. Do GWTC-4 results match expectations?

We start with catalogs simulated according to a SINGLE POWER LAW+2PEAKS+CUTOFF population with a mass cutoff




FIG. 4. Constraints on the mass cutoff for secondary BHs (m_2) inferred from GWTC-4 without GW231123 (black) and from simulated catalogs containing 150 O4-like GW events (blue). The true simulated values for the maximum m_2 are marked with dashed gray lines: $45M_\odot$ in the left panel and $90M_\odot$ in the middle and right panels. Each column represents a different simulated population indicated in the column titles, with ten catalog realizations for each. The top row shows posteriors on the maximum m_2 when the real and simulated catalogs are analyzed with the parametric SINGLE POWER LAW+2PEAKS+CUTOFF population model. The bottom rows show the posterior 99%, 99.9%, and 99.99% population-level m_2 percentiles inferred from each catalog with the nonparametric PIXELPOP model.

at $45M_\odot$. The posteriors on the m_2 cutoff for the ten catalogs when analyzed with that same model are shown in the left-hand panel of the top row in Fig. 4. We find that the posteriors are always consistent with the injected gap location and indeed cluster around the posterior inferred from the catalog of real data, with medians ranging from ≈ 48 – $53M_\odot$ for all but one of our simulations that has median $\approx 75M_\odot$. In all cases the posterior median overestimates the true value, which may suggest that the true cutoff is lower than $\approx 45M_\odot$ implied by inference from GWTC-4. The widths of the 90% posterior CIs are typically ≈ 15 – $30M_\odot$ but for three out of the ten catalogs are $> 60M_\odot$. Compared to a CI of $12M_\odot$ from the analysis of GWTC-4, this implies that the real measurement is consistent with only the best-constrained results among the small number of catalog realizations we consider. However, the simulations also demonstrate that a well-determined mass cutoff is not a guaranteed realization, which

only three out of our ten catalogs constrain to within $\pm 10M_\odot$.

We also analyze these same ten catalogs with the SINGLE POWER LAW+2PEAKS and BROKEN POWER LAW+2PEAKS models—neither of which allow for independent m_2 and m_1 maxima—to test whether a Bayesian model comparison correctly favors the SINGLE POWER LAW+2PEAKS+CUTOFF model from which the catalogs were simulated. None of our simulations produce Bayes factors close to the GWTC-4 result, with the correct model favored at most by $\log_{10} \mathcal{B} \approx 0.8$; cf. Table I. Four out of ten catalogs have no preference within the evidence uncertainties reported from nested sampling (± 0.1) and for one catalog the BROKEN POWER LAW+2PEAKS model is actually preferred over the true model but only by $\log_{10} \mathcal{B} \approx 0.3$. The weaker model preference compared to GWTC-4 results is in line with broader posteriors on the m_2 cutoff from our simulations.

B. Can a cutoff be spuriously inferred?

We simulate another ten catalogs, this time with sources drawn from a SINGLE POWER LAW+2PEAKS population. In particular, both m_1 and m_2 maxima are $90M_\odot$ (chosen based on the GWTC-4 constraints in Fig. 1). The m_2 cutoff posteriors inferred when these catalogs are instead analyzed with the SINGLE POWER LAW+2PEAKS+CUTOFF model allowing for independent m_1 and m_2 maxima are shown in the middle panel in the top row of Fig. 4. None of these measurements are qualitatively compatible with the analysis of GWTC-4, implying that we are unlikely to spuriously infer a PISN cutoff as low as $\sim 45M_\odot$ when it does not exist. The lowest inferred posterior median is $\approx 66M_\odot$. Note that the *marginal* one-dimensional posteriors for the maximum m_2 are systematically lower than the injected value because the prior on the m_2 maximum ensures it is always lower than that of the m_1 maximum, but the injected value is always consistent with the *joint* posterior on both maxima. In most cases, the SINGLE POWER LAW+2PEAKS is correctly preferred over the nsingle power law + two peaks + cutoff model model, though again not substantially: seven out of ten mock catalogs favor it with $0.2 \lesssim \log_{10} \mathcal{B} \lesssim 0.6$, the other three cases being undecided within the estimated evidence error (± 0.1).

C. Can a declining density appear as a cutoff?

We also test whether a more steeply declining m_1 power law above a break mass $m_1 > 34M_\odot$ (this mass chosen to be consistent with our GWTC-4 analyses; see Appendix B) is more likely to lead to spurious measurement of an m_2 cutoff, simulating ten catalogs from a BROKEN POWER LAW+2PEAKS population. The posteriors for the m_2 maximum inferred when using the incorrect SINGLE POWER LAW+2PEAKS+CUTOFF model are shown in the right-hand panel in the top row of Fig. 4. These results are similar to the SINGLE POWER LAW+2PEAKS simulations, except several are broader and skewed toward lower masses, as expected. From the catalog producing the posterior with the lowest median, we measure the m_2 maximum to be $63^{+48}_{-16}M_\odot$, which overlaps more with the GWTC-4 result than the results from the other BROKEN POWER LAW+2PEAKS catalogs and those from the SINGLE POWER LAW+2PEAKS catalogs. On the other hand, from the SINGLE POWER LAW+2PEAKS+CUTOFF simulations, there is one catalog that yields a similar m_2 cutoff measurement of $75^{+50}_{-25}M_\odot$, suggesting that an m_2 cutoff could have been more easily confused with a declining mass spectrum had GWTC-4 been realized differently. But for the actual realization, the power-law slope of the underlying population at high masses would have to be even steeper than simulated (< -4) for the two scenarios to be less distinguishable than in Fig. 4. That said, four of these catalogs do not show model preference within the estimated error, so while the misspecification we introduce does not lead to incorrectly

favoring an m_2 truncation, the catalogs also do not always confidently favor the correct model either.

D. Are nonparametric constraints as expected?

Next, we reanalyze all 30 catalogs above but with the nonparametric PIXELPOP population model. In the bottom rows of Fig. 4, we show posteriors for the same m_2 population percentiles as in Fig. 1 (see also Fig. 9 later for the population distributions). We find that the percentiles inferred from the SINGLE POWER LAW+2PEAKS+CUTOFF catalogs are most similar to those inferred from GWTC-4, particularly 99% and 99.9%, as summarized in Table II. The 99.9% m_2 values most closely recover the injected m_2 maximum of $45M_\odot$, with medians $45 - 50M_\odot$ similar to GWTC-4 ($49M_\odot$) but smaller 90% CIs $12 - 16M_\odot$ than GWTC-4 ($23M_\odot$); for the analyses of the SINGLE POWER LAW+2PEAKS and BROKEN POWER LAW+TWO peaks catalogs, the respective medians— $56 - 68M_\odot$ and $55 - 65M_\odot$ —differ from the GWTC-4 result but have similar uncertainties— $13 - 25M_\odot$ and $20 - 28M_\odot$. The constraints at the 99.99% percentile from these latter two populations are more similar to the GWTC-4 result, while those for the SINGLE POWER LAW+2PEAKS+CUTOFF catalogs are dissimilar and show the least variation, suggesting the constraints on the very tail of this m_2 distribution depend more on prior choices due to the lack of detections above the $45M_\odot$ cutoff.

As in Sec. II C, we compare the slope of the merger rate over higher and lower m_2 ranges, with results summarized in Table III. In general, the uncertainties from our simulated catalogs are larger than from GWTC-4. The posterior medians of the slopes for $10 < m_2/M_\odot < 20$ are consistent with the GWTC-4 result for all three simulated populations, though the BROKEN POWER LAW+2PEAKS model is simulated with a shallower slope in this mass range than the other models, which is correctly recovered by PIXELPOP. Over $30 < m_2/M_\odot < 60$, the median slopes inferred from the SINGLE POWER LAW+2PEAKS+CUTOFF catalogs are typically steeper than from GWTC-4, but those from the other catalogs are typically shallower than from GWTC-4. Over the broader range $30 < m_2/M_\odot < 150$, the GWTC-4 median is typically shallower than results from catalogs of all three simulated populations. From the SINGLE POWER LAW+2PEAKS+CUTOFF catalogs, PIXELPOP identifies that the m_2 slope over $30 - 60M_\odot$ is steeper than over $10 - 20M_\odot$ with posterior probabilities ($\gtrsim 92\%$) consistently more similar to the GWTC-4 result (97%) than the other simulated catalogs ($\gtrsim 60\%$); on the other hand, PIXELPOP reports posterior probabilities for a steeper slope over $30 - 150M_\odot$ than $10 - 20M_\odot$ compatible with the GWTC-4 result for all simulated catalogs.

Overall, these results suggest that models without an m_2 cutoff may be missing a declining density of sources between 30 and $60M_\odot$ but that a strict cutoff may lead to too steep a decline, in agreement with our observations in

Secs. II C and II D. In Appendix D, we assess the impact of single-event measurement uncertainty on identifying the m_2 cutoff in catalogs from the SINGLE POWER LAW+2PEAKS+CUTOFF population with PIXELPOP.

IV. PROJECTIONS FOR THE END OF O4

Moving from current to future GW observations, we project constraints that may be possible by the end of the LVK O4 observing run. We use the same ten catalog realizations as in Sec. III for each simulated population but add a further 150 events to each, roughly compatible with the total number of confident binary BH mergers that may be detected by the end of O4 [90].

A. Populations with a pair-instability cutoff

In Fig. 5, we show the posteriors on the maximum secondary mass when catalogs drawn from the SINGLE POWER LAW+2PEAKS+CUTOFF population are analyzed using that same model. In contrast to the catalogs of 150 events, by the time the catalog has again doubled in size there are no longer such ambiguously broad posteriors. Among the ten catalog realizations, the posterior medians range between 46 and $54M_\odot$ (cf. the simulated value of $45M_\odot$), with 90% CIs $10 - 20M_\odot$, reflecting reductions in uncertainty up to $\lesssim 85\%$ and $\gtrsim 20\%$.

When these catalogs are instead analyzed with the other two models in Table I, we find the correct model is always favored, in both cases now with only one out of the ten catalogs having preference $\log_{10} \mathcal{B} \approx 0.2$ within the estimated error (± 0.2). Compared to the SINGLE POWER LAW

+2PEAKS (BROKEN POWER LAW+2PEAKS) model, five (four) catalogs correctly prefer the SINGLE POWER LAW+2PEAKS+CUTOFF model with $\log_{10} \mathcal{B} \gtrsim 1$, but the largest Bayes factor $\log_{10} \mathcal{B} \approx 1.5$ out of all these analyses is still much less than in our analyses of GWTC-4.

We also show the derived posteriors on the S factor S_{300} in Fig. 5. Note that for the larger catalogs, the posterior is counterintuitively broadened, but at the same time becomes more consistent with the simulated value. This is because the edges of the mass gap are not altered much at higher $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction rates; cf. Fig. 5 of Ref. [25]. Therefore, if the pair-instability cutoff indeed falls at $\sim 45M_\odot$, future GW catalogs will likely only be able to constrain a lower bound. Our ten realizations of O4-like catalogs containing 150 events span $S_{300} > 69\text{--}104$ (68–90) keV b at 90% (99%) credibility, compared to $S_{300} > 89\text{--}126$ (71–98) keV b for our catalogs of 300 events.

B. Populations without a pair-instability cutoff

Although the results of Sec. III suggest that the measured m_2 cutoff is unlikely to be a spurious false measurement, we next consider the catalogs simulated from the SINGLE POWER LAW+2PEAKS population, for which both BH maxima are $90M_\odot$, and analyze them with (SINGLE POWER LAW+2PEAKS+CUTOFF) and without (SINGLE POWER LAW+2PEAKS) an independent m_2 truncation. For most catalogs, the latter model is correctly preferred according to the Bayes factor, with largest preference $\log_{10} \mathcal{B} \approx 0.8$, except in two cases for which it is within the estimated error (± 0.2).




FIG. 5. Posterior distributions for the maximum secondary BH mass (left panel) and the S factor for the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction at a temperate 300 keV (right panel), assuming this maximum mass is the lower edge of the PISN mass gap, inferred using the SINGLE POWER LAW+2PEAKS+CUTOFF model. Results from GWTC-4 (solid black) are used to select a simulated population for this model (dashed gray), from which ten catalogs are analyzed, each of 150 (blue) and 300 (red) O4-like events.




FIG. 6. Posterior distributions for the maximum primary (m_1) and secondary (m_2) BH masses inferred by the SINGLE POWER LAW+2PEAKS+CUTOFF model from GWTC-4 (solid black) and from ten simulated catalogs, each with 150 (blue) and 300 (red) O4-like events from the SINGLE POWER LAW+2PEAKS population that has a common maximum mass of $90M_\odot$ (dashed gray). The bottom-left panel shows the 90% posterior credible regions for the joint posteriors, with the gray region excluded by the constraint $m_1 \geq m_2$. The top-left and bottom-right panels show the individual marginal posteriors.

In Fig. 6, we show the results of the former analysis. In particular, we find that as the maximum primary and secondary BH mass are equal in the true population, the posteriors indeed tend to cluster at equal maxima (the diagonal in the bottom-left panel of Fig. 6). This is in contrast to the clearly distinct constraints from the GWTC-4 analysis. For all catalogs, we find posterior medians $\gtrsim 75M_\odot$ for the maximum secondary mass and 90% CIs up to $\sim 60\%$ narrower than the catalogs of 150 events but one that is $\sim 25\%$ broader. Nonetheless, comparing the inferred maximum m_2 between Figs. 5 and 6 confirms that a continued measurement at $\sim 45M_\odot$ is increasingly unlikely to be spurious.

We also consider the catalogs drawn from the BROKEN POWER LAW+2PEAKS population to test how well we might measure the properties of the mass distribution if the true population were closer to this model in reality. Out of the ten catalogs of 300 events, six favor the correct model with $\log_{10} \mathcal{B} \lesssim 1$, the remaining cases being undecided within the estimated error (± 0.2). In Fig. 7, we show the posteriors for the location m_b of the break in the m_1 power law, the power-law slope α_b above the break (higher m_1), and the slope α below the break (lower m_1) from analyses of GWTC-4 and simulated catalogs of 150 and 300 events. The GWTC-4 results are typically better constrained than




FIG. 7. Posterior distributions for parameters of the BROKEN POWER LAW+2PEAKS model: the location m_b of the break in the primary-mass power law, the power-law slope α_b above the break, and the slope α below the break. Results from GWTC-4 (black) are used to select a simulated population (dashed gray) from which ten catalogs with 150 (blue) and 300 (red) O4-like events are drawn and analyzed with that same model. Panels on the diagonal show marginal posteriors and other panels show 90% credible regions for two-dimensional posteriors. Equal power-law slopes $\alpha = \alpha_b$ (solid gray) are marked for reference in the middle-left panel.

our simulations of similar catalogs, especially the power-law slopes, which may suggest those constraints are driven by features in the data not accounted for by the model. Even when the model is correctly specified, we are unable to confidently constrain a BROKEN POWER LAW in catalogs of 300 events, as indicated by the broad m_b posteriors (bottom-right panel of Fig. 7)—e.g., the 90% posterior CIs have widths $\sim 20M_\odot$, compared to the prior width of $25M_\odot$ —and degeneracy between α and α_b (middle-left panel of Fig. 7)—e.g., we only find $\alpha > \alpha_b$ at posterior credibility 28%–94%.

C. Cosmology projections

For each of the three populations above, we analyze all simulated catalogs again but now also inferring the Hubble parameter H_0 as in Sec. II B. In each case, the true population has $H_0 = 67.74 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [81] and the model used for inference matches the simulated population (except its parameters are inferred). The results are shown in Fig. 8.

The posteriors from our simulated catalogs of 150 events are comparable to the GWTC-4 posteriors, with the widths of the 90% posterior CIs among each of the ten catalog realizations ranging between 45 and 89, 51 and 82, and 56 and 95 $\text{km s}^{-1} \text{ Mpc}^{-1}$ for the SINGLE POWER LAW+2PEAKS




FIG. 8. Posterior distributions for the Hubble parameter H_0 inferred using the GW-only spectral-siren method with the SINGLE POWER LAW+2PEAKS+CUTOFF (left), SINGLE POWER LAW+2PEAKS (middle), and BROKEN POWER LAW+2PEAKS (right) parametric population models. All simulated values are $H_0 = 67.74 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [81] (dashed gray). Other elements are as in previous figures.




FIG. 9. Population distributions of primary (top row) and secondary (bottom row) BH masses inferred by PIXELPOP in terms of 90% CIs for probability densities of logarithmic mass. Each column shows results for analyses of ten catalogs drawn from simulated populations (dashed gray), based on fits to GWTC-4, each containing 150 (blue) or 300 (red) O4-like events. The PIXELPOP results for GWTC-4 (solid black) are included for reference.

+CUTOFF, SINGLE POWER LAW+2PEAKS, and BROKEN POWER LAW+2PEAKS populations, respectively; cf. 56, 63, and $63 \text{ km s}^{-1} \text{ Mpc}^{-1}$ for the analyses of GWTC-4.

These uncertainties are reduced and span $38\text{--}54$, $37\text{--}53$, and $44\text{--}57 \text{ km s}^{-1} \text{ Mpc}^{-1}$, respectively, from analyses of our simulated catalogs of 300 events, suggesting a

reduction of up to $\sim 50\%$ in the H_0 uncertainty from spectral-siren analyses alone by the end of O4. The presence of an m_2 cutoff in the SINGLE POWER LAW+2PEAKS+CUTOFF catalogs does not necessarily result in better H_0 constraints than from the SINGLE POWER LAW+2PEAKS catalogs without one, but the best H_0 constraints

from both these populations are typically better than from BROKEN POWER LAW+2PEAKS catalogs. However, the relative uncertainty (ratio of posterior 90% CI to median) in H_0 can still be as large as $\mathcal{O}(100\%)$.

D. Nonparametric projections

Finally, we repeat the analyses of the three populations with the nonparametric PIXELPOP model (and fixed cosmological model). In Fig. 9, we show the population distributions of primary and secondary BH masses inferred from the simulated catalogs of 150 and 300 events and compare also to the GWTC-4 result. As in Sec. III D, the predicted PIXELPOP uncertainties are typically larger than found from GWTC-4, though the overall shape of the population is compatible with our simulations. The single power

law + two peaks and BROKEN POWER LAW+2PEAKS catalogs without an independent secondary-mass cutoff tend to overpredict the number of mergers with secondary masses above $40 - 50 M_\odot$ compared to GWTC-4. On the other hand, the m_2 cutoff in the SINGLE POWER LAW+2PEAKS+CUTOFF population can cause a local overdensity at $m_2 \sim 45 M_\odot$ in the inferred distribution that is not seen in the GWTC-4 result and leads to an underestimate of the number of mergers with secondary masses above $60 - 70 M_\odot$ compared to GWTC-4. None of our populations are simulated with a significant peak at $\sim 30 M_\odot$ in the

m_2 distribution, but such peaks can nonetheless be present in the inferred populations.

The impact of doubling the catalog size on the inferred merger rate across BH masses is quantified in Fig. 10. The widths of the 90% CIs for merger rates in each primary and secondary mass bin typically reduce by $\sim 20\%$ and up to $\lesssim 40\%$ but in some instances increase by $\lesssim 20\%$. Catalogs of 300 events simulated from the SINGLE POWER LAW+2PEAKS+CUTOFF population typically produce more consistent constraints and improvements thereof over catalogs of 150 events than when simulated from the SINGLE POWER LAW+2PEAKS and BROKEN POWER LAW+2PEAKS populations (but this may also be due to our simulation and resampling procedure described in Appendix B).

Our projected constraints on the upper m_2 percentiles are again summarized in Table II. From the SINGLE POWER LAW+2PEAKS+CUTOFF catalogs, we find 99.9% m_2 posterior medians ranging between 43 and $47 M_\odot$ with 90% CI widths $9 - 13 M_\odot$, reflecting uncertainties reduced by $\lesssim 30\%$ and again demonstrating that this percentile tracks well the simulated m_2 cutoff when present. For the single power law + two peaks (BROKEN POWER LAW+2PEAKS) model, the median 99.9% m_2 values span $60 - 75 M_\odot$ ($58 - 68 M_\odot$) with 90% CI widths $12 - 19 M_\odot$ ($12 - 24 M_\odot$), which are again clearly distinguished from the catalogs with an independent m_2 cutoff. Projections for the merger-rate slopes over m_2 result in similar conclusions




FIG. 10. Ratios of 90% CIs for binary BH merger rates in primary (top row) and secondary (bottom row) mass bins between simulated catalogs of 150 and 300 events across ten catalog realizations (blue), as inferred by PIXELPOP. Ratios below unity (gray) indicate reduced uncertainty with increased catalog size.

as in Sec. III D and are included in Table III for completeness.

V. CONCLUSIONS

The ever-growing catalog of GW observations is enabling deeper insights into the underlying source population. Readily accessible and reusable public data [74,91,92] have lowered the barrier to enter for GW population studies and proliferated independent reanalyses, allowing new insights from GW data. However, the statistical significances of these findings are rarely assessed due to the technical and computational complexities involved in simulating realistic GW catalogs for mock population studies.

Following previous work [73,84], we demonstrated the feasibility of large-scale simulation and projection studies for GW population inference using full Bayesian PE for individual sources. In particular, we successfully repurposed previous PE for simulated sources drawn from other populations [73,84] (see Appendix B), providing a tractable route for robust and comprehensive GW population simulations without additional computationally expensive PE runs or relying on simplified prescriptions for mock PE. Though we may be able to conduct such studies using mock PE [69,93] when the physics of interest is simple, like measuring a mass cutoff, this may not be true when probing the population in finer detail. However, we still neglected some physics like higher modes in the waveform model—which may impact population-level measurement of cuts in the mass spectrum [94]—and made simplifying approximations for the detection of GW signals. A pertinent future data product could therefore be PE for simulated signals in real detector noise that were successfully recovered by the search pipelines used in LVK analyses, which could be repurposed for various target populations of interest as in this work. The development of machine-learning methods for population inference may also offer a promising avenue to efficiently build more significant backgrounds [95,96].

Turning to our astrophysical results, we emphasized that interpretation of a mass gap inferred from GWTC-4 relies on prior assumptions about the source of GW231123 [56] (the heaviest BH merger observed by the LVK) but that the presence of a feature at $\sim 45M_{\odot}$ possibly associated with PISNe is collectively informed by the whole catalog. In our models (similar to Ref. [54]), this feature is the maximum secondary BH mass, but inferring the extrema of distributions can be prone to significant statistical fluctuations [57]. Our mock catalogs—simulated with the efficient procedure described above—revealed that the GWTC-4 constraints on this mass cutoff are compatible with expectations if the underlying population of secondary masses really was truncated at $45M_{\odot}$ and, furthermore, that spuriously inferring a cutoff as low as $\lesssim 50M_{\odot}$ is unlikely if it is not present in reality. In light of this constraint on the lower edge of the PISN gap at $40 - 50M_{\odot}$, it may be interesting to revisit the possible association of the peak in the distribution of BH

masses at $\approx 35M_{\odot}$ with pulsational PISNe, which was previously suggested as an unlikely origin [28,41].

On the other hand, the Bayes factor in favor of a mass cutoff in GWTC-4 was much stronger than for any of our simulated catalogs. One possible reason is that our simulated populations were limited to a maximum mass up to $90M_{\odot}$ (see Appendix B); this is compatible with the GWTC-4 constraints for the SINGLE POWER LAW+2PEAKS (see Fig. 1) and BROKEN POWER LAW+2PEAKS models, whereas the SINGLE POWER LAW+2PEAKS+CUTOFF model prefers a maximum $m_1 > 90M_{\odot}$. Our simulations are therefore conservative, in the sense that they make a lower independent m_2 cutoff more difficult to distinguish from the global m_1 cutoff. There may also be features in the population indicative of the mass gap that we did not include in our models and simulations, e.g., mass-spin correlations [45,48–51,63–66]. These features could impact the measurement of the mass gap, leading to tighter constraints from the real catalog than our simulations.

In contrast to Ref. [17], from GWTC-4 we did not find preference for more subtle structure in the distribution of BH masses, such as breaks in the power-law continuum. Our projections for catalogs similar to that which may be available by the end of O4 suggested that these will continue to be poorly constrained relative to constraints on more clear-cut features such as the mass cutoff. Similarly, our simulations for the end of O4 did not result in especially improved GW-only constraints on the Hubble parameter. However, recent work suggests the application of more flexible population models to spectral-siren analyses may improve constraints [97–99], for which we plan to test the nonparametric PIXELPOP population model in future work.

With our larger simulated catalogs, constraints on the mass cutoff inferred from the population of BH masses alone continued to improve, with spurious measurements even less probable, implying confidence in the physical origin of the real feature will increase. However, if the mass cutoff indeed falls at $40 - 50M_{\odot}$, constraining the uncertain $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction rate using current and near-future GW data alone is perhaps less promising than anticipated, as a large range of reaction rates all map to this same small range of maximum BH mass [25,26]. However, a sufficiently stringent lower bound alone may lead to tension with laboratory measurements or theoretical models [28,35,60,61,67].

Reference [68] recently suggested that evidence for the PISN cutoff or gap is overestimated by parametric population models. Our analyses of GWTC-4 using PIXELPOP also did not find a strict mass cutoff but rather a steeply declining merger rate. This, however, was consistent with analyses of catalogs simulated with a cutoff, though the inferred high-mass slopes tended to be steeper than found from the real catalog. It is possible the true population lies somewhere in between. Our nonparametric analyses also revealed that current parametric population models are likely mischaracterizing secondary-mass peaks at $\sim 10M_{\odot}$

and $\sim 30M_{\odot}$, and thus these models may need to be revised for future analyses. However, our nonparametric constraints did not indicate the presence of additional features in the distribution of BH masses $\sim 60\text{--}70M_{\odot}$, as suggested by Refs. [67,98,100].

We focused on inference of the lower edge of the pair-instability mass gap, as the constraint on this lower cutoff from GWTC-4 is much more robust than a possible upper edge [54]. Interesting future studies are therefore to quantify the overall significance of a mass-gap detection from GW catalogs, including events below, above [101,102], and within the gap; test the results of correlated BH mass–spin models [50,51]; and investigate whether origins directly tracing stellar collapse can be distinguished from other astrophysical formation processes.

ACKNOWLEDGMENTS

We thank Fabio Antonini, Simona Miller, and Hui Tong for discussions. M. M. is supported by a Research Fellowship from the Royal Commission for the Exhibition of 1851 and by the LIGO Laboratory through the National Science Foundation Award No. PHY-2309200. J. H., C. P., and N. E. W. are supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 2141064. S. A.-L. and S. V. are supported by the National Science Foundation Grant No. PHY-2045740. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This research has made use of data or software obtained from the Gravitational Wave Open Science Center, a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA, and is based upon work supported by NSF’s LIGO Laboratory, which is a major facility fully funded by the National Science Foundation.

DATA AVAILABILITY

The data that support the findings of this article are openly available [73,74,84].

APPENDIX A: MODELS

We follow the implementation of the population likelihood and regularization thereof from Ref. [17], including estimation of survey sensitivity [91,103] and use of PE samples [74,92]; we limit the relative Monte Carlo uncertainty of the likelihood estimator due to finite sampling to a maximum of unity [103–107]. Inference is performed using GWAX [108], with DYNesty [109] for nested sampling [110,111] through BILBY [112]. The mass models are summarized in Table I. The parameters of our models and the Bayesian priors we use for analyses of both real and simulated GW catalogs are given in Table IV.

Though Ref. [17] now finds dimensionless BH spin magnitudes are better modeled as independently and identically distributed (IID) according to truncated Gaussians, as in Ref. [16] and for reasons we specify in more detail in Appendix B, we instead use IID nonsingular beta distributions [113]. The polar spin tilt angles between the BH spins and binary orbital angular momenta are fit with IID truncated Gaussian distributions in cosine, with free location and scale parameters [114–116], while the azimuthal angles are assumed to be distributed uniformly. The merger rate over is modeled as evolving over redshift z according to a power law in $1+z$ [117].

Reference [17] used several models for the population of binary BH masses, taking as their default a distribution of primary BH masses m_1 with a BROKEN POWER LAW and two Gaussian peaks, alongside a single power law for the distribution of binary mass ratios $0 < q \leq 1$. The primary and secondary mass ($m_2 = qm_1$) distributions have independent minima and low-mass smoothing, but both have fixed maxima of $300M_{\odot}$. We employ different models. We assume that the m_1 and m_2 distributions have the same minimum and low-mass smoothing; we use a cubic smoothing function that results in population models with closed-form normalization constants (similar to Ref. [118]) and normalize each component individually, such that the inferred branching fractions can be interpreted directly (unlike in Ref. [17]). We allow the maximum m_1 and m_2 to be free parameters that are either equal to each other or independent, the latter case allowing for a truncation in m_2 due to PISN that does not appear in the m_1 distribution, as found by Ref. [54] (note that this is the absolute maximum of m_2 and a given binary must also satisfy $m_2 \leq m_1$). Alongside the two Gaussian peaks, we consider both single and BROKEN POWER LAW models in m_1 .

We also use PIXELPOP [75] to model the merger rate as piecewise constant over fixed two-dimensional $m_1 - m_2$ bins. We use 100 bins per dimensions logarithmically spaced between 3 and $150M_{\odot}$. A log-normal prior is set for the rate in each bin and conditionally depends only on neighboring bins, thus enforcing a minimal smoothing prior that is much more efficient than generic Gaussian process priors [52,119–122]. This allows us to assess the measurability of population-level features for BH masses without making assumptions about which features are present *a priori*, as has been successfully demonstrated on real [76] and simulated GW data [77].

APPENDIX B: SIMULATIONS

The parameter values used for our simulated populations are given in Table IV. To simulate GW catalogs, we need to draw binary BHs with source properties θ from these populations, generate GW signals $s(\theta)$ added to detector noise n , and run PE on each event with data $d = n + s(\theta)$ that passes a detection threshold (“det”) to generate a set of posterior samples Θ . Written probabilistically, the process

TABLE IV. Parameters of the three parametric models we consider, with the priors chosen for inference and the values used for simulated populations. Prior intervals indicated the ranges for uniform distributions and “Dir.” indicates a flat symmetric Dirichlet distribution. The top section includes parameters governing redshift and BH spin distributions common to all parametric models. The second section includes BH mass parameters common to all models. The third contains the primary-mass power-law slope for the SINGLE POWER LAW+2PEAKS and SINGLE POWER LAW+2PEAKS+CUTOFF models. The fourth contains the additional maximum secondary-mass parameters for the single power law + two peaks + cutoff model, such that $m_2 < \min(m_1, m_{\text{gap}})$ with $m_{\text{gap}} < m_{\text{max}}$. The fifth section contains the parameters unique to the BROKEN POWER LAW+2PEAKS model. The bottom section includes cosmological parameters for the spectral-siren population analyses.

Symbol	Description	Prior	Value
γ	z power law index	$[-10, 10]$	3
μ_a	Spin magnitude mode	$[0, 1]$	0.3
σ_a	Spin magnitude width	$[0, 1]$	0.2
μ_θ	Cosine spin tilt mode	$[-1, 1]$	1
σ_θ	Cosine spin tilt width	$[0, 4]$	1.5
μ_1	Low m_1 peak [$M_\odot$]	$[5, 15]$	10
μ_2	High m_1 peak [$M_\odot$]	$[25, 50]$	34
σ_1	Low m_1 peak width [$M_\odot$]	$[0, 10]$	0.5
σ_2	High m_1 peak width [$M_\odot$]	$[0, 10]$	4
f_0	Fraction in m_1 power law	Dir.	0.4
f_1	Fraction in low m_1 peak	Dir.	0.5
f_2	Fraction in high m_1 peak	Dir.	0.1
δ_{min}	Low-mass smoothing [$M_\odot$]	$[0, 10]$	5
m_{min}	Minimum mass [$M_\odot$]	$[3, 10]$	5
m_{max}	Maximum mass [$M_\odot$]	$[30, 150]$	90
β	m_2 power law index	$[-10, 10]$	1
α	m_1 power law index	$[-10, 10]$	-2.5
m_{gap}	Maximum m_2 [$M_\odot$]	$[30, 150]$	45
α	m_1 index below break	$[-10, 10]$	-1.5
α_b	m_1 index above break	$[-10, 10]$	-4
m_b	m_1 break mass [$M_\odot$]	$[25, 50]$	34
H_0	Hubble parameter [$\text{km s}^{-1} \text{Mpc}^{-1}$]	$[20, 200]$	67.74
Ω_m	Matter density	$= 0.3075$	0.3075

of generating the catalog $\{\Theta_n, d_n, \theta_n\}_{n=1}^N$ of PE results for detected sources can be written [105]

$$p(\{\Theta_n, d_n, \theta_n\}) = \prod_{n=1}^N p(\Theta_n | d_n) P(\text{det} | d_n) p(d_n | \theta_n) p(\theta_n), \quad (\text{B1})$$

where $p(\theta)$ is the target population, $p(d|\theta)$ is the GW likelihood [123–125], $P(\text{det}|d)$ is the (discrete) selection function, $p(\Theta|d) = \prod_i p(\Theta^i|d)$ is the PE posterior distribution corresponding to data d as represented by a finite set of independent samples $\Theta = \{\Theta^i\}$, and N is the number of detections in the catalog. The sequence of steps can be understood reading Eq. (B1) right to left.

We follow the procedures in Refs. [73,84] to perform these simulations. In summary: GW signals are generated using the IMRPhenomXP waveform model [126]; Gaussian noise for the LIGO Hanford, LIGO Livingston, and Virgo detectors is generated according to their O4

design power spectral densities [127,128]; data containing signal and noise are classed as detected if the matched-filter signal-to-noise ratio (SNR) [129] is greater than 11 [130,131] (note that this is technically an unphysical selection criterion [132] as the matched-filter SNR is computed using the true signal parameters, which in reality are not directly accessible from the data, and this also means that the population likelihood function is not self consistent, but in several previous works [73,75,77,84,133] and indeed in the present work we found that this approximation does not lead to biased results in practice); for PE, we use the same waveform model and draw posterior samples using nested sampling [110,111] as implemented in DYNesty [109] and BILBY [112], with Bayesian priors set as in Ref. [84]; we estimate the survey sensitivity using the injection distribution from Ref. [84].

Even when using likelihood acceleration techniques [134–138], large-scale PE studies such as required here are very computationally expensive. We avoid having to do new PE runs and instead reuse the posterior samples

produced in Refs. [73,84], which each simulated 1600 GW events but from different populations, say, $q_1(\theta)$ and $q_2(\theta)$. We use a mixture $q(\theta) = f q_1(\theta) + (1-f) q_2(\theta)$ as a proposal distribution from which we reweigh an already simulated catalog to our target populations of interest. The mixing fraction f is chosen such that there is an equal number of detections from the two contributing populations as this maximizes the sample efficiency of reweighed catalogs, i.e.,

$$f = \frac{q_2(\text{det})}{q_1(\text{det}) + q_2(\text{det})}, \quad (\text{B2})$$

where

$$q_i(\text{det}) = \int dd d\theta P(\text{det}|d) p(d|\theta) q_i(\theta) \quad (\text{B3})$$

is the fraction of detectable sources from the population $q_i(\theta)$. As the two population q_1 and q_2 from Refs. [73,84] are similar and the total number of detections considered in each are the same, we have $q_1(\text{det}) \approx q_2(\text{det})$ and consequently $f \approx 1 - f \approx 0.5$.

Denoting by $q(\Theta, d, \theta) = p(\Theta|d) P(\text{det}|d) p(d|\theta) q(\theta)$ the probabilistic model for a single event from the combined catalog $\{\Theta_m, d_m, \theta_m\}_{m=1}^M$ from Refs. [73,84], where $M > N$ is their total number of simulated events, we reweigh to our target populations by downselecting (without repetition) to N events with each drawn according to

$$p(\Theta, d, \theta) = q(\Theta, d, \theta) \frac{p(\theta)}{q(\theta)}. \quad (\text{B4})$$

In words, for each event in the original simulated catalog $\{\Theta_m, d_m, \theta_m\}_{m=1}^M$ drawn from population $q(\theta)$, we randomly select a subset of size N according to probabilities $\propto p(\theta_m)/q(\theta_m)$. For the data-level posterior predictive checks in Sec. II D, the only additional step is selecting the highest-likelihood sample from Θ for each event.

We verified that the population posteriors inferred from these resampled catalogs are consistent with the simulated parameters listed in Table IV. Note, however, that this resampling procedure means each of our catalogs are bootstrapped realizations instead of being independent; i.e., some of the events included in one catalog may also be included another. For example, as we have $M = 3200$ simulated events from Refs. [73,84] available, even without reweighing we would only be able to produce ten completely independent catalogs of 300 events each, which is why we only consider ten different bootstrap realizations for each of our analyses. Nonetheless, the apparent statistical variations (e.g., seen in Fig. 4) demonstrate that our simulations indeed cover a range of plausible catalog realizations and results therefrom. Similarly, the relatively small number (3200) of simulated events with PE implies that the posterior spread in the highest-likelihood predictive test in Fig. 2 may be underestimated, but since the observed distribution is already consistent with the posterior predictions, a wider spread would only make it more so.

Also note that the supports of the target populations $p(\theta)$ must be covered by the support of the proposal population $q(\theta)$. As Refs. [73,84] simulated sources with spin magnitudes drawn from nonsingular beta distributions, which have no support at zero and maximal spins, the populations we consider must also respect this constraint and therefore we cannot simulate from truncated Gaussians as used in Ref. [17], so we instead keep nonsingular beta distributions. Similarly, we cannot reweigh to spin population models from Refs. [51,52,59] as the proposal spin magnitude distribution does not produce sufficiently many sources with large positive or negative effective spin. We are also limited by the maximum BH mass simulated in Refs. [73,84], which is $\approx 90 M_\odot$.

APPENDIX C: IMPACT OF GW231123

In Fig. 11, we show the results when repeating the PIXELPOP analysis of GWTC-4 in Sec. II C but now including GW231123. The inferred primary-mass




FIG. 11. The same as the bottom row of Fig. 1, now comparing PIXELPOP analyses of GWTC-4 with (purple) and without (orange) GW231123 included in the catalog.

distribution is essentially unchanged. The secondary-mass distribution has more support at higher masses where the lighter source BH of GW231123 is inferred to lie based on LVK PE results. The population 99% m_2 mass when including GW231123 ($33^{+4}_{-3}M_\odot$) is almost unchanged from the result with GW231123 excluded ($33^{+3}_{-3}M_\odot$). The 99.9% secondary mass $54^{+46}_{-10}M_\odot$ has substantially increased uncertainty from $49^{+16}_{-7}M_\odot$, while the 99.99% value has similarly broad uncertainty as before but is shifted upward to $111^{+30}_{-39}M_\odot$ from $82^{+48}_{-23}M_\odot$.

APPENDIX D: IMPACT OF SINGLE-EVENT UNCERTAINTY

To test the impact of single-event measurement uncertainty on identifying a mass cutoff with a nonparametric population model, we repeat the PIXELPOP analyses of the 150-event catalogs from the SINGLE POWER LAW+2PEAKS+CUTOFF population but assuming the true source parameters of each event are known exactly. The inferred mass distributions are shown in Fig. 12. We also show the ratios

between the 90% CI widths of the merger rate in each mass bin when single-event uncertainty is and is not included.

Unsurprisingly, assuming point estimates for the source properties of each event tends to reduce uncertainty in the inferred population, particular at BH masses $\sim 10M_\odot$ and $30 - 40M_\odot$ coincident with identifiable features in the simulated population, where the widths of the 90% CIs of the merger rate can be reduced by $\lesssim 70\%$. For some catalogs, however, the point estimates cause larger bin-to-bin fluctuations in the inferred merger rate, especially for secondary masses near the edges of parameter space where the 90% CIs can widen by $\lesssim 70\%$ relative to the same catalogs with single-event uncertainty included. Using again the 99.9% percentile of the secondary-mass population as a proxy for the mass cutoff, we find its 90% CI can be as narrow as $8M_\odot$, compared to the smallest width of $12M_\odot$ when single-event uncertainty is included (see Table II). This suggests that it is not just the flexibility of the nonparametric PIXELPOP population model but also the measurement uncertainty inherent to GW data that impacts our ability to identify a possible mass cutoff with minimal model assumptions.




FIG. 12. Population distributions (top row) of primary (left column) and secondary (right column) BH masses in terms of the 90% CI for each mass bin inferred by PIXELPOP from ten catalogs (blue) of 150 events simulated from the SINGLE POWER LAW+2PEAKS+CUTOFF population (dashed gray) without measurement uncertainty. GWTC-4 results (black) are included for reference. The ratios of the widths of each CI when single-event uncertainty is and is not included is shown for each mass bin in the bottom row, which ratios below unity (solid gray) indicating reduced population measurement uncertainty when single-event measurement uncertainties are neglected.

- [1] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
- [2] E. Capote *et al.*, *Phys. Rev. D* **111**, 062002 (2025).
- [3] S. Soni *et al.* (LIGO Collaboration), *Classical Quantum Gravity* **42**, 085016 (2025).
- [4] F. Acernese *et al.* (Virgo Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
- [5] T. Akutsu *et al.* (KAGRA Collaboration), *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2021).
- [6] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 031040 (2019).
- [7] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **11**, 021053 (2021).
- [8] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **109**, 022001 (2024).
- [9] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), *Phys. Rev. X* **13**, 041039 (2023).
- [10] A. G. Abac *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), *arXiv:2508.18082*.
- [11] M. Mapelli, *Front. Astron. Space Sci.* **7**, 38 (2020).
- [12] M. Mapelli, Formation channels of single and binary stellar-mass black holes, in *Handbook of Gravitational Wave Astronomy*, edited by C. Bambi, S. Katsanevas, and K. D. Kokkotas (Springer, Singapore, 2020), pp. 1–65.
- [13] I. Mandel and F. S. Broekgaarden, *Living Rev. Relativity* **25**, 1 (2022).
- [14] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **882**, L24 (2019).
- [15] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **913**, L7 (2021).
- [16] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaboration), *Phys. Rev. X* **13**, 011048 (2023).
- [17] A. G. Abac *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), *arXiv:2508.18083*.
- [18] W. A. Fowler and F. Hoyle, *Astrophys. J. Suppl. Ser.* **9**, 201 (1964).
- [19] Z. Barkat, G. Rakavy, and N. Sack, *Phys. Rev. Lett.* **18**, 379 (1967).
- [20] G. Rakavy and G. Shaviv, *Astrophys. J.* **148**, 803 (1967).
- [21] G. S. Fraley, *Astrophys. Space Sci.* **2**, 96 (1968).
- [22] S. E. Woosley, *Astrophys. J.* **836**, 244 (2017).
- [23] K. Belczynski *et al.*, *Astron. Astrophys.* **594**, A97 (2016).
- [24] S. Stevenson, M. Sampson, J. Powell, A. Vigna-Gómez, C. J. Neijssel, D. Szécsi, and I. Mandel, *Astrophys. J.* **882**, 121 (2019).
- [25] R. Farmer, M. Renzo, S. de Mink, M. Fishbach, and S. Justham, *Astrophys. J. Lett.* **902**, L36 (2020).
- [26] A. K. Mehta, A. Buonanno, J. Gair, M. C. Miller, E. Farag, R. J. deBoer, M. Wiescher, and F. X. Timmes, *Astrophys. J.* **924**, 39 (2022).
- [27] S. E. Woosley and A. Heger, *Astrophys. J. Lett.* **912**, L31 (2021).
- [28] D. D. Hendriks, L. A. C. van Son, M. Renzo, R. G. Izzard, and R. Farmer, *Mon. Not. R. Astron. Soc.* **526**, 4130 (2023).
- [29] F. Gabrielli *et al.*, *Mon. Not. R. Astron. Soc.* **534**, 151 (2024).
- [30] R. Farmer, M. Renzo, S. E. de Mink, P. Marchant, and S. Justham, *Astrophys. J.* **887**, 53 (2019).
- [31] M. Mapelli, M. Spera, E. Montanari, M. Limongi, A. Chieffi, N. Giacobbo, A. Bressan, and Y. Bouffanais, *Astrophys. J.* **888**, 76 (2020).
- [32] M. Renzo, R. J. Farmer, S. Justham, S. E. de Mink, Y. Göteborg, and P. Marchant, *Mon. Not. R. Astron. Soc.* **493**, 4333 (2020).
- [33] P. Marchant and T. Moriya, *Astron. Astrophys.* **640**, L18 (2020).
- [34] G. Costa, A. Bressan, M. Mapelli, P. Marigo, G. Iorio, and M. Spera, *Mon. Not. R. Astron. Soc.* **501**, 4514 (2021).
- [35] R. J. deBoer *et al.*, *Rev. Mod. Phys.* **89**, 035007 (2017).
- [36] K. Takahashi, *Astrophys. J.* **863**, 153 (2018).
- [37] E. Farag, M. Renzo, R. Farmer, M. T. Chidester, and F. X. Timmes, *Astrophys. J.* **937**, 112 (2022).
- [38] C. Talbot and E. Thrane, *Astrophys. J.* **856**, 173 (2018).
- [39] F. Antonini, M. Gieles, F. Dosopoulou, and D. Chattopadhyay, *Mon. Not. R. Astron. Soc.* **522**, 466 (2023).
- [40] J. Golomb, M. Isi, and W. M. Farr, *Astrophys. J.* **976**, 121 (2024).
- [41] S. K. Roy, L. A. C. van Son, and W. M. Farr, *Classical Quantum Gravity* **42**, 225008 (2025).
- [42] O. Sridhar, A. Ray, and V. Kalogera, *arXiv:2511.22093*.
- [43] Y.-Z. Wang, S.-P. Tang, Y.-F. Liang, M.-Z. Han, X. Li, Z.-P. Jin, Y.-Z. Fan, and D.-M. Wei, *Astrophys. J.* **913**, 42 (2021).
- [44] E. J. Baxter, D. Croon, S. D. McDermott, and J. Sakstein, *Astrophys. J. Lett.* **916**, L16 (2021).
- [45] M. Mould, D. Gerosa, and S. R. Taylor, *Phys. Rev. D* **106**, 103013 (2022).
- [46] D. Gerosa and M. Fishbach, *Nat. Astron.* **5**, 749 (2021).
- [47] C. Karathanasis, S. Mukherjee, and S. Mastrogiovanni, *Mon. Not. R. Astron. Soc.* **523**, 4539 (2023).
- [48] Y.-Z. Wang, Y.-J. Li, J. S. Vink, Y.-Z. Fan, S.-P. Tang, Y. Qin, and D.-M. Wei, *Astrophys. J. Lett.* **941**, L39 (2022).
- [49] Y.-J. Li, Y.-Z. Wang, S.-P. Tang, and Y.-Z. Fan, *Phys. Rev. Lett.* **133**, 051401 (2024).
- [50] G. Pierra, S. Mastrogiovanni, and S. Perriès, *Astron. Astrophys.* **692**, A80 (2024).
- [51] F. Antonini, I. M. Romero-Shaw, and T. Callister, *Phys. Rev. Lett.* **134**, 011401 (2025).
- [52] F. Antonini, T. Callister, F. Dosopoulou, I. M. Romero-Shaw, and D. Chattopadhyay, *Phys. Rev. D* **112**, 063040 (2025).
- [53] I. Magaña Hernandez and A. Palmese, *arXiv:2508.19208*.
- [54] H. Tong *et al.*, *arXiv:2509.04151*.
- [55] N. Guttman, E. Payne, P. D. Lasky, and E. Thrane, *Astrophys. J.* **996**, 144 (2026).
- [56] A. G. Abac *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), *Astrophys. J. Lett.* **993**, L25 (2025).
- [57] I. Mandel, *Astrophys. J. Lett.* **996**, L4 (2026).
- [58] R. Tenorio and D. Gerosa, *arXiv:2601.02467*.
- [59] F. Antonini, I. Romero-Shaw, T. Callister, F. Dosopoulou, D. Chattopadhyay, M. Gieles, and M. Mapelli, *arXiv:2509.04637*.
- [60] Y. Shen *et al.*, *Astrophys. J.* **945**, 41 (2023).
- [61] A. M. Mukhamedzhanov, *Phys. Rev. C* **113**, 025803 (2026).
- [62] S. Banagiri, E. Thrane, and P. D. Lasky, *arXiv:2509.15646*.

- [63] H. Tong, T. A. Callister, M. Fishbach, E. Thrane, F. Antonini, S. Stevenson, I. M. Romero-Shaw, and F. Dosopoulou, [arXiv:2511.05316](#).
- [64] C. Plunkett, T. Callister, M. Zevin, and S. Vitale, [arXiv:2601.07908](#).
- [65] A. M. Farah, A. Vijaykumar, and M. Fishbach, [arXiv:2601.03456](#).
- [66] A. Vijaykumar, A. M. Farah, and M. Fishbach, *Astrophys. J. Lett.* **999**, L30 (2026).
- [67] Y.-Z. Wang, Y.-J. Li, S.-J. Gao, S.-P. Tang, and Y.-Z. Fan, [arXiv:2510.22698](#).
- [68] A. Ray and V. Kalogera, *Astrophys. J. Lett.* **998**, L20 (2026).
- [69] A. M. Farah, B. Edelman, M. Zevin, M. Fishbach, J. M. Ezquiaga, B. Farr, and D. E. Holz, *Astrophys. J.* **955**, 107 (2023).
- [70] A. Toubiana, M. L. Katz, and J. R. Gair, *Mon. Not. R. Astron. Soc.* **524**, 5844 (2023).
- [71] S. J. Miller, Z. Ko, T. Callister, and K. Chatziioannou, *Phys. Rev. D* **109**, 104036 (2024).
- [72] S. Biscoveanu, *Astrophys. J.* **996**, 95 (2026).
- [73] S. Vitale and M. Mould, *Phys. Rev. D* **112**, 083015 (2025).
- [74] A. G. Abac *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), [arXiv:2508.18079](#).
- [75] J. Heinzl, M. Mould, S. Álvarez-López, and S. Vitale, *Phys. Rev. D* **111**, 063043 (2025).
- [76] J. Heinzl, M. Mould, and S. Vitale, *Phys. Rev. D* **111**, L061305 (2025).
- [77] S. Alvarez-Lopez, J. Heinzl, M. Mould, and S. Vitale, [arXiv:2506.20731](#).
- [78] R. Farmer, M. Renzo, S. de Mink, M. Fishbach, and S. Justham, [10.5281/zenodo.4281044](#) (2020).
- [79] S. R. Taylor, J. R. Gair, and I. Mandel, *Phys. Rev. D* **85**, 023535 (2012).
- [80] A. G. Abac *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), [arXiv:2509.04348](#).
- [81] P. A. R. Ade *et al.* (Planck Collaboration), *Astron. Astrophys.* **594**, A13 (2016).
- [82] A. M. Farah, M. Fishbach, and D. E. Holz, *Astrophys. J.* **962**, 69 (2024).
- [83] I. M. Romero-Shaw, E. Thrane, and P. D. Lasky, *Pub. Astron. Soc. Aust.* **39**, e025 (2022).
- [84] N. E. Wolfe, M. Mould, J. Heinzl, and S. Vitale, [arXiv:2510.06220](#).
- [85] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **125**, 101102 (2020).
- [86] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **900**, L13 (2020).
- [87] R. Essick and M. Fishbach, On reweighing single-event posteriors with population priors, Technical Report T1900895-v2, LIGO, 2019.
- [88] T. Callister, Reweighting single event posteriors with hyperparameter marginalization, Technical Report T2100301-v3, LIGO, 2021.
- [89] C. J. Moore and D. Gerosa, *Phys. Rev. D* **104**, 083008 (2021).
- [90] <http://gracedb.ligo.org/superevents/public/O4>.
- [91] R. Essick *et al.*, *Phys. Rev. D* **112**, 102001 (2025).
- [92] G. Ashton, [arXiv:2510.11197](#).
- [93] M. Fishbach, W. M. Farr, and D. E. Holz, *Astrophys. J. Lett.* **891**, L31 (2020).
- [94] M. K. Singh, S. J. Kapadia, A. Vijaykumar, and P. Ajith, *Astrophys. J.* **971**, 23 (2024).
- [95] K. Leyde, S. R. Green, A. Toubiana, and J. Gair, *Phys. Rev. D* **109**, 064056 (2024).
- [96] J.-Q. Jiang, H.-L. Huang, J. He, Y.-T. Wang, and Y.-S. Piao, [arXiv:2505.15530](#).
- [97] H. Tong, M. Fishbach, and E. Thrane, *Astrophys. J.* **985**, 220 (2025).
- [98] G. Pierra and A. Papadopoulos, [arXiv:2601.03257](#).
- [99] M. Tagliazucchi, M. Moresco, N. Borghi, and C. Ciapetti, [arXiv:2601.03347](#).
- [100] I. Magaña Hernandez and A. Palmese, *Phys. Rev. D* **111**, 083031 (2025).
- [101] J. M. Ezquiaga and D. E. Holz, *Astrophys. J. Lett.* **909**, L23 (2021).
- [102] G. Franciolini, K. Kritos, L. Reali, F. Broekgaarden, and E. Berti, *Phys. Rev. D* **110**, 023036 (2024).
- [103] V. Tiwari, *Classical Quantum Gravity* **35**, 145009 (2018).
- [104] W. M. Farr, *Res. Not. AAS* **3**, 66 (2019).
- [105] R. Essick and W. Farr, [arXiv:2204.00461](#).
- [106] C. Talbot and J. Golomb, *Mon. Not. R. Astron. Soc.* **526**, 3495 (2023).
- [107] J. Heinzl and S. Vitale, [arXiv:2509.07221](#).
- [108] M. Mould, <http://github.com/mdmould/gwax> (2025).
- [109] J. S. Speagle, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [110] J. Skilling, *AIP Conf. Proc.* **735**, 395 (2004).
- [111] J. Skilling, *Bayesian Analysis* **1**, 833 (2006).
- [112] G. Ashton *et al.*, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [113] D. Wysocki, J. Lange, and R. O’Shaughnessy, *Phys. Rev. D* **100**, 043012 (2019).
- [114] S. Vitale, R. Lynch, R. Sturani, and P. Graff, *Classical Quantum Gravity* **34**, 03LT01 (2017).
- [115] C. Talbot and E. Thrane, *Phys. Rev. D* **96**, 023012 (2017).
- [116] S. Vitale, S. Biscoveanu, and C. Talbot, *Astron. Astrophys.* **668**, L2 (2022).
- [117] M. Fishbach, D. E. Holz, and W. M. Farr, *Astrophys. J. Lett.* **863**, L41 (2018).
- [118] V. De Renzi, F. Iacovelli, D. Gerosa, M. Mancarella, and C. Pacilio, *Phys. Rev. D* **111**, 044048 (2025).
- [119] I. Mandel, W. M. Farr, A. Colonna, S. Stevenson, P. Tiño, and J. Veitch, *Mon. Not. R. Astron. Soc.* **465**, 3254 (2017).
- [120] A. Ray, I. Magaña Hernandez, S. Mohite, J. Creighton, and S. Kapadia, *Astrophys. J.* **957**, 37 (2023).
- [121] T. A. Callister and W. M. Farr, *Phys. Rev. X* **14**, 021005 (2024).
- [122] A. M. Farah, T. A. Callister, J. M. Ezquiaga, M. Zevin, and D. E. Holz, *Astrophys. J.* **978**, 153 (2025).
- [123] E. Thrane and C. Talbot, *Pub. Astron. Soc. Aust.* **36**, e010 (2019); **37**, e036(E) (2020).
- [124] A. G. Abac *et al.* (LIGO Scientific, KAGRA, and Virgo Collaborations), *Astrophys. J. Lett.* **995**, L18 (2025).
- [125] C. Talbot *et al.*, *Classical Quantum Gravity* **42**, 235023 (2025).
- [126] G. Pratten *et al.*, *Phys. Rev. D* **103**, 104056 (2021).
- [127] B. P. Abbott *et al.* (KAGRA, LIGO Scientific, and Virgo Collaborations), *Living Rev. Relativity* **19**, 1 (2016).
- [128] LIGO–Virgo–KAGRA Collaboration, Noise curves used for Simulations in the update of the Observing Scenarios Paper, Technical Report T2000012-v2, LIGO, 2020.

- [129] A. G. Abac *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), [arXiv:2508.18081](#).
- [130] R. Essick, [Phys. Rev. D **108**, 043011 \(2023\)](#).
- [131] M. Mould, C. J. Moore, and D. Gerosa, [Phys. Rev. D **109**, 063013 \(2024\)](#).
- [132] R. Essick and M. Fishbach, [Astrophys. J. **962**, 169 \(2024\)](#).
- [133] M. Mould, N. E. Wolfe, and S. Vitale, [Phys. Rev. D **111**, 123049 \(2025\)](#).
- [134] N. J. Cornish, [arXiv:1007.4820](#).
- [135] N. J. Cornish, [Phys. Rev. D **104**, 104054 \(2021\)](#).
- [136] B. Zackay, L. Dai, and T. Venumadhav, [arXiv:1806.08792](#).
- [137] N. Leslie, L. Dai, and G. Pratten, [Phys. Rev. D **104**, 123030 \(2021\)](#).
- [138] K. Krishna, A. Vijaykumar, A. Ganguly, C. Talbot, S. Biscoveanu, R. N. George, N. Williams, and A. Zimmerman, [arXiv:2312.06009](#).