

The population of merging compact binaries inferred using gravitational waves through GWTC-3

The LIGO Scientific Collaboration, The Virgo Collaboration, and The KAGRA Scientific Collaboration
(Dated: November 5, 2021)

We report on the population properties of 76 compact binary mergers detected with gravitational waves below a false alarm rate of 1 per year through the cumulative Gravitational Wave Transient Catalog 3 (GWTC-3). The catalog contains three classes of binary mergers: binary black hole (BBH), binary neutron star (BNS), and neutron star–black hole (NSBH) mergers. We infer the BNS merger rate to be between $13 \text{ Gpc}^{-3} \text{ yr}^{-1}$ and $1900 \text{ Gpc}^{-3} \text{ yr}^{-1}$ and the NSBH merger rate to be between $7.4 \text{ Gpc}^{-3} \text{ yr}^{-1}$ and $320 \text{ Gpc}^{-3} \text{ yr}^{-1}$, assuming a constant rate density versus comoving volume and taking the union of 90% credible intervals for methods used in this work. Accounting for the BBH merger rate to evolve with redshift, we find the BBH merger rate to be between $17.3 \text{ Gpc}^{-3} \text{ yr}^{-1}$ and $45 \text{ Gpc}^{-3} \text{ yr}^{-1}$ at a fiducial redshift ($z = 0.2$). Using both binary neutron star and neutron star–black hole binaries, we obtain a broad, relatively flat neutron star mass distribution extending from $1.2_{-0.2}^{+0.1} M_{\odot}$ to $2.0_{-0.2}^{+0.3} M_{\odot}$. We can confidently identify a rapid decrease in merger rate versus component mass between neutron star-like masses and black-hole-like masses, but there is no evidence that the merger rate increases again before $10M_{\odot}$. We also find the binary black hole mass distribution has localized over- and under-densities relative to a power law distribution, with peaks emerging at chirp masses of $7.8_{-0.5}^{+0.4} M_{\odot}$ and $26.6_{-1.2}^{+1.3} M_{\odot}$. While we continue to find the mass distribution of a binary’s more massive component strongly decreases as a function of primary mass, we observe no evidence of a strongly suppressed merger rate above $\approx 60M_{\odot}$, which would highlight the presence of an upper mass gap. The rate of BBH mergers is observed to increase with redshift at a rate proportional to $(1+z)^{\kappa}$ with $\kappa = 2.7_{-1.9}^{+1.8}$ for $z \lesssim 1$. Observed black hole spins are small, with half of spin magnitudes below $\chi_i \approx 0.26$. We observe evidence of negative aligned spins in the population, and an increase in spin magnitude for systems with more unequal mass ratio. We also observe evidence of misalignment of spins relative to the orbital angular momentum.

I. INTRODUCTION

We analyze the population properties of black holes (BHs) and neutron stars (NSs) in compact binary systems using data through the end of the third observing run of LIGO–Virgo (O3). Gravitational Wave Transient Catalog 3 (GWTC-3) [1] combines observations from the first three observing runs (O1, O2 [2] and O3 [1, 3, 4]) of the Advanced LIGO [5] and Advanced Virgo [6] gravitational-wave observatories. Counting only events with false alarm rate (FAR) of $< 0.25 \text{ yr}^{-1}$, we have two binary neutron star (BNS) events, two neutron star–black hole (NSBH) events and 63 confident binary black hole (BBH) events. We distinguish between NSs and BHs using prior information about the maximum NS mass, obtained from constraints on the dense-matter equation of state [7–9]. We use the observed population of events to infer the properties of the astrophysical BNS, NSBH and BBH populations. In particular, we infer the mass and spin distributions of the NS and BH populations, the overall merger rate, and investigate their cosmological evolution.

The population includes a number of exceptional events, notably the discovery in O3 data of two NSBH binaries: GW200105_162426 and GW200115_042309 [10]. In these two systems, the primary mass m_1 is larger than the maximum mass allowed by the NS equation of state, and the secondary mass m_2 is consistent with known NS masses. Here and throughout the paper, the primary mass m_1 refers to the larger of the two component masses

in the binary, while the secondary mass m_2 refers to the smaller of the two. The inclusion of NSBH events enables the first joint analysis of the full BNS–NSBH–BBH population, including identification of sub-populations of binaries and any mass gaps between them. We also perform an analysis of the NS population properties using both BNS and NSBH systems.

The increased number of BBH observations allows for a more detailed investigation of the mass and spin distributions of BH. Overall, our new observations and results are consistent with the expectations about the mass and spin distribution of BBHs derived with our previous observations through GWTC-2 [11], which capture broad features on larger parameter scales than those emphasized in this study, and which we henceforth denote *coarse-grained features*. We demonstrate the use of non-parametric or broadly modeled methods to characterize the BBH distribution and use these to identify structure in the mass distribution. Another feature of our sample is the accumulation of more BBHs with preferentially negatively aligned spins relative to their orbital angular momentum (e.g., GW191109_010717, GW200225_060421), albeit at a significance that could occur by chance in our large catalog. Finally, the larger sample size allows for more detailed investigations of correlations between black hole masses and spins.

In this work, we adopt a high-purity set of candidate events whose selection biases we understand. Our chosen FAR threshold both ensures a sufficiently pure sample for the analyses performed in this work, particularly of the four binaries containing NS in our sample. How-

ever, due to the higher observed rate of BBH mergers, even at a less stringent threshold of $< 1 \text{ yr}^{-1}$ the relative proportion of background events remains below 10% for analyses of BBH; we therefore adopt this less stringent threshold for analyses of solely the BBH population. Both sensitivity thresholds omit several candidates of moderate significance identified in recent work, including candidates identified by our own search [1, 4], which have required the probability of an event being of astrophysical origin, $p_{\text{astro}} > 0.5$ [12]. For example, our chosen FAR threshold excludes some of the most massive events identified in GWTC-3 [1] (e.g., GW190403.051519 and GW200220.061928). We briefly discuss these events, and those identified by other groups, in the context of our reconstructed populations.

The remainder of this paper is organized as follows. In Section II we summarize observations we reported through O3, then highlight our key conclusions about them obtained in this study. In Section III we describe the hierarchical method used to fit population models to data, and to validate their results. In Section IV we describe analyses for the whole compact binary population, including both BHs and NSs. In Section V we describe our results for binaries containing one or more NSs. In Section VI and VII we describe our results for BBH masses and spins respectively. In Section VIII we discuss results obtained with other searches or selection criteria, comparing to the populations identified in this work. In Section IX we discuss the astrophysical interpretation of our observations and population inferences. In Section X we comment on prospects for future searches for the stochastic background of gravitational radiation from all compact binary mergers on our past light cone during the next observing run. We conclude in Section XI with the significance of our results. In our Appendices, we provide the details of how we estimate sensitivity to compact binary mergers (Appendix A), a comprehensive description of the population models used in this work (Appendix B), methods we used to validate our study

against prominent sources of systematic error (Appendix C), and additional details of the BBH results (Appendix D). In Appendix E, we provide revised posterior distributions for all events used in this work, each reassessed using information obtained from an estimate for the full population.

II. SUMMARY OF OBSERVATIONS AND RESULTS

A total of 90 compact binary coalescences (CBCs) have been detected in the first three observing runs [1]. The threshold used in GWTC-3 requires a probability of astrophysical origin of at least 50%. For the population analysis presented here, it is preferable to work with a different threshold to ensure lower contamination from signals of non-astrophysical origin, and to reduce the model dependence in assessing probabilities of astrophysical origin. Consequently, for the majority of analyses presented in this paper, we require a FAR of $< 0.25 \text{ yr}^{-1}$ in at least one of the search analyses in GWTC-3. This threshold limits the number of events to 67; at this threshold, we expect approximately one event not to be of astrophysical origin. For BBH focused analyses, we loosen the threshold to a FAR $< 1 \text{ yr}^{-1}$ due to the higher observed rate of BBH mergers, giving 76 events with available parameter estimates, of which approximately 4.6 are expected to be non-astrophysical. This significantly expands the number of observations since GWTC-2, which included 50 events, of which 47 had FAR of $< 1 \text{ yr}^{-1}$ and were used in our previous population analysis [11]. Table I shows selected properties of all events used to infer the astrophysical population of binary mergers in the Universe. The table contains all events with FAR $< 1 \text{ yr}^{-1}$, with less significant events having FAR between 1 yr^{-1} and 0.25 yr^{-1} which are excluded from all but the BBH analyses clearly identified. Henceforth, we abbreviate candidate names by omitting the last six digits when unambiguous.

Name	FAR _{min} (yr ⁻¹)	p_{astro}	m_1/M_{\odot}	m_2/M_{\odot}	\mathcal{M}/M_{\odot}	χ_{eff}	First appears in
GW150914	$< 1 \times 10^{-5}$	$> .99$	$35.6^{+4.7}_{-3.1}$	$30.6^{+3.0}_{-4.4}$	$28.6^{+1.7}_{-1.5}$	$-0.01^{+0.12}_{-0.13}$	[13]
GW151012	7.92×10^{-3}	$> .99$	$23.2^{+14.9}_{-5.5}$	$13.6^{+4.1}_{-4.8}$	$15.2^{+2.1}_{-1.2}$	$0.05^{+0.31}_{-0.20}$	[14]
GW151226	$< 1 \times 10^{-5}$	$> .99$	$13.7^{+8.8}_{-3.2}$	$7.7^{+2.2}_{-2.5}$	$8.9^{+0.3}_{-0.3}$	$0.18^{+0.20}_{-0.12}$	[15]
GW170104	$< 1 \times 10^{-5}$	$> .99$	$30.8^{+7.3}_{-5.6}$	$20.0^{+4.9}_{-4.6}$	$21.4^{+2.2}_{-1.8}$	$-0.04^{+0.17}_{-0.21}$	[16]
GW170608	$< 1 \times 10^{-5}$	$> .99$	$11.0^{+5.5}_{-1.7}$	$7.6^{+1.4}_{-2.2}$	$7.9^{+0.2}_{-0.2}$	$0.03^{+0.19}_{-0.07}$	[17]
GW170729	1.80×10^{-1}	0.98	$50.2^{+16.2}_{-10.2}$	$34.0^{+9.1}_{-10.1}$	$35.4^{+6.5}_{-4.8}$	$0.37^{+0.21}_{-0.25}$	[2]
GW170809	$< 1 \times 10^{-5}$	$> .99$	$35.0^{+8.3}_{-5.9}$	$23.8^{+5.1}_{-5.2}$	$24.9^{+2.1}_{-1.7}$	$0.08^{+0.17}_{-0.17}$	[2]
GW170814	$< 1 \times 10^{-5}$	$> .99$	$30.6^{+5.6}_{-3.0}$	$25.2^{+2.8}_{-4.0}$	$24.1^{+1.4}_{-1.1}$	$0.07^{+0.12}_{-0.12}$	[18]
GW170817	$< 1 \times 10^{-5}$	$> .99$	$1.46^{+0.12}_{-0.10}$	$1.27^{+0.09}_{-0.09}$	$1.186^{+0.001}_{-0.001}$	$0.00^{+0.02}_{-0.01}$	[19]
GW170818	$< 1 \times 10^{-5}$	$> .99$	$35.4^{+7.5}_{-4.7}$	$26.7^{+4.3}_{-5.2}$	$26.5^{+2.1}_{-1.7}$	$-0.09^{+0.18}_{-0.21}$	[2]
GW170823	$< 1 \times 10^{-5}$	$> .99$	$39.5^{+11.2}_{-6.7}$	$29.0^{+6.7}_{-7.8}$	$29.2^{+4.6}_{-3.6}$	$0.09^{+0.22}_{-0.26}$	[2]
GW190408.181802	$< 1 \times 10^{-5}$	$> .99$	$24.6^{+5.1}_{-3.4}$	$18.4^{+3.3}_{-3.6}$	$18.3^{+1.9}_{-1.2}$	$-0.03^{+0.14}_{-0.19}$	[3]
GW190412.053044	$< 1 \times 10^{-5}$	$> .99$	$30.1^{+4.7}_{-5.1}$	$8.3^{+1.6}_{-0.9}$	$13.3^{+0.4}_{-0.3}$	$0.25^{+0.08}_{-0.11}$	[20]

GW190413.134308	1.81×10^{-01}	0.99	$47.5^{+13.5}_{-10.7}$	$31.8^{+11.7}_{-10.8}$	$33.0^{+8.2}_{-5.4}$	$-0.03^{+0.25}_{-0.29}$	[3]
GW190421.213856	2.83×10^{-03}	> .99	$41.3^{+10.4}_{-6.9}$	$31.9^{+8.0}_{-8.8}$	$31.2^{+5.9}_{-4.2}$	$-0.06^{+0.22}_{-0.27}$	[3]
GW190425.081805	3.38×10^{-02}	0.78	$2.0^{+0.6}_{-0.3}$	$1.4^{+0.3}_{-0.3}$	$1.44^{+0.02}_{-0.02}$	$0.06^{+0.11}_{-0.05}$	[21]
GW190503.185404	$< 1 \times 10^{-5}$	> .99	$43.3^{+9.2}_{-8.1}$	$28.4^{+7.7}_{-8.0}$	$30.2^{+4.2}_{-4.2}$	$-0.03^{+0.20}_{-0.26}$	[3]
GW190512.180714	$< 1 \times 10^{-5}$	> .99	$23.3^{+5.3}_{-5.8}$	$12.6^{+3.6}_{-2.5}$	$14.6^{+1.3}_{-1.0}$	$0.03^{+0.12}_{-0.13}$	[3]
GW190513.205428	$< 1 \times 10^{-5}$	> .99	$35.7^{+9.5}_{-9.2}$	$18.0^{+7.7}_{-4.1}$	$21.6^{+3.8}_{-1.9}$	$0.11^{+0.28}_{-0.17}$	[3]
GW190517.055101	3.47×10^{-04}	> .99	$37.4^{+11.7}_{-7.6}$	$25.3^{+7.0}_{-7.3}$	$26.6^{+4.0}_{-4.0}$	$0.52^{+0.19}_{-0.19}$	[3]
GW190519.153544	$< 1 \times 10^{-5}$	> .99	$66.0^{+10.7}_{-12.0}$	$40.5^{+11.0}_{-11.1}$	$44.5^{+6.4}_{-7.1}$	$0.31^{+0.20}_{-0.22}$	[3]
GW190521.030229	$< 1 \times 10^{-5}$	> .99	$95.3^{+28.7}_{-18.9}$	$69.0^{+22.7}_{-23.1}$	$69.2^{+17.0}_{-10.6}$	$0.03^{+0.32}_{-0.39}$	[22]
GW190521.074359	1.00×10^{-02}	> .99	$42.2^{+5.9}_{-4.8}$	$32.8^{+5.4}_{-6.4}$	$32.1^{+3.2}_{-2.5}$	$0.09^{+0.10}_{-0.13}$	[22]
GW190527.092055	2.28×10^{-01}	0.85	$36.5^{+16.4}_{-9.0}$	$22.6^{+10.5}_{-8.1}$	$24.3^{+9.1}_{-4.2}$	$0.11^{+0.28}_{-0.28}$	[3]
GW190602.175927	$< 1 \times 10^{-5}$	> .99	$69.1^{+15.7}_{-13.0}$	$47.8^{+14.3}_{-17.4}$	$49.1^{+9.1}_{-8.5}$	$0.07^{+0.25}_{-0.24}$	[3]
GW190620.030421	1.12×10^{-02}	0.99	$57.1^{+16.0}_{-12.7}$	$35.5^{+12.2}_{-12.3}$	$38.3^{+8.3}_{-6.5}$	$0.33^{+0.22}_{-0.25}$	[3]
GW190630.185205	$< 1 \times 10^{-5}$	> .99	$35.1^{+6.9}_{-5.6}$	$23.7^{+5.2}_{-5.1}$	$24.9^{+2.1}_{-2.1}$	$0.10^{+0.12}_{-0.13}$	[3]
GW190701.203306	5.71×10^{-03}	0.99	$53.9^{+11.8}_{-8.0}$	$40.8^{+8.7}_{-12.0}$	$40.3^{+5.4}_{-4.9}$	$-0.07^{+0.23}_{-0.29}$	[3]
GW190706.222641	$< 1 \times 10^{-5}$	> .99	$67.0^{+14.6}_{-16.2}$	$38.2^{+14.6}_{-13.3}$	$42.7^{+10.0}_{-7.0}$	$0.28^{+0.26}_{-0.29}$	[3]
GW190707.093326	$< 1 \times 10^{-5}$	> .99	$11.6^{+3.3}_{-1.7}$	$8.4^{+1.4}_{-1.7}$	$8.5^{+0.6}_{-0.5}$	$-0.05^{+0.10}_{-0.08}$	[3]
GW190708.232457	3.09×10^{-04}	> .99	$17.6^{+4.7}_{-2.3}$	$13.2^{+2.0}_{-2.7}$	$13.2^{+0.9}_{-0.6}$	$0.02^{+0.10}_{-0.08}$	[3]
GW190720.000836	$< 1 \times 10^{-5}$	> .99	$13.4^{+6.7}_{-3.0}$	$7.8^{+2.3}_{-2.2}$	$8.9^{+0.5}_{-0.8}$	$0.18^{+0.14}_{-0.12}$	[3]
GW190727.060333	$< 1 \times 10^{-5}$	> .99	$38.0^{+9.5}_{-6.2}$	$29.4^{+7.1}_{-8.4}$	$28.6^{+5.3}_{-3.7}$	$0.11^{+0.26}_{-0.25}$	[3]
GW190728.064510	$< 1 \times 10^{-5}$	> .99	$12.3^{+7.2}_{-2.2}$	$8.1^{+1.7}_{-2.6}$	$8.6^{+0.5}_{-0.3}$	$0.12^{+0.20}_{-0.07}$	[3]
GW190803.022701	7.32×10^{-02}	0.94	$37.3^{+10.6}_{-7.0}$	$27.3^{+7.8}_{-8.2}$	$27.3^{+5.7}_{-4.1}$	$-0.03^{+0.24}_{-0.27}$	[3]
GW190814.211039	$< 1 \times 10^{-5}$	> .99	$23.2^{+11.1}_{-1.0}$	$2.59^{+0.08}_{-0.09}$	$6.09^{+0.06}_{-0.06}$	$0.00^{+0.06}_{-0.06}$	[23]
GW190828.063405	$< 1 \times 10^{-5}$	> .99	$32.1^{+5.8}_{-4.0}$	$26.2^{+4.6}_{-4.8}$	$25.0^{+3.4}_{-2.1}$	$0.19^{+0.15}_{-0.16}$	[3]
GW190828.065509	$< 1 \times 10^{-5}$	> .99	$24.1^{+7.0}_{-7.2}$	$10.2^{+3.6}_{-2.1}$	$13.3^{+1.2}_{-1.0}$	$0.08^{+0.16}_{-0.16}$	[3]
GW190910.112807	2.87×10^{-03}	> .99	$43.9^{+7.6}_{-6.1}$	$35.6^{+6.3}_{-7.2}$	$34.3^{+4.1}_{-4.1}$	$0.02^{+0.18}_{-0.18}$	[3]
GW190915.235702	$< 1 \times 10^{-5}$	> .99	$35.3^{+9.5}_{-6.4}$	$24.4^{+5.6}_{-6.1}$	$25.3^{+3.2}_{-2.7}$	$0.02^{+0.20}_{-0.25}$	[3]
GW190924.021846	$< 1 \times 10^{-5}$	> .99	$8.9^{+7.0}_{-2.0}$	$5.0^{+1.4}_{-1.9}$	$5.8^{+0.2}_{-0.2}$	$0.03^{+0.30}_{-0.09}$	[3]
GW190925.232845	7.20×10^{-03}	0.99	$21.8^{+6.1}_{-3.3}$	$15.2^{+2.7}_{-3.1}$	$15.8^{+0.8}_{-0.6}$	$0.09^{+0.1}_{-0.09}$	[4]
GW190929.012149	1.55×10^{-01}	0.87	$80.8^{+33.0}_{-33.2}$	$24.1^{+19.3}_{-10.6}$	$35.8^{+14.9}_{-8.2}$	$0.01^{+0.34}_{-0.33}$	[3]
GW190930.133541	1.23×10^{-02}	> .99	$12.3^{+12.4}_{-2.3}$	$7.8^{+1.7}_{-3.3}$	$8.5^{+0.5}_{-0.5}$	$0.14^{+0.31}_{-0.15}$	[3]
GW191105.143521	1.18×10^{-02}	> .99	$10.7^{+3.7}_{-1.6}$	$7.7^{+1.4}_{-1.9}$	$7.82^{+0.61}_{-0.45}$	$-0.02^{+0.13}_{-0.09}$	[1]
GW191109.010717	1.80×10^{-04}	> .99	65^{+11}_{-11}	47^{+15}_{-13}	$47.5^{+9.6}_{-7.5}$	$-0.29^{+0.42}_{-0.31}$	[1]
GW191127.050227	2.49×10^{-01}	0.49	53^{+47}_{-20}	24^{+17}_{-14}	$29.9^{+11.7}_{-9.1}$	$0.18^{+0.34}_{-0.36}$	[1]
GW191129.134029	$< 1 \times 10^{-5}$	> .99	$10.7^{+4.1}_{-2.1}$	$6.7^{+1.5}_{-1.7}$	$7.31^{+0.43}_{-0.28}$	$0.06^{+0.16}_{-0.08}$	[1]
GW191204.171526	$< 1 \times 10^{-5}$	> .99	$11.9^{+3.3}_{-1.8}$	$8.2^{+1.4}_{-1.6}$	$8.55^{+0.38}_{-0.27}$	$0.16^{+0.08}_{-0.05}$	[1]
GW191215.223052	$< 1 \times 10^{-5}$	> .99	$24.9^{+7.1}_{-4.1}$	$18.1^{+3.8}_{-4.1}$	$18.4^{+2.2}_{-1.7}$	$-0.04^{+0.17}_{-0.21}$	[1]
GW191216.213338	$< 1 \times 10^{-5}$	> .99	$12.1^{+4.6}_{-2.3}$	$7.7^{+1.6}_{-1.9}$	$8.33^{+0.22}_{-0.19}$	$0.11^{+0.13}_{-0.06}$	[1]
GW191222.033537	$< 1 \times 10^{-5}$	> .99	$45.1^{+10.9}_{-8.0}$	$34.7^{+9.3}_{-10.5}$	$33.8^{+7.1}_{-5.0}$	$-0.04^{+0.20}_{-0.25}$	[1]
GW191230.180458	5.02×10^{-02}	0.95	$49.4^{+14.0}_{-9.6}$	37^{+11}_{-12}	$36.5^{+8.2}_{-5.6}$	$-0.05^{+0.26}_{-0.31}$	[1]
GW200105.162426	2.04×10^{-01}	0.36	$8.9^{+1.2}_{-1.5}$	$1.9^{+0.3}_{-0.2}$	$3.41^{+0.08}_{-0.07}$	$-0.01^{+0.11}_{-0.15}$	[10]
GW200112.155838	$< 1 \times 10^{-5}$	> .99	$35.6^{+6.7}_{-4.5}$	$28.3^{+4.4}_{-5.9}$	$27.4^{+2.6}_{-2.1}$	$0.06^{+0.15}_{-0.15}$	[1]
GW200115.042309	$< 1 \times 10^{-5}$	> .99	$5.9^{+2.0}_{-2.5}$	$1.44^{+0.85}_{-0.29}$	$2.43^{+0.05}_{-0.07}$	$-0.15^{+0.24}_{-0.42}$	[10]
GW200128.022011	4.29×10^{-03}	> .99	$42.2^{+11.6}_{-8.1}$	$32.6^{+9.5}_{-9.2}$	$32.0^{+7.5}_{-5.5}$	$0.12^{+0.24}_{-0.25}$	[1]
GW200129.065458	$< 1 \times 10^{-5}$	> .99	$34.5^{+9.9}_{-3.2}$	$28.9^{+3.4}_{-9.3}$	$27.2^{+2.1}_{-2.3}$	$0.11^{+0.11}_{-0.16}$	[1]
GW200202.154313	$< 1 \times 10^{-5}$	> .99	$10.1^{+3.5}_{-1.4}$	$7.3^{+1.1}_{-1.7}$	$7.49^{+0.24}_{-0.20}$	$0.04^{+0.13}_{-0.06}$	[1]
GW200208.130117	3.11×10^{-04}	> .99	$37.8^{+9.2}_{-6.2}$	$27.4^{+6.1}_{-7.4}$	$27.7^{+3.6}_{-3.1}$	$-0.07^{+0.22}_{-0.27}$	[1]
GW200209.085452	4.64×10^{-02}	0.95	$35.6^{+10.5}_{-6.8}$	$27.1^{+7.8}_{-7.8}$	$26.7^{+6.0}_{-4.2}$	$-0.12^{+0.24}_{-0.30}$	[1]
GW200219.094415	9.94×10^{-04}	> .99	$37.5^{+10.1}_{-6.9}$	$27.9^{+7.4}_{-8.4}$	$27.6^{+5.6}_{-3.8}$	$-0.08^{+0.23}_{-0.29}$	[1]
GW200224.222234	$< 1 \times 10^{-5}$	> .99	$40.0^{+6.9}_{-4.5}$	$32.5^{+5.0}_{-7.2}$	$31.1^{+3.2}_{-2.6}$	$0.10^{+0.15}_{-0.15}$	[1]
GW200225.060421	$< 1 \times 10^{-5}$	> .99	$19.3^{+5.0}_{-3.0}$	$14.0^{+2.8}_{-3.5}$	$14.2^{+1.5}_{-1.4}$	$-0.12^{+0.17}_{-0.28}$	[1]
GW200302.015811	1.12×10^{-01}	0.91	$37.8^{+8.7}_{-8.5}$	$20.0^{+8.1}_{-5.7}$	$23.4^{+4.7}_{-3.0}$	$0.01^{+0.25}_{-0.26}$	[1]
GW200311.115853	$< 1 \times 10^{-5}$	> .99	$34.2^{+6.4}_{-3.8}$	$27.7^{+4.1}_{-5.9}$	$26.6^{+2.4}_{-2.0}$	$-0.02^{+0.16}_{-0.20}$	[1]

GW200316.215756	$< 1 \times 10^{-5}$	$> .99$	$13.1^{+10.2}_{-2.9}$	$7.8^{+1.9}_{-2.9}$	$8.75^{+0.62}_{-0.55}$	$0.13^{+0.27}_{-0.10}$	[1]
GW190413.052954	8.17×10^{-01}	0.93	$34.7^{+12.6}_{-8.1}$	$23.7^{+7.3}_{-6.7}$	$24.6^{+5.5}_{-4.1}$	$-0.01^{+0.29}_{-0.34}$	[3]
GW190426.152155	9.12×10^{-01}	0.14	$5.7^{+3.9}_{-2.3}$	$1.5^{+0.8}_{-0.5}$	$2.41^{+0.08}_{-0.08}$	$-0.03^{+0.32}_{-0.30}$	[4]
GW190719.215514	6.31×10^{-01}	0.92	$36.5^{+18.0}_{-10.3}$	$20.8^{+9.0}_{-7.2}$	$23.5^{+6.5}_{-4.0}$	$0.32^{+0.29}_{-0.31}$	[3]
GW190725.174728	4.58×10^{-01}	0.96	$12.1^{+9.6}_{-3.4}$	$6.0^{+2.4}_{-2.2}$	$7.3^{+0.6}_{-0.5}$	$-0.06^{+0.36}_{-0.22}$	[24]
GW190731.140936	3.35×10^{-01}	0.78	$41.5^{+12.2}_{-9.0}$	$28.8^{+9.7}_{-9.5}$	$29.5^{+7.1}_{-5.2}$	$0.06^{+0.24}_{-0.24}$	[3]
GW190805.211137	6.28×10^{-01}	0.95	$48.4^{+19.4}_{-12.2}$	$29.8^{+11.7}_{-9.7}$	$32.5^{+7.9}_{-5.5}$	$0.44^{+0.23}_{-0.33}$	[4]
GW190917.114630	6.56×10^{-01}	0.77	$9.1^{+3.6}_{-3.6}$	$2.2^{+1.1}_{-0.5}$	$3.7^{+0.2}_{-0.2}$	$-0.14^{+0.27}_{-0.4}$	[4]
GW191103.012549	4.58×10^{-01}	0.94	$11.8^{+6.2}_{-2.2}$	$7.9^{+1.7}_{-2.4}$	$8.34^{+0.66}_{-0.57}$	$0.21^{+0.16}_{-0.10}$	[1]
GW200216.220804	3.50×10^{-01}	0.77	51^{+22}_{-13}	30^{+14}_{-16}	$32.9^{+9.3}_{-8.5}$	$0.10^{+0.34}_{-0.36}$	[1]

TABLE I: A table of GW events that meet the criteria for inclusion in this work. Events are separated by a horizontal line into sections of $\text{FAR}_{\min} < 0.25 \text{ yr}^{-1}$ and $\text{FAR}_{\min} \geq 0.25 \text{ yr}^{-1}$ (lower), where FAR_{\min} is the smallest FAR reported over all pipelines. Within these sections, events are listed by the date they were detected. Columns provide the FAR, p_{astro} (from the pipeline with the smallest FAR), and previously-reported estimates of selected parameters. The low-significance event GW190531 is not included, lacking parameter inferences.

Fig. 1 shows the properties of the new observations included in this analysis [1]. The shaded regions show two-dimensional marginal distributions for individual events. For reference, the black contours show expected two-dimensional marginal distribution for observed BBH events deduced in our previous analysis of GWTC-2 (the POWERLAW+PEAK model from [11]). In these plots and henceforth, we define $q = m_2/m_1$ and chirp mass

$$\mathcal{M} = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}. \quad (1)$$

The dimensionless spin of each black hole is denoted $\chi_i = \mathbf{S}_i/m_i^2$ and the effective inspiral spin parameter [26]

$$\chi_{\text{eff}} = \frac{(m_1 \chi_1 + m_2 \chi_2) \cdot \hat{\mathbf{L}}}{m_1 + m_2}, \quad (2)$$

where $\hat{\mathbf{L}}$ is the instantaneous orbital angular momentum direction. Finally, z is the redshift of the event, inferred from the measured luminosity distance using $H_0 = 67.9 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.3065$ [27]. From these plots, we make several observations that motivate the investigations and results presented in the remainder of the paper.

Neutron star-black hole binaries. The two NSBH binary observations GW200105 and GW200115 [10] are apparent in Fig. 1 as two of the lowest-mass new sources. Prior to O3, gravitational wave and Galactic observations had not identified any NSBH binaries [10]. We now know that these objects exist and merge, occupying a previously unexplored region in the mass and merger rate parameter space. NSBHs form a distinct population from the BNS and most BBHs, motivating the detailed multi-component analyses pursued in Sec. IV. For the first time, we are able to present rates for BNS, NSBH and BBH inferred jointly from an analysis of all observations. The NSBH merger rate is substantially larger than the BBH merger rate. As a result, our joint analyses produce a marginal mass distribution $p(m_1)$ which

differs substantially from our previous work, and from analyses in this work based solely on BBHs: the NSBH merger rate overwhelms the BBH rate at low mass.

Lower Mass Gap. We identify a relative dearth of observations of binaries with component masses between $3M_\odot$ and $5M_\odot$. This underabundance is visible in the spectrum of observed primary masses plotted in Fig. 1. Gravitational wave and Galactic observations through O3a were consistent with a mass gap for compact objects between the heaviest NSs and the least massive BHs [28–31]. The gap was thought to extend from roughly $3M_\odot$ to $5M_\odot$, potentially due to the physics of core-collapse supernova explosions [32–36]. Both Galactic and gravitational wave observations made contemporaneously with O3 challenge this assumption [23, 37, 38]. Most notably, the secondary in GW190814 sits just above the maximum mass that the dense-matter equation of state is expected to support [23]. The primary of GW200115 [1, 10] may also lie above the maximum NS mass but below $5M_\odot$. Due to considerable uncertainty in their mass ratio, several binaries’ secondaries may also hail from this gap region between $3M_\odot$ and $5M_\odot$. We investigate the prospect of a mass gap in Sec. IV C, treating all compact objects equivalently.

NS mass distribution. The observation of NSBH binaries enables a detailed study of the observed mass distribution of NSs, combining results from both BNSs and NSBHs. We discuss this in detail in Sec. V, comparing source classifications informed by the NS equation of state (EOS) as well as the inferred location of the lower mass gap. The inferred NS mass distribution, albeit based upon a limited sample of observations, does not exhibit a peak at $1.35M_\odot$; in contrast, radio observations of Galactic BNS favor such a peak [39–41]. We investigate the impact of outliers in the mass distribution in Sec. V C, particularly GW190814 whose secondary mass lies above the otherwise inferred NS mass range.

Additional substructure in the BBH mass dis-




FIG. 1. New observations since GWTC-2. The measured properties of new CBC candidates announced since GWTC-2 with $\text{FAR} < 1/\text{yr}$ and reported parameters (blue shaded regions), compared to the expected population of detected BBHs (black contours) as inferred from past analysis of GWTC-2 with the same FAR threshold [25]. The left hand plot shows the inferred primary mass m_1 and mass ratio q ; the center plot shows the effective spin χ_{eff} and chirp mass \mathcal{M} and the right plot shows redshift z and primary mass. The least-massive sources in this sample include NSBH events GW200105 and GW200115.




FIG. 2. Illustrating substructure in the chirp mass distribution for BBH (with $\text{FAR} < 1 \text{ yr}^{-1}$, excluding GW190814, as in Sec. VI). *Top* The individual-event observations versus chirp mass (grey) and an inferred distribution of the *observed* chirp mass distribution (black solid) using an adaptive kernel density estimator [42, 43]. The kernel bandwidth is optimized for the local event density and a 90% confidence interval (black dashed) is obtained by bootstrapping [44]. *Bottom* The solid curve is the predicted chirp mass distribution obtained using the flexible mixture model framework (FM); see Sec. III for details. The distribution shows three clusters at low masses and a lack of mergers in the chirp-mass range $10 - 12 M_{\odot}$.

tribution. The observed masses of BBH binaries are clumped. This is most visible on the central panel in Fig. 1, where overdensities in the chirp mass distribution from 8 to $10M_{\odot}$ and around $30M_{\odot}$ are visible. In Fig. 2, we show the one-dimensional chirp mass distribution for BBH events. The top panel shows the observations for individual events, overlaid with the observed distribution. The observations cluster in chirp mass, with about one-eighth of observed events having chirp masses within 8– $10.5M_{\odot}$. Compared to chirp mass accuracy for these events $\lesssim 1M_{\odot}$, this region is well-separated from the next most massive binaries in chirp mass. There is also a significant overdensity at $\mathcal{M} \approx 30M_{\odot}$ and a weaker feature at $15M_{\odot}$. These features were previously identified using only GWTC-2 [45–48]. In the bottom panel of Fig. 2, we show the inferred astrophysical distribution of chirp mass, as recovered by the same FLEXIBLE MIXTURES (FM) approach that first identified these modulations [45, 49]. The features in the observed distribution are mirrored in the astrophysical one. In Section VI we show that these features are robustly identified by several independent analyses, and demonstrate that the observed structure in the mass distribution is highly significant. Since strong features correlated with chirp mass, but independent of mass ratio, are *a priori* astrophysically unlikely, these significant overdensities suggest the two-dimensional marginal distribution of the BBH population should also have significant substructure and localized overdensities. We explore this in detail in Sec. VIB

BBH Rate evolution with redshift. We find that the merger rate density increases with redshift. The right plot in Figure 1 shows the distribution of events as a function of redshift. While there is a clear evolution of the observed mass distribution with redshift, this arises from the detectors’ greater sensitivity to higher mass systems. Consequently, from Fig. 1 alone, we are not able to draw inferences about the evolution of the population or merger rate with redshift. We explore these issues in detail in Sec. VID, where we show that there is no evidence for the evolution of the mass distribution with redshift. However, the merger rate density does increase with redshift. Modeling the rate as $\propto (1+z)^{\kappa}$, we find that $\kappa = 2.7_{-1.9}^{+1.8}$. Our analysis strongly disfavors the possibility that the merger rate does not evolve with redshift.

Low BBH spins. The BBH detections exhibit effective inspiral spins concentrated about $\chi_{\text{eff}} \approx 0$, with the highest inferred spins below 0.6. The spread is consistent with expectations from GWTC-2. The events include individual candidates that probably have negative effective inspiral spin, consistent with our previous conclusion that the spin distribution contains events with $\chi_{\text{eff}} < 0$.

III. METHODS

A. Data and event selection

We consider candidate events identified by our search analyses for compact binary mergers using archival data, comprising results from the GSTLAL [50–52], PYCBC [53–58], and MBTA [59] analyses using template-based matched filtering techniques, and the CWB [60, 61] analysis using an excess-energy search that does not assume a physically parameterized signal model. Details of these analyses and the configurations used for O3 data are given in [1, 3, 4]. Out of the thousands of candidates produced, only a small minority correspond to astrophysical merger signals, most being caused by instrumental noise. While methods are emerging for performing a joint population analysis including both signal and noise events [12, 62–64], here we largely follow a simple procedure [11, 25] of imposing a significance threshold to identify events for our population analysis and implicitly treating all events passing the threshold as true signals. The choice of threshold will then limit the expected level of noise contamination.

The analyses calculate a ranking statistic for all candidate events, which is used as the basis for estimating the events’ FARs. The ranking statistic allows for sources over a broad parameter space of binary component masses and spins to be detected, without making strong assumptions on the form of the source distribution (except in the case of PYCBC-BBH, specialized for comparable-mass BBH mergers). The analyses additionally calculate an estimate of the probability of astrophysical (signal) origin, p_{astro} , using analysis-specific assumptions on the form of the signal distribution (detailed in [1, 4]). Since, in this work, we explore a range of different assumptions and models for the binary merger population, we define our event set by imposing a threshold on FAR values, rather than on p_{astro} [11].

Our searches and event validation techniques for gravitational wave transients have so far identified 76 candidates with FAR below 1 yr^{-1} in LIGO and Virgo data through O3. Table I presents these events. In our analysis here, we remove candidates with probable instrumental origin (e.g., 200219.201407 [1]). Assuming our analyses produce noise triggers independently, we expect $\sum_k \mathcal{R}T_k \simeq 4.6$ false events in our sample, where \mathcal{R} is the false alarm rate and T_k is an estimate of the time examined by the k th search. For the population studies presented here, the event list can be further restricted by additional FAR thresholds to identify a high-purity list of candidates and to assess the stability of our results to changes in threshold. The choice of FAR threshold to achieve a given level of noise contamination will depend on the number of significant event candidates (and hence, likely signals) considered for an analysis. The most prominent difference concerns analyses for binaries with one or more NS components, in Sections IV and V, as opposed to analyses which only consider BBH sys-

tems, in Sections VI and VII. While our data set contains many tens of confidently detected BBH mergers, there is only a handful of comparably significant BNS or NSBH events. This leads us to impose a more stringent threshold of $\text{FAR} < 0.25 \text{ yr}^{-1}$ for all analyses considering NS systems.

Because population reconstruction requires careful understanding of search selection biases, we do not include candidates identified by independent analyses [65–71] of the publicly released LIGO and Virgo data [72, 73]. Previous studies [63, 64] suggest that our results are unlikely to change significantly with the inclusion of these events. We similarly omit any triggers produced from our focused IMBH or eccentric binary searches [74], as we have not assessed their sensitivity to the full mass range investigated here using the consistent framework adopted for our primary results. These searches also did not yield any additional significant detections. Future analyses may be able to include events from multiple independent catalogs with a unified framework for calculating event significance independently of specific search methods [75, 76].

Parameter estimation results for each candidate event [3] were obtained using the LALINFERENCE [77], RIFT [78, 79], or BILBY [80, 81] analyses. The parameter estimation analyses use Bayesian sampling methods to produce fair draws from the posterior distribution function of the source parameters, conditioned on the data and a given model for the signal and noise [82]. Unless otherwise noted, we use previously-published samples for each event through GWTC-2.1 [2–4], as well as for GW200105 and GW200115 [10]. For previously-reported events through GWTC-2, we adopt the same parameter and event choices reported in our previous population study [11]. For O1 events, we use published samples which equally weight analyses with SEOBNRv3 [83, 84] and IMRPhenomPv2 [85] waveforms, and for new events reported in the GWTC-2 update [3], we use published samples with higher order modes, selected by equally weighting all available higher-order mode analyses (PrecessingIMRPHM). The higher-mode analyses associated with GWTC-2 do not include calibration uncertainty. Regarding new events presented in GWTC-2.1, we use the fiducial analysis reported in that work (unless otherwise noted) comprised of merged posterior samples equally drawn from SEOBNRv4PHM [86, 87] and IMRPhenomXPHM [88]. Both models implement precession and include beyond-quadrupole radiation for asymptotically quasicircular orbits. For O3b events newly reported in GWTC-3, we use a set of parameter estimation samples that are equivalent but not identical to the publicly released C01: Mixed samples from [1], which equally weigh two analyses with the models SEOBNRv4PHM [87] and IMRPhenomPv4PHM [88]. In comparison to the public GWTC-3 samples, these samples do not include final downsampling and backwards spin evolution to 20 Hz for SEOBNRv4PHM RIFT runs. Additionally, they lack the im-

pact of calibration error on the SEOBNRv4PHM analyses for GW200316, GW200129, and GW200112, and use preliminary parameter estimation runs for GW191216 and GW200208_13. Only the samples for GW191216, GW200129, and GW200208_13 exhibit noticeable differences in their measured source-frame masses, distances, or effective spin (Jensen-Shannon divergence greater than 0.005 nats) compared to their counterparts presented in [1], but these differences are only percent-level shifts in the measured parameters that are far below the statistical uncertainties, so we do not expect them to affect the population inferences herein. A more complete description of the parameter estimation methods and waveform models used can be found in Section V of [3]. To avoid ambiguity where multiple versions of these samples exist, our input posterior samples adopt the D_L^2 prior on luminosity distance D_L and have reference spins specified at 20 Hz. In the case of the BNS events GW170817 and GW190425 and the NSBH events GW200105 and GW200115, two versions of the samples are available: one that assumes component spins $\chi_{1,2} < 0.05$ for putative NS, and a less restrictive but event-dependent bound otherwise (e.g., $\chi_{1,2} < 0.99$ for GW200105 and GW200115). We use the latter (high-spin) samples here.

The transfer function between the observed strain and astrophysical strain is subject to a systematic calibration uncertainty. Our parameter inferences incorporate our best estimates of calibration uncertainty, as reported in previous work. Since calibration uncertainty has been incorporated independently for each event, we have implicitly assumed any consistent systematic bias applied to all events is small; we estimate less than 0.54% (1.74%) effect for LIGO (Virgo) respectively [89, 90]. For O3a, the amplitude uncertainty was $\lesssim 3\%$ [91]. Because we assume the secular calibration error is much smaller than the calibration error envelope applied when analyzing individual events, we do not incorporate this calibration uncertainty into our estimates of network sensitivity. In O3, this calibration uncertainty implies $\lesssim 10\%$ systematic uncertainty in the sensitive spacetime volume and the inferred merger rate, which is subdominant to our uncertainties from Poisson counting error for most source classes and mass regions.

Each foreground event in O3 has been rigorously validated [1]. Out of the 108 triggers examined in O3 (including events not included in final search results for this or our companion papers), only 4 were rejected due to the presence of instrumental noise artifacts. The number of vetoed events is comparable to or less than the expected number of false events for our fiducial analysis threshold, and far smaller than the number of events examined in this study.

B. Population analysis framework

To infer the parameters describing population models, we adopt a hierarchical Bayesian approach, in which we

marginalize over the uncertainty in our estimate of individual event parameters; see, e.g., [92–94]. Given a set of data, $\{d_i\}$, from N_{det} gravitational-wave detections, we model the total number of events as an inhomogeneous Poisson process, giving the likelihood of the data given population parameters Λ as [92, 93, 95]

$$\mathcal{L}(\{d\}, N_{\text{det}}|\Lambda, N_{\text{exp}}) \propto N^{N_{\text{det}}} e^{-N_{\text{exp}}} \prod_{i=1}^{N_{\text{det}}} \int \mathcal{L}(d_i|\theta) \pi(\theta|\Lambda) d\theta. \quad (3)$$

Here, N_{exp} is the expected number of detections over the full duration of an observation period for the population model Λ , $N = N_{\text{exp}}/\xi(\Lambda)$ is the expected number of mergers over the observation period, with $\xi(\Lambda)$ the fraction of mergers that are detectable for a population with parameters Λ . The term $\mathcal{L}(d_i|\theta)$ is the individual event likelihood for the i th event in our data set that is described by a set of parameters θ . The conditional prior $\pi(\theta|\Lambda)$ governs the population distribution on event parameters θ (e.g., the masses, spins, and redshifts) given a specific population model and set of hyperparameters Λ to describe the model. Constraining the population hyperparameters describing the distribution of gravitational-wave signals according to different models is one of the primary goals of this paper. A notable simplification results if a log-uniform prior is imposed on $N \equiv N_{\text{exp}}/\xi(\Lambda)$, the total number of events (detectable or not): one can then marginalize Eq. (3) over N to obtain [93, 94, 96]

$$\mathcal{L}(\{d\}|\Lambda) \propto \prod_{i=1}^{N_{\text{det}}} \frac{\int \mathcal{L}(d_i|\theta) \pi(\theta|\Lambda) d\theta}{\xi(\Lambda)}. \quad (4)$$

To evaluate the single-event likelihood $\mathcal{L}(d_i | \theta)$, we use posterior samples that are obtained using some default prior $\pi_{\varnothing}(\theta)$. In this case, we can calculate the integrals over the likelihood with importance sampling over the discrete samples where we denote weighted averages over posterior samples as $\langle \dots \rangle$. Equation (4), for example, becomes

$$\mathcal{L}(\{d\}|\Lambda) \propto \prod_{i=1}^{N_{\text{det}}} \frac{1}{\xi(\Lambda)} \left\langle \frac{\pi(\theta|\Lambda)}{\pi_{\varnothing}(\theta)} \right\rangle, \quad (5)$$

where the factor of $\pi_{\varnothing}(\theta)$ serves to divide out the prior used for initial parameter estimation. The likelihoods are implemented in a variety of software including GW-POPULATION [97], POPMODELS [98], SODAPOP [99], and Vamana [49]. Each code evaluates one of the likelihoods described above for population models, building a posterior with one of the EMCEE, DYNESTY, or STAN packages [100–103]. Appendix A describes how we estimate search sensitivity using synthetic sources.

In this paper, we refer to both the *astrophysical* distribution of a parameter – the version as it appears in nature – and the *observed* distribution of a parameter

– what appears in our detectors due to selection effects. The *posterior population distribution* for a given model represents our best guess for the *astrophysical* distribution of some source parameter θ , averaged over the posterior for population parameters Λ .

$$p_{\Lambda}(\theta) = \int \pi(\theta|\Lambda) p(\Lambda|\{d\}) d\Lambda. \quad (6)$$

The subscript Λ indicates that we have marginalized over population parameters. Meanwhile, the *posterior predictive distribution* refers to the population-averaged distribution of source parameters θ *conditioned on detection*.

C. Population models used in this work

In this section, we briefly summarize some of the tools and ingredients we use to generate phenomenological models $\pi(\theta|\Lambda)$ in this work. Appendix B provides a comprehensive description of the population models used in this work, including their functional form and prior assumptions.

1. Parametric mass models

Neutron star mass models: In the analyses that focus exclusively on the NS-containing events, we model the NS mass distribution as either a power law or a Gaussian with sharp minimum and maximum mass cutoffs. The latter shape is inspired by the Galactic double NS mass distribution [39–41]. In both models, which we call POWER and PEAK respectively, we assume that the components of BNSs are drawn independently from the common NS mass distribution. For NSBHs, we assume a uniform BH mass distribution and random pairing with NSs.

Fiducial population mass and redshift analysis: In the fiducial power law plus peak (POWER LAW + PEAK (PP)) model [96, 104], the mass-redshift distribution (per unit comoving volume and observer time) was assumed to be of the form $p(m_1, q, z) \propto q^\beta p(m_1)(1+z)^{\kappa-1}$, with $p(m_1)$ a mixture model containing two components: a power law with some slope and limits; and a Gaussian with some mean and variance. [In practice, this model as applied to GWTC-2 also usually included additional smoothing parameters for the upper and lower limit of the power law.] The merger rate normalization is chosen such that the source-frame merger rate per comoving volume at redshift z is given by

$$\begin{aligned} \mathcal{R}(z) &= \frac{dN}{dV_c dt}(z) \\ &= \mathcal{R}_0(1+z)^\kappa, \end{aligned} \quad (7)$$

where \mathcal{R}_0 is the local merger rate density at $z = 0$ and κ is a free parameter governing the evolution of $R(z)$ with

higher redshift. The corresponding redshift distribution of BBHs (per unit redshift interval) is [96]

$$p(z|\kappa) \propto \frac{1}{1+z} \frac{dV_c}{dz} (1+z)^\kappa, \quad (8)$$

where the leading factor of $(1+z)^{-1}$ converts time increments from the source frame to the detector frame. Past analyses generally *fixed* the redshift distribution of binaries, assuming a source-frame merger rate that is constant and uniform in comoving volume; this choice corresponds to $\kappa = 0$. Our previous population studies [11, 25] additionally considered an *evolving* merger rate with variable κ .

POWER LAW + DIP + BREAK model (PDB): To fit the distribution of BH and NS masses, we use a parameterized model described in [105] and [106], consisting of a broken power law with a notch filter. The variable depth of this notch filter allows for a dearth of events between two potential subpopulations at low and high mass. It also uses a low-pass filter at high masses to allow for a potential tapering of the mass distribution at high BH masses. The component mass distribution is then

$$p(m|\lambda) = n(m|M_{\text{low}}^{\text{gap}}, M_{\text{high}}^{\text{gap}}, A) \times l(m|m_{\text{max}}, \eta) \times \begin{cases} m^{\alpha_1} & \text{if } m < M_{\text{high}}^{\text{gap}} \\ m^{\alpha_2} & \text{if } m > M_{\text{high}}^{\text{gap}} \\ 0 & \text{if } m > m_{\text{max}} \text{ or } m < m_{\text{min}} \end{cases}. \quad (9)$$

Here, $l(m|m_{\text{max}}, \eta)$ is the low pass filter with powerlaw η applied at mass m_{max} , $n(m|M_{\text{low}}^{\text{gap}}, M_{\text{high}}^{\text{gap}}, A)$ is the notch filter with depth A applied between $M_{\text{low}}^{\text{gap}}$ and $M_{\text{high}}^{\text{gap}}$. In this model, the primary and secondary masses are fit by the same parameters and are related by a pairing function [107, 108]. Two pairing functions are considered. The first is random pairing: primary and secondary masses take independent values so long as $m_2 < m_1$. This model takes the form

$$p(m_1, m_2|\Lambda) \propto p(m = m_1|\Lambda) p(m = m_2|\Lambda) \times \Theta(m_2 < m_1), \quad (10)$$

where Θ is the Heaviside step function that enforces primary masses are greater than secondary masses and Λ is the full set of eight hyperparameters. The second is a power-law-in-mass-ratio pairing function, as in [107]. The full mass distribution in the power-law-in-mass-ratio model is thus described by

$$p(m_1, m_2|\Lambda) \propto p(m = m_1|\Lambda) p(m = m_2|\Lambda) \times q^\beta \Theta(m_2 < m_1). \quad (11)$$

Unless otherwise stated, the results from the random pairing model are presented in this work.

2. Spin models

Fiducial population spin analyses: Compact binary spins may be parameterized in several different ways. In addition to the dimensionless spin magnitudes χ_i ($i \in \{1, 2\}$)

and the polar tilt angles θ_i between each spin vector and a binary's orbital angular momentum [109], we often appeal to the *effective* spin parameters χ_{eff} and χ_{p} . The effective inspiral spin χ_{eff} characterizes a mass-averaged spin angular momentum in the direction parallel to the binaries orbital angular momentum. The effective precessing spin χ_{p} , meanwhile, corresponds approximately to the degree of *in-plane* spin, and phenomenologically parametrizes the rate of relativistic precession of the orbital plane [110]:

$$\chi_{\text{p}} = \max \left[\chi_1 \sin \theta_1, \left(\frac{3+4q}{4+3q} \right) q \chi_2 \sin \theta_2 \right]. \quad (12)$$

We leverage these two descriptions to explore the nature of BBH spins in two complementary ways. First, we use the DEFAULT spin model [111] to directly measure the distribution of BBH component spin magnitudes and tilts. We model component spin magnitudes as being independently and identically drawn from a Beta distribution [110], with

$$p(\chi_i|\alpha_\chi, \beta_\chi) \propto \chi_i^{\alpha_\chi-1} (1-\chi_i)^{\beta_\chi-1}. \quad (13)$$

Values of the shape parameters α_χ and β_χ are restricted to $\alpha_\chi > 1$ and $\beta_\chi > 1$ to ensure a *non-singular* component spin distribution. We describe component spin tilts, in turn, via a mixture between two sub-populations, one with isotropically oriented tilts and another with tilts preferentially concentrated about $\theta_i = 0$ [111]:

$$p(\cos \theta_i|\zeta, \sigma_t) = \frac{1}{2} (1-\zeta) + \zeta \mathcal{N}_{[-1,1]}(\cos \theta_i; 1, \sigma_t). \quad (14)$$

Here, $\mathcal{N}_{[-1,1]}(\cos \theta_i; 1, \sigma_t)$ is a normal distribution truncated to the interval $-1 \leq \cos \theta_i \leq 1$, centered at 1 with a standard deviation σ_t . The mixing parameter ζ governs the relative fraction of systems drawn from each sub-population. The form of Eq. (14) is motivated by a desire to capture the behavior of BBHs originating from both dynamical and isolated evolution channels, which are expected to yield preferentially isotropic and aligned spin orientations, respectively. *Perfect* spin-orbit alignment across the BBH population would correspond to $\zeta = 1$ or $\sigma_t = 0$, which our prior analysis on GWTC-2 ruled out at high confidence [11]. This default spin model is characterized by two parameters characterizing the spin magnitude distribution (e.g., α, β) and two parameters characterizing the spin misalignment mixture model (i.e., ξ, σ_t). In part because this parameterization approaches isotropy in two independent limits ($\sigma_t = 0$ or $\zeta = 1$), it assigns high prior weight to nearly-isotropic spin distributions.

Gaussian spin model: Our second approach is to instead seek to measure the distribution of effective spin parameters χ_{eff} and χ_{p} . In this case, we phenomenologically model the joint $\chi_{\text{eff}}\text{-}\chi_{\text{p}}$ distribution as a bivariate Gaussian [112, 113]:

$$p(\chi_{\text{eff}}, \chi_{\text{p}}|\mu_{\text{eff}}, \sigma_{\text{eff}}, \mu_{\text{p}}, \sigma_{\text{p}}, r) \propto \mathcal{N}(\chi_{\text{eff}}, \chi_{\text{p}}|\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (15)$$

centered at $\boldsymbol{\mu} = (\mu_{\text{eff}}, \mu_{\text{p}})$ and with a covariance matrix

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_{\text{eff}}^2 & r\sigma_{\text{eff}}\sigma_{\text{p}} \\ r\sigma_{\text{eff}}\sigma_{\text{p}} & \sigma_{\text{p}}^2 \end{pmatrix}. \quad (16)$$

Equation (15) is truncated to the intervals $-1 \leq \chi_{\text{eff}} \leq 1$ and $0 \leq \chi_{\text{p}} \leq 1$ over which the effective spin parameters are defined. This second model has five parameters for spin: two mean values and three parameters describing the covariance.

3. Multi-source mixture model

MULTI SOURCE model (MS): MS models all source categories in a mixture model, with one subpopulation for BNS, NSBH, and BBH. The BBH subpopulation follows the MULTISPIN model introduced in [11]. This model features a power law continuum $q^\beta m_1^\alpha$, plus a peak modeled as a bivariate Gaussian in m_1, m_2 . Consequently, the mass distribution is similar to the PP model. However, the spin distribution in the power law and Gaussian subpopulations are independent, as are the primary and secondary spins, with each of the four scenarios following the DEFAULT spin model, with $\zeta \equiv 1$.

New to MS are two additional bivariate Gaussian subpopulations, characterizing BNS and NSBH mergers. The BH component of NSBH follow an independent Gaussian mass distribution. As with BBH, these BH follow an independent DEFAULT spin model with $\zeta = 1$. All three types of NS (two in BNS and one in NSBH) are assumed to follow the same Gaussian mass distribution. Each type of NS follows an independent DEFAULT spin model, except here the spin magnitudes are scaled down to $\chi_{\text{max}} = 0.05$, and $\zeta \equiv 0$ since tilts are not well measured.

4. Nonparametric models

POWER LAW + SPLINE model (PS): The PS model parameterizes perturbations to a simpler phenomenological primary mass model, that is modeled as a cubic spline function.

$$p_{\text{PS}}(m_1|\Lambda, \{f_i\}) \propto p(m_1|\Lambda) \exp[f(m_1|\{f_i\})]. \quad (17)$$

Here, $f(m_1|\{f_i\})$ is the perturbation function interpolated from a set of n knots, fixed uniformly in $\log m_1$ space, and with heights $\{f_i\}$ [46]. In this work, we choose as a base model a truncated power law [11, 114] with a low mass taper, similar to our fiducial model but lacking a Gaussian peak in $p(m_1)$. This model has all the parameters of the truncated model in mass and spin, as well as an additional parameter that characterizes the low mass tapering and n more describing the heights of the cubic spline knots.

FLEXIBLE MIXTURES model (FM): Vamana, the FM model, characterizes the population as a mixture model,

summing over individually separable components describing the distribution of chirp mass, mass ratio, and $\chi_{i,z}$ [49]. Each component is composed of a Gaussian to model the chirp mass, another Gaussian to model the aligned-spin component, and a power law to model the mass ratio distribution. The weights follow a uniform prior and are proposed using a Dirichlet distribution. We choose eleven components. This choice maximizes the marginal likelihood; however, our results are robust against selecting different numbers of components.

BINNED GAUSSIAN PROCESS model (BGP): We also model the two-dimensional mass distribution as a binned Gaussian Process based on methods outlined in [115, 116]. In this approach, while the redshift and spin distribution are fixed (here, to uniform in comoving volume and isotropic and uniform in magnitude, respectively), we assume the merger rate over distinct mass bins is related via a Gaussian process that correlates the merger rates of neighboring bins. We use conventional techniques provided by PYMC3 [117] to explore the hyperparameters of the Gaussian process, in particular its covariance, to optimally reproduce our data.

IV. BINARY MERGER POPULATION ACROSS ALL MASSES

In this Section, we jointly analyze the masses of all events in Table I for several reasons. First, it allows for the inclusion of all events regardless of their inferred source type. This eliminates issues of ambiguity in source classification for a number of events in O3. Second, it makes possible the detection and characterization of additional features such as a lower mass gap between the lowest-mass objects (likely though not necessarily NS) and the more massive BH populations [105], or multiple subpopulations [116]. Third, it facilitates a self-consistent calculation of merger rates in different regions of the mass spectrum without explicitly counting the number of events in each category [12, 118]. Last, it naturally produces an overall rate of compact binary coalescences that does not require combining rates produced by disjoint models which may have differing systematics. We choose a detection threshold of $\text{FAR} < 0.25 \text{ yr}^{-1}$ which ensures even sub-populations and features driven by a few events are not contaminated by our background.

When searching for features in the population of compact binary coalescences, we want to draw robust conclusions, stable to different choices of model and approach. We, therefore, fit three independent population models, described in Sec. III. The PDB model uses a parametrized dip in the mass distribution to characterize modulations of a simple broken power law at low mass. The BGP model is a nonparametric method allowing considerable flexibility in the mass distribution, constrained only weakly by certain smoothness priors. The BGP and PDB models assume an isotropically-oriented, uniform-in-magnitude spin distribution for simplicity. For most

merging binaries and particularly those with component masses below $10M_{\odot}$, spin effects have a highly subdominant impact on our sensitivity and thus on our inferences about the compact binary merger rate distribution versus mass. By contrast, the MS model uses a multi-component mixture model, treating the mass, rate, and spin parameters of each component almost entirely independently. However, to be directly comparable to the BBH-only analyses presented in Sec. VI, our MS analysis omits the outlier event GW190814. To ensure consistent estimates of spin selection effects, the MS analysis presented here only employs O3 events; however, in Appendix C, we have demonstrated that our many analyses produce comparable results when including or excluding pre-O3 results.

A. Merger Rates

Models spanning all source classifications allow us to self-consistently measure the merger rates for all detected CBCs, both overall and subdivided into astrophysically interesting mass ranges, assuming they are independent of redshift. Moreover, because events can be classified into each category using mass limits with relatively high confidence, this approach also provides our fiducial BNS, NSBH, and BBH merger rates. Specifically, taking NS masses to lie between 1 and $2.5M_{\odot}$ and BH masses to be between 2.5 and $100M_{\odot}$ and taking the lowest 5% and highest 95% credible interval out of all three models, we infer merger rates between $13.0 \text{ Gpc}^{-3} \text{ yr}^{-1} - 1900.0 \text{ Gpc}^{-3} \text{ yr}^{-1}$ for BNS, $7.4 \text{ Gpc}^{-3} \text{ yr}^{-1} - 320.0 \text{ Gpc}^{-3} \text{ yr}^{-1}$ for NSBH, and $16.0 \text{ Gpc}^{-3} \text{ yr}^{-1} - 130.0 \text{ Gpc}^{-3} \text{ yr}^{-1}$ for BBH. Our choice of $2.5M_{\odot}$ as a boundary between BH and NS, albeit different than the nominal threshold of $3M_{\odot}$ adopted in GWTC-3, is consistent with our subsequent classification, based both on EOS and merger rate. Table II provides the rate estimates obtained for the three models used in this section and, in addition, shows rates for events in the mass gap, as discussed in detail in Section IV C.

For most categories, our merger rate estimates are consistent with previously published estimates. For example, following GWTC-2 we inferred a binary black hole merger rate to be $23.9_{-8.6}^{+14.9} \text{ Gpc}^{-3} \text{ yr}^{-1}$. Our knowledge of the coarse-grained mass spectrum has not significantly evolved since our previous analysis, and we find the inferred BBH rate is consistent with the previously reported rate, which also omitted GW190814.

For BNS, the inferred merger rate depends on the presumed mass distribution. With few observations to pin down their behavior at low mass, the three approaches adopted in this Section arrive at different compact binary mass distributions between $1M_{\odot}$ and $2.5M_{\odot}$. Because the merger rate in this region scales $\propto \langle VT \rangle^{-1} \simeq \langle \mathcal{M}^{15/6} \rangle^{-1}$ (where VT denotes the sensitive 4-volume for a specific binary and the angled brackets denote averaging over objects less than $2.5M_{\odot}$, the upper boundary

used in this Section for NS masses), the three methods used in this Section arrive at merger rates differing by factors of up to approximately three. We previously reported a BNS merger rate of $320_{-240}^{+490} \text{ Gpc}^{-3} \text{ yr}^{-1}$ [21].

For NSBH we previously inferred a merger rate of $45_{-33}^{+75} \text{ Gpc}^{-3} \text{ yr}^{-1}$ assuming the observed NSBH are representative of the population or $130_{-69}^{+112} \text{ Gpc}^{-3} \text{ yr}^{-1}$ assuming a broad NSBH population [10]. In this work, each of our joint analyses recovers and adopts different mass spectra, producing a broadly consistent rate (between $7.4 \text{ Gpc}^{-3} \text{ yr}^{-1}$ and $320.0 \text{ Gpc}^{-3} \text{ yr}^{-1}$, including systematics). Combined, our results for the NSBH and BNS merger rates highlight the important role of modeling systematics when drawing inferences about populations with few confident members.

To further highlight the impact of model systematics on inferred merger rates, in Table II, we present our deduced merger rates across the mass space using all three models presented in this section. For simplicity, we label mass bins with NS and BH based solely on a boundary at $2.5M_{\odot}$. We also provide a rate for events in the mass gap between 2.5 and $5M_{\odot}$, in a binary with either NS or BH. The bin intervals here are chosen for ease of use to roughly capture features in the mass spectrum but do not reflect our methods for event classification nor our inference on features such as the maximum NS mass or edges of any potential mass gaps.

The models used in this Section do not model the redshift evolution of the merger rate, and instead report a constant in comoving volume merger rate density [i.e. $\kappa = 0$ in Eq. (7)]. For most of the mass intervals considered, our surveys to date extend to only modest redshift, so rate evolution versus redshift can be safely neglected. However, for high-mass binary black holes, our network has cosmologically significant reach, over which the merger rate should evolve. Furthermore, as discussed in Sec. II, we observe structure in the mass distribution for black hole binaries. Therefore, in Sec. VI we provide a more detailed description of BBH merger rates, incorporating both redshift and mass dependence.

B. Identifying sub-populations of CBCs

As discussed in Sec. II, electromagnetic observations had previously suggested a mass gap between black holes and NSs. On the one hand, astrophysical EOS inferences limit nonrotating NS masses to be below the Tolman–Oppenheimer–Volkoff (TOV) mass, $M_{\text{max,TOV}} \sim 2.2 - 2.5M_{\odot}$ [8, 9, 119–121], and studies of GW170817’s remnant limit them to $\lesssim 2.3M_{\odot}$ [122–127]. On the other hand, until recently [23, 37, 38] black holes had not been observed below $\sim 5M_{\odot}$. The sparsity of observations between $\sim 2.5M_{\odot}$ and $\sim 5M_{\odot}$ suggested a potential lower mass gap [28–31].

Figure 3 shows the two-dimensional merger rate versus component masses for the three models used in this Section, as well as the results of FM model applied to BBH.

	BNS	NSBH	BBH	NS-Gap	BBH-gap	Full
	$m_1 \in [1, 2.5]M_\odot$	$m_1 \in [2.5, 50]M_\odot$	$m_1 \in [2.5, 100]M_\odot$	$m_1 \in [2.5, 5]M_\odot$	$m_1 \in [2.5, 100]M_\odot$	$m_1 \in [1, 100]M_\odot$
	$m_2 \in [1, 2.5]M_\odot$	$m_2 \in [1, 2.5]M_\odot$	$m_2 \in [2.5, 100]M_\odot$	$m_2 \in [1, 2.5]M_\odot$	$m_2 \in [2.5, 5]M_\odot$	$m_2 \in [1, 100]M_\odot$
PDB (pair)	960^{+1700}_{-700}	59^{+81}_{-38}	25^{+10}_{-7}	41^{+69}_{-30}	$9.3^{+19.0}_{-7.6}$	1100^{+1700}_{-750}
PDB (ind)	250^{+640}_{-200}	170^{+150}_{-89}	22^{+9}_{-6}	29^{+55}_{-23}	10^{+15}_{-8}	470^{+830}_{-300}
MS	470^{+1400}_{-410}	57^{+120}_{-42}	42^{+88}_{-20}	$3.7^{+20}_{-3.4}$	$0.17^{+56}_{-0.16}$	650^{+1600}_{-460}
BGP	99^{+260}_{-86}	32^{+62}_{-25}	33^{+16}_{-10}	$2.1^{+33}_{-2.1}$	$5.1^{+12}_{-4.0}$	180^{+260}_{-110}
MERGED	13 – 1900	7.4 – 320	16 – 130	0.029 – 84	0.01 – 56	71 – 2200

TABLE II. Merger rates in $\text{Gpc}^{-3} \text{yr}^{-1}$ for the various mass bins, assuming merger rates per unit comoving volume are redshift-independent. BNS, NSBH and BBH regions are based solely upon component masses, with the split between NS and BH taken to be $2.5M_\odot$. We also provide rates for binaries with one component in the purported mass gap between $2.5M_\odot$ and $5M_\odot$. For all but the last row, merger rates are quoted at the 90% credible interval. For the last row, we provide the union of 90% credible intervals for the preceding three rows, as our most conservative realistic estimate of the merger rate for each class accounting for model systematics. The PDB (pair) model is distinct from the other three models due to its use of a pairing function [107] and is therefore excluded from the union of credible intervals in the final row. In Sec. VI we estimate the merger rate for BBH alone, accounting for variation in merger rate versus redshift.

This representation emphasizes the importance of asymmetric binaries to the overall merger rate $d\mathcal{R}/dm_1$ for masses between $1M_\odot$ and $10M_\odot$. The inferred merger rates further illustrate a falloff in event rate at masses above the BNS scale, with additional peaks associated with both unequal mass binaries consistent with NSBH systems as well as approximately equal mass BBH binaries. The rate of events with at least one component between $2.5\text{--}5M_\odot$ (i.e. in the purported mass gap) is constrained to be lower than the rate of BNS-like events, but is consistent with the rate of BBH-like events. As further emphasis, Fig. 4 shows the merger rate versus mass for all binaries and also restricting to binaries with $q \simeq 1$ (e.g., the diagonal bins in the BGP model). The rate for approximately equal mass binaries is significantly lower. In other words, because asymmetric mergers like NSBH occur at a much higher rate than BBH but a much lower rate than BNS, in a joint analysis they significantly impact the marginal merger rate $d\mathcal{R}/dm_1$ at the lowest masses.

This result highlights another feature: the compact binary population has (at least) three dominant populations: BNS-like systems; significantly asymmetric binaries with small m_2 , comparable to the typical masses of NSs (i.e., including NSBHs as well as GW190814); and the main BBH population with q preferentially more symmetric than $1/4$ (i.e., including GW190412 but not GW190814).

For binaries containing lower mass gap-scale objects, our inferences about the merger rate and its dependence on mass are consistent despite considerable modeling uncertainty. For binaries containing objects between $2.5M_\odot$ and $5M_\odot$ and having massive BH-scale primaries ($> 5M_\odot$), the mass distribution and merger rate is informed by a few events (GW190814 in particular), thus subject to considerable uncertainty in the inferred component mass distributions. Likewise, for binaries containing objects between $2.5M_\odot$ and $5M_\odot$ and having NS-

scale companions, the merger rate is marginally informed by a few events that may not be associated with this region (i.e., GW200115), exacerbating uncertainty in the inferred NS and BH mass distributions. Providing multiple results for these two source classes explores our systematic uncertainty. The models presented in this section are subject to different sources of systematic uncertainty. For example, MS employs a Gaussian to model components in BNSs, whereas PDB uses a single power law with a sharp turn on at low masses to model all objects below the inferred lower edge of the mass gap. These differences result in considerably different BNS rates due to the limited number of detections in the NS mass range. In particular, differences in pairing function shift the rate inference and add statistical uncertainty in the BNS region. MS and BGP both assume independent pairing of component masses, while PDB allows for both independent as well as power law-in-mass-ratio pairing. The power law in mass ratio is dominated by the BBH mergers, which tend to prefer equal masses. However, the observation of unequal mass binaries containing NS (e.g., GW200105 and GW200115) implies that PDB could increase the inferred rate of BNS mergers to account for the lack of equal mass mergers detected in the NS range. Though this change seems to significantly increase the total merger rate, because these hypothesized mergers occur at low mass, they don't imply PDB would drastically overpredict BNS observations, because of the relatively low sensitivity at the lowest allowed NS masses.

C. Characterizing suppressed merger rates between NS and BH Masses

Figures 3 and 4 show a reduction in the rate above NS masses. It was shown using GWTC-2 that the merger rate between $3M_\odot$ and $7M_\odot$ is suppressed relative to an unbroken power-law extending from higher masses




FIG. 3. Rate density versus component masses for different models inferred from events with $\text{FAR} < 0.25 \text{ yr}^{-1}$, illustrating consistency on large, coarse-grained scales, but some disagreement and systematics in areas with few observed events. *Top left panel:* Rate density computed with the FM model assuming no redshift evolution, for binary black holes only. Modulations along lines of constant chirp mass are apparent. *Top right panel:* Rate density inferred with the BGP model using all compact objects. This model can reproduce observations with localized regions of relatively enhanced rate density. In the binary black hole region, some regions of enhanced density are commensurate with the FM result. *Bottom left panel:* Rate density inferred with MS. For mergers involving NS, this model reproduces observations, consistent with smoothing the BGP result. For mergers involving typical BH, this model strongly favors equal-mass mergers. *Bottom right panel:* Rate density inferred with PDB. This model is also consistent with smoothing the FM result, producing features similar to MS, albeit with less structure in the mass ratio distribution for BBH, and by construction lacking a peak near $30M_{\odot}$.

[11]. With additional observations, as well as models and sensitivity estimates that span the full mass range of CBCs, we can now produce a comprehensive perspective on merger rates versus mass throughout the low-mass interval $1\text{--}10M_{\odot}$. In so doing, we find a dropoff in merger rates above NS-scale masses. As a result, in the detection-weighted population, objects with NS-scale mass components are well-separated from objects with BH-scale masses. However, we are unable to confidently infer an absence or presence of a subsequent rise in merger rates from lower mass gap masses. The purported lower mass gap [28–31] between the NS and black hole populations would produce such a rise, such that

the mass gap produces an extended local minimum in the merger rate versus mass. We therefore neither find evidence for nor rule out the existence of a two-sided lower mass gap.

Figure 5 shows the differential rate as a function of component mass inferred from all three models. The PDB model infers the location of this drop-off to occur at $M_{\text{low}}^{\text{gap}} = 2.1_{-0.6}^{+0.7} M_{\odot}$, as shown by black vertical lines. While the other models do not explicitly infer the location of the drop-off, they do clearly show a reduction in the rate at a similar location. The prominence of this drop-off can be characterized by comparing the rate of mergers with both masses below $2.5M_{\odot}$ (BNS) to that of




FIG. 4. Impact of asymmetric binaries on the primary mass distribution, illustrating how the depth and extent of a mass gap in any one-dimensional distribution depends on the choice of slicing or marginalization over the remaining dimension. Differential merger rate as a function of primary mass for the BGP model when considering only the diagonal $q \simeq 1$ bins in Fig. 3 i.e. $m_1 \frac{dN}{dm_1 dq dv_c dt} \Big|_{q \simeq 1}$ and the population of compact binaries across all mass bins. The rate for approximately equal mass binaries is significantly lower highlighting the contribution of asymmetric mergers like NSBHs to the marginal distribution over primary mass. The plot uses the BGP population model inferred from events passing a FAR threshold of $< 0.25 \text{ yr}^{-1}$. Solid curves represent the median rate densities and shaded areas denote 90% credible regions.

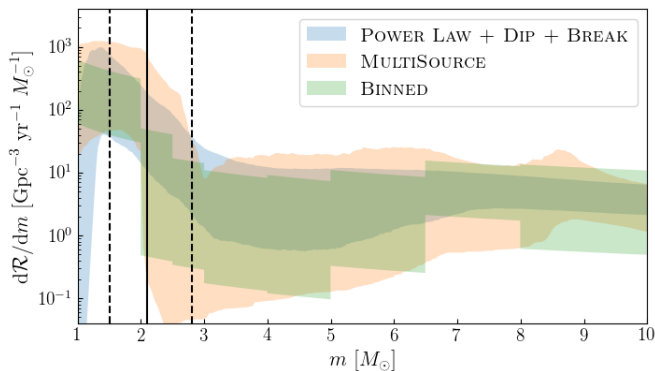


FIG. 5. Differential merger rate as a function of component mass for the PDB, MS, and BGP model. Three independent methods with different modeling assumptions agree on the merger rate versus mass, while illustrating the importance of modeling systematics on the overall rate for objects with NS-scale masses. Shaded areas denote 90% credible regions, while vertical black lines denote the median (*solid*) and 90% credible intervals (*dashed*) of the lower boundary of the mass gap, $M_{\text{low}}^{\text{gap}}$, in the PDB model rate droppoff location.

mergers with at least one component mass between 2.5 and $5M_{\odot}$ (in the mass gap). For this comparison, we find that the differential merger rate of systems with at least one component in the mass gap is one to two orders of magnitude lower than the BNS rate. Thus, even in the absence of any prior knowledge of the difference between NSs and BHs, the gravitational-wave data suggest two

distinct populations of compact objects. This is consistent with results initially found for GWTC-1 by [105].

A subsequent rise in the mass distribution above the putative mass gap is less clearly discernible. The PDB model explicitly parametrizes the mass gap with both low and high-mass transitions $M_{\text{low/high}}^{\text{gap}}$ and a gap depth A (where $A = 0$ corresponds to no gap and $A = 1$ to a lower mass gap containing no events). While the posterior on A peaks around 0.82, i.e. corresponding to a relatively empty mass gap, it has broad support between 0 and 1, indicating an inability to unambiguously differentiate between the presence or absence of a lower mass gap. Additionally, the Bayes factor for a model with no gap ($A = 0$) or a completely empty gap ($A = 1$), relative to the parametrized model, are 0.056 and 1.6, respectively. This lack of clear preference indicates an inability to resolve the absence or existence of a clear gap-like feature in this part of the mass spectrum.

A subsequent rise in the mass distribution at $M_{\text{high}}^{\text{gap}}$ is also less clear to discern. The models infer mass distributions with similar support for both a mildly pronounced gap and a flat transition above $M_{\text{high}}^{\text{gap}}$. Both of these are consistent with the finding in [11] of a deviation from a single power law below primary masses of $\sim 7M_{\odot}$.

We find that if a lower mass gap does exist, it may not be totally empty. While the merger rates show a fall-off above around $2.5M_{\odot}$ in Fig. 5, the rate does not fall to zero. Furthermore, the component masses of 6 events have at least 5% posterior support between $M_{\text{low}}^{\text{gap}}$ and $M_{\text{high}}^{\text{gap}}$ when using a population-informed prior [106]. GW190814 stands out as having considerable support for its secondary being within the mass gap or below the droppoff in the rate at $M_{\text{low}}^{\text{gap}}$: $P(m_{2,\text{GW190814}} \in [M_{\text{low}}^{\text{gap}}, M_{\text{high}}^{\text{gap}}]) = 0.81$. This event has a mass ratio $q = 0.112^{+0.008}_{-0.009}$ [23], hinting at either a potential subpopulation of low- q , low- m_2 BBHs, or a handful of NSBHs with high NS masses. The former possibility is examined in Sec. VI E, and the latter is discussed in Sec. V C. For both the NSBH systems, there is a $\sim 10\%$ probability of the secondary lying in the mass gap and, for GW200115 the primary has a 70% probability of $m_1 < M_{\text{high}}^{\text{gap}}$. Finally, GW190924-021846, which is the BBH event with the lowest total mass, we find roughly equal support for the secondary being either within ($m_2 < M_{\text{high}}^{\text{gap}}$) or above ($m_2 < M_{\text{high}}^{\text{gap}}$) the mass gap.

The inferred depth of the gap does depend heavily on the assumed pairing function: a model in which objects are randomly paired with other objects regardless of mass ratio predicts a more prominent gap than one with a power-law-in-mass-ratio pairing function as in Eq. (11). Similarly, a change of the pairing function will impact the classification of various components as below, in or above the mass gap. Consequently, we do not rely on this methodology for event classification in Sec. V, and instead use EOS-informed limits on the maximum allowed NS mass, and perform leave-one-out analyses with respect to known subpopulations. The lower mass gap-

related results stated here are obtained using a random pairing model.

Though we report on our analysis with $\text{FAR} < 0.25 \text{ yr}^{-1}$, to assess the stability of our results to threshold choices we have repeated our analyses using all events with previously reported parameter inferences below 1 yr^{-1} (i.e., excluding GW190531). Even though such an analysis includes all five candidate NSBH, our key conclusions remain largely unchanged: the derived merger rates versus mass are consistent with the error bars shown in Figs. 4, and 5, and the merger rates reported in Table II are consistent. In particular, we draw similar conclusions about the merger rate between $2M_{\odot}$ and $10M_{\odot}$: suppressed but likely filled, without evidence for or against a true two-sided mass gap.

V. MASS DISTRIBUTION OF NEUTRON STARS IN BINARIES

In this section, we characterize the astrophysical population of NSs using data from the gravitational wave events that are likely to contain at least one NS. Because of the paucity of low-mass compact binary mergers observed to date, and the difficulty in ascertaining the presence of a NS in these systems, modeling the NS population observed in gravitational waves has been challenging. In our previous population analysis through GWTC-2 [11], the rate density of BNS and NSBH mergers was estimated, but the shape of the mass distribution of the NSs in these compact binaries was not inferred. The BNS events GW170817 and GW190425 were included in a joint study of the Galactic and gravitational wave populations of BNSs in [128], which linked the two observed populations via a bimodal birth mass distribution. The confident BNS and NSBH detections made to date were analyzed in a study of the gravitational wave population in [99], which found the observed NS masses to be consistent with a uniform distribution.

We begin by classifying the observed low-mass compact binaries as BNSs, NSBHs or BBHs. The classifications are based on a comparison of their component masses with an EOS-informed estimate of the maximum NS mass, and are corroborated against the location of the lower mass gap between NSs and BHs as inferred in the previous section. Then, adopting these source classifications as definite and considering the BNS and NSBH detections below a FAR threshold of 0.25 yr^{-1} , we infer the shape of the NS mass distribution in compact binaries. In contrast to Sec. IV, we do not attempt to determine the overall rate of such mergers, nor do we attempt to infer the mass distribution of BHs in coalescing NSBH systems. Our analysis makes a comparison with the observed Galactic population of NSs, and we additionally investigate the impact on the population of the event GW190814, a lower mass-gap merger whose secondary may possibly be a NS, but is more likely a low-mass BH.

A. Events containing NSs

The gravitational-wave signal of a compact binary merger involving a NS differs from that of a BBH due to matter effects in the waveform, most notably the phasing of the gravitational waveform during the inspiral due to the tidal deformation of the NS [129]. Since none of the observations in O3b [1] yield an informative measurement of tidal deformability, the gravitational-wave data do not identify which sources contain a NS. Nonetheless, we can establish whether their components are consistent with NSs by comparing their masses to the maximum NS mass, M_{max} , following the method described in [7].

The precise value of M_{max} is unknown because of uncertainty in the NS EOS. Mass measurements for the heaviest known pulsars [130, 131] set a lower bound of $\sim 2 M_{\odot}$ on M_{max} , while basic causality considerations imply that $M_{\text{max}} \lesssim 3 M_{\odot}$ [132, 133]. While individual nuclear theory models for the EOS can produce maximum masses as large as $\sim 3 M_{\odot}$, astrophysical inferences of the EOS generally predict that the maximum mass of a nonrotating NS, the TOV mass $M_{\text{max,TOV}}$, is between $2.2M_{\odot}$ and $2.5 M_{\odot}$ [9, 119–121, 134]. Similarly, studies of GW170817’s merger remnant suggest that $M_{\text{max,TOV}} \lesssim 2.3 M_{\odot}$ [122–127]. Rapid rotation can sustain a maximum mass up to $\sim 20\%$ larger than $M_{\text{max,TOV}}$ [135]. However, the astrophysical processes that form compact binaries may prevent the EOS-supported M_{max} from being realized in the population.

We can therefore identify objects as NS candidates based on their mass using estimates of M_{max} , as long as we assume a clean separation between the NS and BH mass spectra. Of the events with FAR less than 0.25 yr^{-1} , five have at least one component mass with support below $3 M_{\odot}$, making them potentially consistent with a BNS or NSBH merger. These events are listed in Table III, and their component mass posteriors are compared to two estimates of M_{max} in Fig. 6.

For each of these observed low-mass events, we calculate in Table III the probability that at least one of the component masses is less than the maximum NS mass, marginalizing over statistical uncertainties and assuming a uniform component mass prior. We adopt a threshold probability of 50% for classification as a NS. Our fiducial maximum NS mass estimate is taken to be $M_{\text{max,TOV}}$ from the EOS inference of [9], which is based on pulsar timing, gravitational wave and x-ray observations of NSs. That study finds $M_{\text{max,TOV}} = 2.21^{+0.31}_{-0.21} M_{\odot}$, and the corresponding posterior distribution is shown for comparison in Fig. 6. Four of the $\text{FAR} < 0.25 \text{ yr}^{-1}$ events have $P(m < M_{\text{max,TOV}}) > 0.5$ for at least one component, and we deem them either BNSs (if $m_1 < M_{\text{max,TOV}}$) or NSBHs (if only $m_2 < M_{\text{max,TOV}}$). The fifth event, GW190814, has $P(m < M_{\text{max,TOV}}) = 0.06$ and is therefore classified as a BBH. These source classifications do not change if, instead of $M_{\text{max,TOV}}$, we compare against the rotating NS maximum mass, $M_{\text{max}}(\chi)$, as calculated from an empirical relation involving the TOV mass and

the component spin χ [136]. This allows for the possibility that one or more of the low-mass components is rapidly rotating.

We draw similar conclusions about each event if we interpret the sharp decrease in merger rate near $2.5M_\odot$ seen in the PDB analysis as the separation between NS and BH mass ranges. (This interpretation does not imply that $M_{\max,\text{TOV}}$ and $M_{\text{low}}^{\text{gap}}$ need to agree: $M_{\text{low}}^{\text{gap}}$ could be below $M_{\max,\text{TOV}}$ if the heaviest NSs the EOS can support are not realized in nature, or $M_{\text{low}}^{\text{gap}}$ could be above $M_{\max,\text{TOV}}$ if the lower mass gap occurs within the BH mass spectrum.) Following [106], we compare the component mass measurements against the inferred $M_{\text{low}}^{\text{gap}}$ parameter from the PDB model, as shown in Fig. 6, and list the probabilities $P(m < M_{\text{low}}^{\text{gap}})$ in Table III. The same four events are consistent with BNSs or NSBHs.

Fig. 6 also plots the component mass posteriors for two FAR $< 1 \text{ yr}^{-1}$ events from Table I that may contain NSs, if astrophysical in origin. In particular, GW190426 and GW190917 have masses consistent with NSBH systems [3, 4]. This classification is confirmed by the $P(m < M_{\max,\text{TOV}})$ and $P(m < M_{\text{low}}^{\text{gap}})$ probabilities calculated for them in Table III.

B. Mass distribution

Using the FAR $< 0.25 \text{ yr}^{-1}$ events classified as BNSs or NSBHs in Table III, we infer the mass distribution of NSs in merging compact binaries. We adopt the POWER and PEAK parametric mass models described in Sec. III and implement a selection function based on a semi-analytic approximation of the integrated network sensitivity VT , fixing the redshift evolution of the population and ignoring spins when estimating the detection fraction. The population hyper-parameters are sampled from uniform prior distributions, subject to the condition $m_{\min} \leq \mu \leq m_{\max}$ in the PEAK model, except that we assume that the maximum mass in the NS population, m_{\max} , does not exceed $M_{\max,\text{TOV}}$. This is consistent with our use of the nonrotating maximum NS mass to classify the events, and amounts to an assumption that the NSs observed via inspiral gravitational waves are not rotationally supported. In practice, this means imposing a prior proportional to the cumulative distribution function of $M_{\max,\text{TOV}}$, as shown in the inset of Fig. 7 and detailed in Appendix B 1.

The inferred mass distributions for these two models are plotted in Fig. 7. The posterior population distribution for the POWER model has $\alpha = -2.0^{+5.1}_{-7.0}$, consistent with a uniform mass distribution, although the median distribution is a decreasing function of mass. The power-law hyper-parameter is most strongly constrained relative to the flat $\alpha \in [-12, 4]$ prior on the low end. The two bumps in the 90% credible interval visible in Fig. 7 correspond respectively to the minimum and maximum mass cutoffs of the population model realizations with $\alpha < 0$ and $\alpha > 0$. The median inferred PEAK

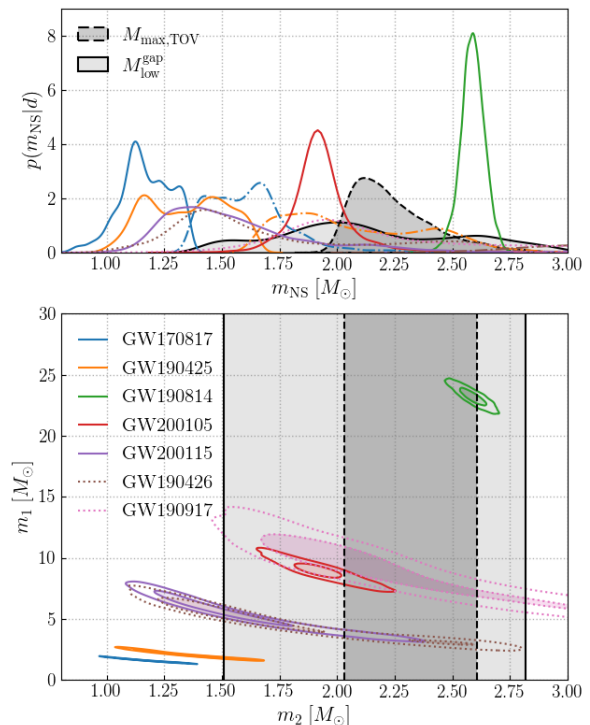


FIG. 6. Masses for events with at least one candidate neutron star. *Upper panel*: one-dimensional posterior distributions for the masses of the candidate NSs, as compared to estimates of the maximum NS mass based on the dense-matter EOS [9] ($M_{\max,\text{TOV}}$) and on the inferred location of the lower mass gap in Sec. IV’s PDB analysis ($M_{\text{low}}^{\text{gap}}$). Primary components are shown dash-dotted. GW190814’s secondary component lies above both estimates of the maximum NS mass. *Lower panel*: two-dimensional 50% (shaded) and 90% (unshaded) credible regions for the binary masses of each candidate NS merger. The marginal events GW190426 and GW190917 are shown dotted. The 90% credible intervals of the maximum NS mass posterior inferred from the EOS and from the lower mass gap location are also plotted. GW190814 occupies a distinct region of the m_1 - m_2 plane compared to the events deemed BNSs or NSBHs.

distribution is relatively flat, and the peak width and location are almost entirely unconstrained relative to the prior: $\sigma = 1.1^{+0.8}_{-0.8} M_\odot$ and $\mu = 1.5^{+0.4}_{-0.4} M_\odot$ for a uniform $\sigma \in [0.01, 2.00] M_\odot$ and $\mu \in [1, 3] M_\odot$ prior subject to $m_{\min} \leq \mu \leq m_{\max}$. Thus, the gravitational wave observations to date do not support a NS mass distribution with a pronounced single peak. This contrasts with the Galactic BNS subpopulation, whose mass distribution is sharply peaked around $1.35 M_\odot$ [39, 41, 137], as shown for comparison in Fig. 7. The mass distribution of NSs observed in gravitational waves is broader and has greater support for high-mass NSs. This latter point

Name	FAR_{\min} (yr^{-1})	$P(m < M_{\max, \text{TOV}})$	$P(m < M_{\text{low}}^{\text{gap}})$	Classification
GW170817	$< 1 \times 10^{-5}$	0.99	0.98	BNS
GW190425	3.38×10^{-02}	0.68	0.73	BNS
GW190814	$< 1 \times 10^{-5}$	0.06	0.19	BBH
GW200105	2.04×10^{-01}	0.94	0.74	NSBH
GW200115	$< 1 \times 10^{-5}$	0.95	0.97	NSBH
GW190426	9.12×10^{-01}	0.82	–	NSBH
GW190917	6.56×10^{-01}	0.56	–	NSBH

TABLE III. Classifications for low-mass events from Table I. The probability that a component is compatible with a NS is measured by the fraction of its mass posterior lying below an estimate [9] of the maximum nonrotating NS mass, $M_{\max, \text{TOV}}$, marginalized over statistical uncertainties. We adopt a 50% threshold for classification as a NS, assuming a clean separation between NS and BH mass spectra. Probabilities are reported relative to a uniform prior on the component mass. They refer to the secondary component of all events except GW170817 and GW190425, in which case the secondary is securely below the maximum NS mass and the probability for the primary is given. The probabilities are similar and the classifications are unchanged when the component masses are compared to $M_{\text{low}}^{\text{gap}}$, the location of the lower mass gap between NSs and BHs inferred from Sec. IV’s PDB analysis of the $\text{FAR} < 0.25 \text{ yr}^{-1}$ events.

is also true compared to the Galactic NS population as a whole, whose mass distribution has a double-peaked shape [138–140].

The minimum NS mass in the gravitational wave population is inferred to be $1.2_{-0.2}^{+0.1} M_{\odot}$ in both the POWER and PEAK models. The lower bound on m_{\min} is a prior boundary motivated by the sensitivity model, as the gravitational-wave searches target sources above $1 M_{\odot}$. The maximum mass in the population is found to be $2.0_{-0.2}^{+0.3} M_{\odot}$ in both models, relative to the EOS-informed m_{\max} prior. This value is consistent with the maximum mass inferred from the Galactic NS population, $2.22_{-0.23}^{+0.85} M_{\odot}$ [139], as can be seen in the inset of Fig. 7. The maximum mass is the best-constrained hyper-parameter in the population models. Its upper bound is more tightly constrained than the Galactic m_{\max} in Fig. 7 as a result of the imposed $m_{\max} \leq M_{\max, \text{TOV}}$ prior, which begins tapering above $2 M_{\odot}$, and the strong selection bias of gravitational-wave observations towards heavier masses, which renders the non-observation of heavier NSs informative. Nonetheless, the statistical uncertainty in m_{\max} remains large, and it is expected that approximately 50 BNS detections will be needed before the maximum mass in the NS population can be measured to within $0.1 M_{\odot}$ [141].

The m_{\max} value inferred from gravitational waves is also as large as $M_{\max, \text{TOV}}$ within statistical uncertainties. This would not be the case if, for instance, the astrophysical processes that form coalescing compact binaries prevented $2 M_{\odot}$ NSs from pairing with other compact objects. Such a scenario is compatible with the EOS-informed m_{\max} prior that we impose. However, we find there is no evidence that the NS mass spectrum observed with gravitational waves is limited by the astrophysical formation channel: NSs as heavy as can be supported by the EOS can end up in merging compact binaries.

Moreover, we infer a consistent maximum mass if we adopt a uniform m_{\max} prior instead of the EOS-informed

one. This relaxes the assumption that the observed NS masses must be below the nonrotating maximum mass, and accounts for the possibility that rapid rotation may cause a NS’s mass to exceed M_{TOV} . Specifically, we find $m_{\max} = 2.2_{-0.3}^{+0.7} M_{\odot}$ in both POWER and PEAK models. The upper error bar on m_{\max} extends to much higher values in this case because it is no longer subject to the tapering EOS-informed prior, which has little support above $2.5 M_{\odot}$. We also obtain consistent results if we expand the event list to include the two marginal NSBH detections listed in Table III, as described in Appendix C 1.

C. Outlier events

The mass-based event classification carried out above deemed GW190814 to be a BBH merger on the basis of the maximum NS mass the EOS can support. We now further demonstrate that it is an outlier from the population of BNSs and NSBHs observed with gravitational waves.

If we dispense with its $M_{\max, \text{TOV}}$ -based classification, and include GW190814 as a NSBH in the population analysis, the inferred maximum mass is shifted up to $2.8_{-0.2}^{+0.2} M_{\odot}$ in the POWER model and $2.7_{-0.2}^{+0.2} M_{\odot}$ in the PEAK model (cf. $m_{\max} = 2.2_{-0.3}^{+0.7} M_{\odot}$ without GW190814 in both models). These values are obtained relative to a uniform m_{\max} prior, since we are no longer consistently enforcing $m \leq M_{\max, \text{TOV}}$; all results in this subsection refer to this prior. The m_{\max} posterior has support up to $3 M_{\odot}$, where the prior truncates and the models’ fixed BH mass distribution begins. The inferred NS mass distributions with GW190814 are similar, but flatter and broader, than those depicted in Fig. 7.

To test whether GW190814 hails from the same population as GW170817, GW190425, GW200105 and GW200115, we examine the PEAK model’s posterior predictive distribution for secondary masses with and




FIG. 7. Inferred neutron star mass distribution. The median mass distribution (solid) and 90% credible interval (shading) inferred for the POWER (respectively, PEAK) population model is shown in blue (orange), as compared to the mass distribution of NSs in Galactic BNSs [40] (dot-dashed black) and the mass distribution of all Galactic NSs [139] (solid black). The inferred gravitational-wave population has a greater prevalence of high-mass NSs. The inset shows the posterior distribution for the maximum mass in the NS population for both models, as compared to the Galactic m_{\max} . The EOS-informed m_{\max} prior, which is proportional to the cumulative distribution function of $M_{\max, \text{TOV}}$, is also shown in the inset (dashed). It enforces $m \leq M_{\max, \text{TOV}}$ using the maximum TOV mass estimate from [9]. The maximum mass in the gravitational-wave population is as large as $M_{\max, \text{TOV}}$ within statistical uncertainties.

without GW190814 in the event list. Figure 8 compares GW190814’s measured $m_2 = 2.59^{+0.08}_{-0.09} M_{\odot}$ against the prediction for the largest observed secondary mass, $\max_5(m_2)$, after two BNS observations and three NSBH observations. That is, we draw two pairs of masses from the posterior predictive distribution for BNSs and three secondary masses from the posterior predictive distribution for NSBHs, take the largest of the five secondaries, and build up the plotted distributions by performing this procedure repeatedly. The probability of observing a secondary mass at least as large as the mean of GW190814’s m_2 in the population is only 0.1% according to the PEAK model fit that excludes GW190814. (We characterize GW190814’s m_2 by its mean, since it is measured so precisely.) The equivalent probability relative to the PEAK model fit that includes GW190814 is 3.5%; we expect a rigorous, fully self-consistent calculation of a p-value to lie between these two numbers [142]. Hence, GW190814’s secondary component is an outlier from the secondaries

in BNS and NSBH systems. In the next section, we also establish GW190814 as an outlier from the BBH population observed in gravitational waves, corroborating our previous analysis [11]. These findings reinforce that it represents a distinct subpopulation of merging compact binaries.




FIG. 8. Comparison between GW190814’s secondary component and the largest secondary mass in the observed BNS and NSBH population. The PEAK model is fit to the population including (respectively, excluding) GW190814. The predicted distribution of the largest secondary mass, $\max_5(m_2)$, observed after five detections—two BNSs and three NSBHs—is shown in orange (blue). The shaded region represents the 90% credible interval of the posterior distribution for the mass of GW190814’s secondary component. GW190814’s m_2 is a 0.1%-level outlier from the rest of the observed population of NS secondaries.

VI. MASS DISTRIBUTION OF BLACK HOLES IN BINARIES

We find two key new conclusions about the black hole mass distribution using the GWTC-3 dataset to infer a population: that the mass distribution has a substructure, reflected in clustering of detected events, and that observations are consistent with a continuous, monotonically decreasing mass distribution at masses $> 50 M_{\odot}$, providing inconclusive evidence for an upper mass gap. Adopting previous coarse-grained models, we find conclusions consistent with our analysis of GWTC-2 [25]. For the purposes of this section, given our large BBH sample, we adopt a FAR threshold of 1 yr^{-1} , but we do not include the previously identified outliers GW190917 (a NSBH) and GW190814 (an extreme mass ratio binary)

in the BBH population unless otherwise noted. Additionally, unlike the redshift-independent results described in Sec. IV, the new analyses described in this section all account for a redshift-dependent BBH merger rate according to Eq. (7). Specifically, in this Section we present results for the PP model to broadly characterize the mass spectrum and corroborate results found in GWTC-2, as well as the the cubic spline power law perturbation PS model and the binned Gaussian process BGP, as both can capture smaller-scale features in the mass distribution. All three models are described in detail in Section III C. We report on the same BGP analysis as performed in Sec. IV, with $\text{FAR} < 0.25 \text{ yr}^{-1}$ and without allowing for redshift dependence; by contrast, the PS, PP, and FM models allow for redshift dependence and use $\text{FAR} < 1 \text{ yr}^{-1}$. Table IV summarizes our results for the overall BBH merger rate, as well as merger rates over restricted mass intervals.

A. Broad features of the mass spectrum

The events from GWTC-3 are broadly consistent with the previously identified population [11]. Figure 9 compares some of the expectations from our previous analysis of GWTC-2 BBHs with the comprehensive sample of GWTC-3 BBH events. The panels compare the observed and expected fractions of all events detected below a threshold in primary mass m_1 , effective inspiral spin χ_{eff} , or source redshift. The panels also show the Wilson score interval [143], a frequentist estimate of the uncertainty in the cumulative distribution F , which is approximately $\pm 1.68 \sqrt{F(1-F)/N_{\text{obs}}}$ when F is significantly different from 0 or 1.

All the cumulative distributions in Figure 9 are broadly consistent with our prior expectations based on coarse-grained models used in our previous work. For this reason, we begin by presenting the inferred coarse-grained mass distribution of black hole binaries, making use of the PP model [11] which best fitted the population from GWTC-2.

Figure 10 shows our inference on the astrophysical primary mass (left) and mass ratio (right) distributions, using the fiducial mass model, compared to what was previously found in GWTC-2 (black). We find a power-law slope for the primary mass, $\alpha = 3.4_{-0.49}^{+0.58}$, supplemented by a Gaussian peak at $34_{-3.8}^{+2.3} M_{\odot}$. On the upper end, the mass of the 99th percentile, $m_{99\%}$, is found to be $45_{-6.2}^{+11} M_{\odot}$. The mass ratio distribution is modelled as a power law q^{β_q} with $\beta_q = 1.1_{-1.3}^{+1.8}$.

In contrast to our GWTC-2 population fit, the inferred mass spectrum decays more rapidly; the $m_{99\%}$ is considerably lower than $60_{-13}^{+14} M_{\odot}$, as was found with GWTC-2. These results are expected, given that the new observations in GWTC-3 contain a greater fraction of lower mass systems (see, e.g., Fig. 1). The fraction of BBH mergers with primary masses within the Gaussian component of the fiducial model is found to be

$\lambda = 0.039_{-0.026}^{+0.058}$ ($0.1_{-0.071}^{+0.14}$ in GWTC-2), but still rules out zero. This result further highlights that the fraction of higher mass binaries has decreased in GWTC-3. Both the mean and the standard deviation of the Gaussian component are consistent with previous inferences. Furthermore, the inferred mass ratio distribution is less peaked towards equal mass binaries ($\beta_q = 1.1_{-1.3}^{+1.8}$) compared to GWTC-2 ($\beta_q = 1.3_{-1.5}^{+2.4}$), a result driven by the discovery of binaries with support for substantially unequal masses (see, e.g., Fig. 9).

We previously used several other phenomenological models to interpret sources in GWTC-2. Using this broader suite of models, we draw similar conclusions to those presented above: the mass distribution is inconsistent with the single power law and has a feature at $\sim 35\text{--}40 M_{\odot}$. The peak's location is also well-separated from the largest black holes predicted by the other component: the over-density and maximum mass are still not associated. The odd ratios discriminating between these models are modest, of order one in three to one in ten. Despite for presentation purposes adopting the PP for illustrating consistency with GWTC-2 results, we cannot decisively differentiate between a peak near $35 M_{\odot}$ versus a more generic transition towards a lower merger rate at higher mass; see Appendix D 1 for details.

In Table IV, we provide BBH merger rates for the full population, as well as split based upon the primary mass at $m_1 < 20 M_{\odot}$, $m_1 \in [20, 50] M_{\odot}$ and $m_1 > 50 M_{\odot}$ to capture the broad features of the mass spectrum: the high rate at low masses, a peak around $35\text{--}45 M_{\odot}$ and the falling merger rate at high masses.

B. Mass distribution has substructure

With new discoveries in O3, we are now confident the mass distribution has substructure, with localized peaks in the component mass distribution. For example, we find overdensities in the merger rate ($> 99\%$ credibility) as a function of primary mass, when compared to a power law, at $m_1 = 10_{-0.59}^{+0.29} M_{\odot}$ and $m_1 = 35_{-3.1}^{+1.5} M_{\odot}$. At best, we have modest confidence (less than 10:1 odds) in additional structure. These signs of substructure were identified in O3a [45] and are corroborated by consistent observations in O3b.

We arrive at these conclusions through multiple independent analyses. Each of these model agnostic approaches attempts to reconstruct the mass distributions with minimal constraints imposed. Specifically, we employ a flexible mixture model (introduced in Section III C and labelled FM in tables and figures), a cubic spline power law perturbation (PS), and a binned Gaussian process (BGP). Figure 11 shows the inferred rate $d\mathcal{R}/dm_1$ as a function of primary mass for each of the non-parametric models. There is a clear presence of structure beyond an unbroken single power law found when using these more flexible models, with a global maximum of the merger rate at larger masses at around $10 M_{\odot}$ followed by a fall




FIG. 9. The empirical cumulative density function $\hat{F} = \sum_k P_k(x)/N$ of observed binary parameter distributions (derived from the single-event cumulative distributions $P_k(x)$ for each parameter x) are shown in black for primary mass (left), effective inspiral spin (center), and redshift (right). All binaries used in this study with $\text{FAR} < 1/4\text{yr}$ are included, and each is analyzed using our fiducial noninformative prior. For comparison, the gray bands show the expected observed distributions, based on our previous analysis of GWTC-2 BBH. Solid lines show the medians, while the shading indicates a 90% credible interval on the empirical cumulative estimate and selection-weighted reconstructed population, respectively. GW190814 is excluded from this analysis.




FIG. 10. The astrophysical BBH primary mass (left) and mass ratio (right) distributions for the fiducial PP model, showing the differential merger rate as a function of primary mass or mass ratio. The solid blue curve shows the posterior population distribution (PPD) with the shaded region showing the 90% credible interval. The black solid and dashed lines show the PPD and 90% credible interval from analyzing GWTC-2 as reported in [11]. The vertical gray band in the primary mass plot shows 90% credible intervals on the location of the mean of the Gaussian peak for the fiducial model.

	$m_1 \in [5, 20]M_\odot$ $m_2 \in [5, 20]M_\odot$	$m_1 \in [20, 50]M_\odot$ $m_2 \in [5, 50]M_\odot$	$m_1 \in [50, 100]M_\odot$ $m_2 \in [5, 100]M_\odot$	All BBH
PP	$23.4^{+12.9}_{-8.6}$	$4.5^{+1.8}_{-1.3}$	$0.2^{+0.1}_{-0.1}$	$28.1^{+14.8}_{-10.0}$
BGP	$20.0^{+10.0}_{-8.0}$	$6.4^{+3.0}_{-2.1}$	$0.74^{+1.2}_{-0.46}$	$33.0^{+16.0}_{-10.0}$
FM	$21.1^{+10.7}_{-8.3}$	$4.1^{+2.0}_{-1.4}$	$0.2^{+0.3}_{-0.1}$	$26.0^{+11.5}_{-8.7}$
PS	$27^{+12}_{-9.4}$	$3.6^{+1.5}_{-1.1}$	$0.2^{+0.18}_{-0.1}$	$32^{+14}_{-9.6}$
MERGED	12.8 – 40	0.098 – 6.3	2.5 – 0.5	17.3 – 45

TABLE IV. Merger rates in $\text{Gpc}^{-3} \text{yr}^{-1}$ for BBH binaries, quoted at the 90% credible interval, for the PP model and for three non-parametric models (BINNED GAUSSIAN PROCESS, FLEXIBLE MIXTURES, POWER LAW + SPLINE). Rates are given for three ranges of primary mass, m_1 as well as for the entire BBH population. Despite differences in methods, the results are consistent among the models. BGP assumes a non-evolving merger rate in redshift. The merger rate for PP, FM, and PS is quoted at a redshift value of 0.2, the value where the relative error in merger rate is smallest.

off to lower rates. Modulating this extended decline, the PS, FM and even BGP show indications of additional structure. As the BGP likely cannot resolve small-scale features, we assess these features’ details and significance with the remaining two models.

Figure 12 shows the results of the spline perturbation model, where 1000 posterior draws of the spline function $f(m_1)$ are illustrated, where $\exp f(m_1)$ modulates an underlying power-law distribution. The inferred perturbation $f(m_1)$ strongly disfavors zero at both the $10M_\odot$ and $35M_\odot$ peak, finding $f \leq 0$ at 0.38%, and $< 0.07\%$ credibilities respectively. Additionally for the drop in merger rate at $14M_\odot$, the PS model finds $f \leq 0$ at 96.4% credibility.

C. Inconclusive evidence for upper mass gap

Stellar evolution models predict a lack of black holes with masses from $50_{-10}^{+20}M_\odot$ to $\sim 120M_\odot$ due to pair production instability [144–150]. The high-mass event GW190521 could have a component lying within this mass gap [22, 151]. Other analyses of this event with independent parameter inferences have argued this event could have both components outside this gap [152–154]. We define a gap as a rapid decline in merger rate at some cutoff mass, followed by a rapid rise in the distribution at a significantly higher mass. Repeating similar analyses with the full O3 data set, we find no evidence for such a gap. Following [154], we extend our PP mass model to allow for masses $> 100M_\odot$, and to include a zero-rate mass interval, parameterized with the lower edge and width of the gap. With this extended model, we find minimal posterior support for the gap to start at $< 75M_\odot$ (4.1% credibility). When it does, the gap width is constrained to be $< 35M_\odot$. The majority of the posterior support has the gap start above $75M_\odot$, consistent with the inferred maximum mass cutoff from the PP model without a gap. This allows both component masses of the most massive BBH in the catalog, GW190521, to fall below the cutoff, leaving no observations with masses larger than the start of the gap. We are not able to determine whether or not the mass distribution exhibits a rise again at higher masses. We find a slight preference ($\ln \mathcal{B} = 0.6$) for the PP model without a gap over one with the gap included, thus we report inconclusive evidence for a zero-rate upper mass gap. Inconclusive support for a zero-rate gap challenges classical conclusions for the pair-instability mass gap. The pair-instability mass gap could start higher than theory expects, or the high-mass binaries in our catalog could be formed in a way that avoids pair-instability.

D. Evolution of rate with redshift

The observation of BBH mergers offers us the means of not only measuring the local merger rate per comoving volume but also the *evolution* of this merger rate as we

look back towards larger redshifts z . Given the limited range of redshift to which our searches are sensitive, we parametrize the merger rate per comoving volume as a simple power law, with $\mathcal{R}(z) \propto (1+z)^\kappa$ [96].

In our previous study [11], the redshift distribution was weakly constrained, exhibiting a preference for a rate that increased with redshift but still consistent with a non-evolving merger rate. Here, in addition to new events observed in O3b, we leverage updated pipelines and our improved sensitivity models to update our inference of κ . As discussed further in Appendix C4, these sensitivity model refinements indicate a lower search sensitivity to high-redshift BBH mergers than previously concluded. We now confidently claim to see evolution of the BBH merger rate with redshift in our population with a FAR $< 1\text{ yr}^{-1}$, inferring that $\kappa > 0$ at 99.4% credibility. While the exact distribution of κ does depend on the chosen mass model, we can rule out a redshift-independent merger rate at similar credible levels when adopting any of the parameterized mass distribution models used in [11].

Figure 13 shows the marginal posterior on κ given GWTC-3 in blue, obtained while using the PP and DEFAULT mass and spin models. The dashed distribution, meanwhile, shows the previously published measurement of κ obtained with GWTC-2. In Fig. 13 we also show our corresponding constraints on $\mathcal{R}(z)$ itself as a function of redshift. The dark blue line traces our median estimate on $\mathcal{R}(z)$ at each redshift, while the dark and light shaded regions show central 90% and 50% credible bounds. Our best measurement of the BBH merger rate occurs at $z \approx 0.2$, at which $\mathcal{R}(z = 0.2) = 19\text{--}41\text{ Gpc}^{-3}\text{ yr}^{-1}$. For comparison, the dashed black line in Fig. 13 is proportional to the Madau–Dickinson star formation rate model [155], whose evolution at low redshift corresponds to $\kappa_{\text{SFR}} = 2.7$. While the rate evolution remains consistent with the Madau–Dickinson star formation rate model, it is not expected for these two rates to agree completely due to the time delays from star formation to merger [156–164].

In most plausible formation scenarios (e.g., if BBHs arise from stellar progenitors), we do not expect $\mathcal{R}(z)$ to continue growing with arbitrarily high z . Instead, we anticipate that $\mathcal{R}(z)$ will reach a maximum beyond which it turns over and falls to zero. Even in cases where the peak redshift z_p at which $\mathcal{R}(z)$ is maximized lies beyond the LIGO–Virgo detection horizon, a sufficiently tight upper limit on the stochastic gravitational-wave background due to distant compact binary mergers [165–167] can be leveraged to bound z_p from above, potentially yielding a joint measurement of κ and z_p [168]. As demonstrated in [169], our current instruments are not yet sensitive enough to enable a meaningful joint constraint on κ and z_p , even with the inclusion of new events in GWTC-3.

As heavy BBHs are primarily believed to arise from low-metallicity stellar progenitors [170–172], one might wonder if more massive BBHs are observed at systematically higher redshifts than less massive systems. More-




FIG. 11. The differential merger rate for the primary mass predicted using three non-parametric models compared to the fiducial PP model. Solid curves are the medians and the colored bands are the 90% credible intervals. These models offer increased flexibility compared to phenomenological models in predicting the population. The PS applies a perturbation to the primary mass in a modified version of our fiducial PP model that does not include the Gaussian peak. FM models the chirp mass, mass ratio, and aligned spin distribution as a weighted sum of mixture components. Both of these models incorporate a single parameter redshift evolution of the merger rate [Eq. (7)]. The BGP models the two-dimensional mass distribution as a binned Gaussian Process which is piecewise constant in $\log m_i$, illustrating the same analysis as presented in Sec. IV with $\text{FAR} < 0.25 \text{ yr}^{-1}$. All three models infer a local maximum in the merger rate at around $10 M_\odot$ and $35 M_\odot$.




FIG. 12. The cubic spline function, $f(m_1)$, describing the perturbations to an underlying power law inferred with the PS model. The thin grey lines show 1000 draws from the posterior while the black lines show the knot locations (vertical) and the 90% credible region of the posterior. The dashed blue lines mark the 90% credible bounds of the Gaussian priors (centered on zero) imposed on each knot's height. The shaded region covers any masses less than the 95th percentile of the marginal posterior distribution on m_{\min} . Because the low mass region of the mass distribution is cut off and there are no observations there, the posterior in this region resembles the prior of the cubic spline function.

over, any metallicity dependence in the physics of stars, such as the maximum black hole mass imposed by pair instability supernovae (PISN) [144, 146, 173], could yield redshift-dependent features in the black hole mass distribution [174, 175]. Such a redshift dependence would confound efforts to leverage the PISN mass gap as a probe of cosmology. Previous investigations [176] demonstrated

using GWTC-2 that redshift dependence of the maximum BBH mass would be *required* to fit the observations if the BBH mass distribution has a sharp upper cutoff. However, if the distribution decays smoothly at high masses, for example as a power-law, the data are consistent with no redshift dependence of the cutoff location.

We revisit this question using the latest BBH detections among GWTC-3, finding that these conclusions remain unchanged. Specifically, by modelling the high-mass tail of the distribution with a separate power-law index, we find no evidence that the distribution is redshift dependent, suggesting that the high-mass structure in the BBH mass distribution remains consistent across redshift.

E. Outliers in the BBH Population

While we inferred the population of most BBH and binaries involving NS, some systems (particularly with high mass ratio) lie at the boundary between these categories [23, 142]. So far, we have simply excluded these events from our BBH analysis. To demonstrate this choice is internally self-consistent and well-motivated, we show that these events are outliers from our recovered BBH population. Specifically, we repeat the population analysis using the PP model, highlighting the extent to which the population changes when including these events.

For a population consisting of all potential BBH events in O3, including GW190917 and GW190814, the mass distribution must extend to lower masses. In Fig. 14 we plot the recovered distribution for the minimum BH




FIG. 13. Constraints on the evolution of the BBH merger rate with redshift. *Top*: Posterior on the power-law index κ governing the BBH rate evolution, which is presumed to take the form $\mathcal{R}(z) \propto (1+z)^\kappa$. The blue histogram shows our latest constraints using GWTC-3 ($\kappa = 2.7^{+1.8}_{-1.9}$), while the dashed distribution shows our previous constraints under GWTC-2. *Bottom*: Central 50% (dark blue) and 90% (light blue) credible bounds on the BBH merger rate $\mathcal{R}(z)$. The dashed line, for reference, is proportional to the rate of cosmic star formation [155]; we infer that $\mathcal{R}(z)$ remains consistent with evolution tracing star formation.

mass, m_{\min} , that characterizes the primary mass scale above which a power-law distribution exists. The minimum mass is $m_{\min} = 2.2^{+0.28}_{-0.22} M_\odot$, with an extremely sharp turn-on of $\delta_m = 0.38^{+1.2}_{-0.35} M_\odot$. By contrast, if we remove the two low-mass events, we find a minimum BH mass of $m_{\min} = 2.5^{+0.67}_{-0.44} M_\odot$, which is consistent with a mass gap, and a broader turn-on of $\delta_m = 7.8^{+1.9}_{-4.0} M_\odot$. It is the secondary masses, m_2 of these events which is in tension with the remainder of the population. This is clearly demonstrated in Fig. 14 where the secondary masses are shown by the shaded regions. A single minimum BH mass is imposed upon all BH, and therefore it is the secondary masses of low-mass or asymmetric binaries have the strongest impact on our inference of m_{\min} .

These analyses imply two key results about the com-




FIG. 14. The posterior distribution on the minimum mass truncation hyper-parameter, m_{\min} , inferred with the PP model. The posteriors are shown both including and excluding the two BH mergers containing low mass secondaries, GW190814 and GW190917. The cutoff at $m_{\min} = 2 M_\odot$ corresponds to the lower bound of the prior distribution. The inclusion of either of these two events significantly impacts the distribution. The shaded regions indicate the 90% credible interval on the m_2 posterior distribution for the two outlier events, GW190814 (purple) and GW190917 (grey).

pact binary population. On the one hand, the binary black hole population excluding highly asymmetric systems GW190814 is well-defined, and the analyses carried out in this section well-suited to characterizing the bulk of the BBH population. On the other hand, the existence of GW190814 implies the existence of a sub-population of highly asymmetric binaries, disconnected from the BBH population but potentially connected to the recently-identified population of NSBH.

VII. SPIN DISTRIBUTION OF BLACK HOLES IN BINARIES

Compared to our previous work [11], we find two key new conclusions for black hole spins: that the spin distribution broadens above $30 M_\odot$, and that the mass ratio and spin are correlated. Adopting previous coarse-grained models, we find consistent conclusions as our analysis of GWTC-2; notably, we still conclude that a fraction of events probably have negative χ_{eff} .

The component spins of binary black holes may offer vital clues as to the evolutionary pathways that produce merging BBHs [177–184]. The magnitudes of BBH spins are expected to be influenced by the nature of angular momentum transport in stellar progenitors [172, 185, 186], processes like tides [183, 187, 188] and mass transfer that operate in binaries, and the environment in which the binary itself is formed. Their directions, meanwhile, may tell us about the physical processes by which binaries are most often constructed; we expect BBHs born from isolated stellar evolution to possess spins preferentially aligned with their orbital angular momenta, while binaries that are dynamically assembled

in dense environments are predicted to exhibit isotropically oriented spins [177, 179].

Figure 15 illustrates our constraints on the component spin magnitudes (left) and spin tilts (right) of BBHs under the DEFAULT spin model. Using GWTC-3, we make similar conclusions regarding the spin magnitude distribution as made previously with GWTC-2. In particular, spin magnitudes appear concentrated below $\chi_i \lesssim 0.4$, with a possible tail extending towards large or maximal values. Our understanding of the spin tilt distribution, in contrast, has evolved with the addition of new BBHs in GWTC-3. As in GWTC-2, we again exclude the case of perfect spin-orbit alignment (corresponding to $\zeta = 1$ and $\sigma_t = 0$). With GWTC-3, however, we more strongly favor a broad or isotropic distribution of spin tilts. This shift is seen in the right-hand side of Fig. 15: whereas the $\cos\theta$ distribution inferred from GWTC-2 was consistent with tilts concentrated preferentially around $\cos\theta = 1$, evidence for this concentration is now diminished, with O3b results preferring a flatter distribution across $\cos\theta$.

Figure 16 illustrates our updated constraints on the χ_{eff} and χ_p distributions under the GAUSSIAN spin model. As above, our previous results obtained with GWTC-2 are shown in blue, while black curves show our updated measurements with O3b. Measurement of the χ_{eff} distribution with GWTC-2 suggested an effective inspiral spin distribution of non-vanishing width centered at $\chi_{\text{eff}} \approx 0.05$, while the χ_p distribution appeared incompatible with a narrow distribution at $\chi_p = 0$, bolstering the conclusion above that the BBH population exhibits a range of non-vanishing spin-tilt misalignment angles. These conclusions are further strengthened when updating our analysis with GWTC-3. We again infer a χ_{eff} distribution compatible with small but non-vanishing spins, with a mean centered at $0.06_{-0.05}^{+0.04}$. Our updated constraints on the effective precessing spin distribution reaffirm the need for non-vanishing χ_p among the BBH population. The χ_p measurements made previously with GWTC-2 were consistent with both a broad underlying distribution or a very narrow distribution centered around $\chi_p \approx 0.3$; this latter possibility is the source of the apparent jaggedness seen in the GWTC-2 result. We draw similar conclusions with GWTC-3, finding that χ_p measurements can be explained either by a broad distribution centered at $\chi_p = 0$, or a narrow distribution centered at $\chi_p \approx 0.2$. If we include GW190814 in our sample (which is otherwise excluded by default from our BBH analyses) support for this second mode is diminished, leaving a zero-centered χ_p distribution with standard deviation $0.16_{-0.08}^{+0.15}$.

In addition to the distributions of effective inspiral spins and component spins χ_1 and χ_2 associated with the more and less massive components of BBHs, respectively, we also explore the distributions of the *more and less rapidly spinning* components among the BBH population [189]. For a given binary, we define $\chi_A = \max_{|\chi|}(\chi_1, \chi_2)$ and $\chi_B = \min_{|\chi|}(\chi_1, \chi_2)$ as the component spins with the larger and smaller magnitudes, re-

spectively. As discussed in Sec. IX, some models for stellar evolution and explosion predict that isolated black holes are born effectively non-rotating and that binary black hole systems primarily acquire spin through tidal spin-up of the secondary component by the first-born (non-spinning) black hole. If this is the case, then we expect to observe a non-vanishing distribution of χ_A but a distribution of χ_B concentrated at or near zero. Figure 17 shows the resulting distributions of these spin-sorted magnitudes χ_A (blue) and χ_B (green), as implied by the DEFAULT model constraints on component spin magnitudes and tilt angles. Light and dark shaded regions show 50% and 90% credible bounds on each parameter, while the dark lines trace the expectation value of $p(\chi)$ as a function of spin-sorted χ . The χ_A distribution, by definition, is concentrated at larger values than the peak seen in Fig. 15 (at $\chi \approx 0.2$). Across the BBH population, these more rapidly spinning components exhibit a distribution that likely peaks near $\chi_A \approx 0.4$, with 1st and 99th percentiles at $0.08_{-0.04}^{+0.05}$ and $0.78_{-0.09}^{+0.09}$, respectively. Less rapidly spinning components, meanwhile, are centered at or below $\chi_B \lesssim 0.2$, with 99% of values occurring below $0.53_{-0.07}^{+0.09}$.

One significant question explored in our previous study [11] was the degree to which BBHs exhibit *extreme* spin-orbit misalignment, with tilt angles exceeding $\theta \geq 90^\circ$ and thus negative effective inspiral spins. Such steeply tilted spins are unlikely for BBH formation from isolated stellar progenitors [190], and hence would serve as a strong indicator of dynamical interaction during BBH evolution. Our GWTC-2 study [11] interpreted the results of the DEFAULT and GAUSSIAN spin analyses as indicating the presence of extremely misaligned spins. As seen in Fig. 15, the component spin-tilt distribution is non-vanishing below $\cos\theta = 0$. Similarly, in Fig. 16 the χ_{eff} distribution has significant support at $\chi_{\text{eff}} < 0$. To check whether this requirement for negative χ_{eff} was a true feature of the data or an extrapolation of the Gaussian population model (which assumes the existence of extended tails), we extended the Gaussian model to truncate the effective inspiral spin on the range $\chi_{\text{eff},\text{min}} \leq \chi_{\text{eff}} \leq 1$ (rather than $-1 \leq \chi_{\text{eff}} \leq 1$) and hierarchically measured the lower truncation bound $\chi_{\text{eff},\text{min}}$. We found $\chi_{\text{eff},\text{min}} < 0$ at 99.1% credibility, concluding that the data required the presence of negative effective inspiral spins. We obtain consistent results if we perform an identical check with GWTC-3; Fig. 18 illustrates our updated posterior on $\chi_{\text{eff},\text{min}}$, now inferred to be negative at 99.8% credibility.

This interpretation was challenged in [191] and [192], which argued that no evidence for extreme spin misalignment exists if BBH spin models are expanded to allow the existence of a secondary subpopulation with vanishingly small spins. Other avenues of investigation are also in tension with the identification of extreme spin-orbit misalignment. When the χ_{eff} distribution is allowed to *correlate* with other BBH parameters, like the binary mass ratio (see Sec. VII B), evidence for negative χ_{eff} values




FIG. 15. The distributions of component spin magnitudes χ (left) and spin-orbit misalignment angles θ (right) among binary black hole mergers, inferred using the DEFAULT component spin model described further in Sect. B 2 a; e.g., both spin magnitudes are drawn from the same distribution. In each figure, solid black lines denote the median and central 90% credible bounds inferred on $p(\chi)$ and $p(\cos\theta)$ using GWTC-3. The light grey traces show individual draws from our posterior distribution on the DEFAULT model parameters, while the blue traces show our previously published results obtained using GWTC-2. As with GWTC-2, in GWTC-3 we conclude that the spin magnitude distribution peaks near $\chi_i \approx 0.2$, with a tail extending towards larger values. Meanwhile, we now more strongly favor isotropy, obtaining a broad $\cos\theta_i$ distribution that may peak at alignment ($\cos\theta_i = 1$) but that is otherwise largely uniform across all $\cos\theta$.




FIG. 16. Left panel: Inferred distribution of χ_{eff} for our latest full analysis in black. For comparison, the blue distribution and interval shows our inferences derived from GWTC2. Right panel: Corresponding result for χ_p . While both panels in this figure are derived using the Gaussian spin model, we find similar conclusions with the other spin models used to analyze GWTC-2.

diminishes [193]. Motivated by the concerns raised in [191] and [192], we repeat our inference of $\chi_{\text{eff},\text{min}}$ but under an expanded model that allows for a narrow sub-population of BBH events with extremely small effective inspiral spins:

$$p(\chi_{\text{eff}}|\mu_{\text{eff}}, \sigma_{\text{eff}}, \chi_{\text{eff},\text{min}}) = \zeta_{\text{bulk}}\mathcal{N}_{[\chi_{\text{eff},\text{min}}, 1]}(\chi_{\text{eff}}|\mu_{\text{eff}}, \sigma_{\text{eff}}) + (1 - \zeta_{\text{bulk}})\mathcal{N}_{[-1, 1]}(\chi_{\text{eff}}|0, 0.01). \quad (18)$$

Here, ζ_{bulk} is the fraction of BBHs in the wide *bulk* population, truncated above $\chi_{\text{eff},\text{min}}$, while $(1 - \zeta_{\text{bulk}})$ is the fraction of events residing in the *vanishing spin* sub-population, which formally extends from -1 to 1 . When repeating our inference of $\chi_{\text{eff},\text{min}}$ under this expanded

model, our data still prefer a negative $\chi_{\text{eff},\text{min}}$ but with lower significance. As seen in Fig. 18, we now infer that $\chi_{\text{eff},\text{min}} < 0$ at 88.4% credibility. This expanded model allows us to additionally investigate evidence for the existence of a sub-population of BBHs with vanishingly small spins. GWTC-3 prefers but does not require such a sub-population to exist. We measure $\zeta_{\text{bulk}} = 0.49^{+0.32}_{-0.22}$, with $\zeta_{\text{bulk}} > 0.18$ at 99% credibility, but also find that our posterior remains consistent with $\zeta_{\text{bulk}} = 1$.




FIG. 17. Distribution of magnitudes of the most (χ_A) and least (χ_B) rapid component spin among BBHs in GWTC-3. Traces show individual draws from our posterior on the spin population under the DEFAULT model, while dark curves bound 90% credible bounds on $p(\chi_A)$ and $p(\chi_B)$.

A. Spin distribution consistent as mass increases

Our previous analysis adopted the same spin distribution at all masses. The spins of low-mass binaries dominate the reconstructed spin distribution. However, the binaries with the most extreme values of spins have heavier masses: observations GW170729, GW190517, GW190519, GW190620, GW190706, GW190805, and GW191109, constitute 70% of the binaries with moderate to high spins. This preponderance of massive binaries with large spin suggests a one-size-fits-all approach might not fully capture how well we can predict black hole spins *given their mass*; conversely, this preponderance can also reflect the increased impact on our search sensitivity at the highest masses. Too, astrophysical formation scenarios often predict correlations between mass and spin, both from isolated and dynamical formation [172, 194]. Using the FM model for the aligned spin components we reconstruct the trend of $|s_z|$ versus mass. Figure 19 shows the aligned spin magnitude distribution versus binary chirp mass. At low masses, the aligned spin is consistent with (and well constrained to be close to) zero, (i.e., maximum aligned spin magnitude averaged over chirp masses $30M_\odot$ or less is 0.37 at 90% credibility). At heavier masses, the aligned spin is still consistent with zero, albeit with larger dispersion (i.e., maximum aligned spin magnitude averaged over chirp masses $30M_\odot$ or more is 0.51 at 90% credibility). This trend is qualitatively consistent with the relative proportion of events versus chirp mass: very few observations have high chirp masses, providing relatively little leverage to constrain spins. At high chirp masses, the spin distribution is poorly constrained by only a handful of measurements, closer to our broad prior assumptions, in contrast to the better-constrained distribution at low mass. We have no




FIG. 18. Cumulative probabilities on the minimum truncation bound on the χ_{eff} distribution as inferred using GWTC-2 and GWTC-3. When modeling the effective inspiral spin distribution as a Gaussian truncated on $\chi_{\text{eff},\text{min}} \leq \chi_{\text{eff}} \leq 1$, we inferred using GWTC-2 that $\chi_{\text{min,eff}} < 0$ at 99.1% credibility, and hence that the data support the existence of BBH mergers with negative effective inspiral spins. Using GWTC-3, this same analysis more strongly infers that $\chi_{\text{min,eff}} < 0$, now at 99.8% credibility. As discussed further below, evidence for negative effective inspiral spins is diminished under an expanded model that allows for a subset of BBHs to possess vanishing effective inspiral spins. When instead modeling the χ_{eff} distribution as a mixture between a broad Gaussian and a narrow Gaussian sub-population centered at $\chi_{\text{eff}} = 0$ (e.g., the second consistent with zero spin), we infer $\chi_{\text{min,eff}} < 0$ at 88.4% credibility.

evidence to support or refute a trend of aligned spin with chirp mass.

Figure 19 suggests aligned spin magnitude remains constrained to be close to zero independently of the most well-identified peaks in the mass distribution, contrary to what would be expected from hierarchical formation scenarios for these peaks [108, 194–198].

B. High spin correlates with asymmetric binaries

BBHs may exhibit an anti-correlation between their mass ratios and spins, such that binaries with $q \sim 1$ favor effective inspiral spin parameters near zero, while binaries with more unequal mass ratios exhibit preferentially positive χ_{eff} values [193]. To evaluate the degree to which q and χ_{eff} are (or are not) correlated, following prior work [193] we adopt a Gaussian model for the χ_{eff} distribution with a mean and standard deviation that are allowed to evolve with q :

$$p(\chi_{\text{eff}}|q) \propto \exp \left[-\frac{(\chi_{\text{eff}} - \mu(q))^2}{2\sigma^2(q)} \right], \quad (19)$$

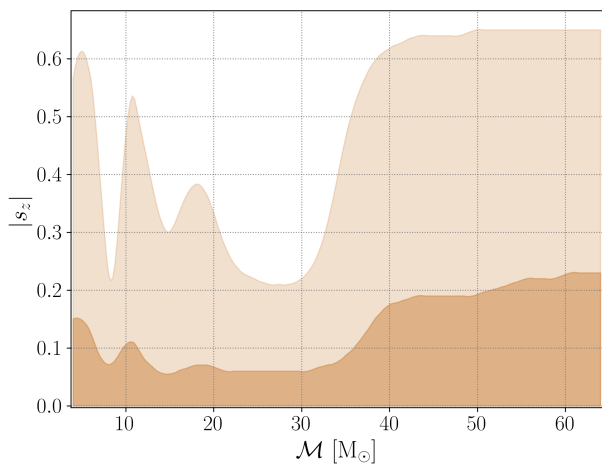


FIG. 19. The dependence of aligned spin magnitude on the chirp mass. The light/dark shaded regions are the aligned spin magnitude at a credibility 90%/50%. The distribution is consistent with small values for lower chirp mass binaries, however, there is an increase in the magnitude for chirp masses of $30 M_{\odot}$ and higher.




FIG. 20. Posterior constraints on the mean (top) and standard deviation (bottom) of the χ_{eff} distribution as a function of mass ratio q . At 98.3% credibility, we find that the mean of the χ_{eff} shifts towards larger values for more unequal mass systems. The grey region in the lower panel shows the area artificially excluded by our prior on the parameters σ_0 and β ; see Eq. (20).

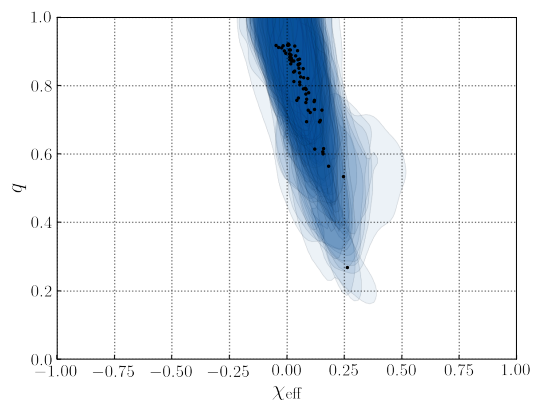


FIG. 21. Posteriors on the mass ratios and effective inspiral spins of BBHs in GWTC-3, reweighted to a population-informed prior allowing for a correlation between q and χ_{eff} . We infer that the mean of the BBH χ_{eff} distribution shifts towards larger values with decreasing mass ratios. Accordingly, reweighted events shift considerably, such that events with $q \sim 1$ contract about $\chi_{\text{eff}} \approx 0$ while events with $q < 1$ shift towards larger effective inspiral spins.

with

$$\mu(q) = \mu_0 + \alpha(q - 1) \quad (20a)$$

$$\log_{10} \sigma(q) = \log_{10} \sigma_0 + \beta(q - 1). \quad (20b)$$

The new hyperparameters α and β measure the extent to which the *location* or *width* of the χ_{eff} distribution changes as a function of mass-ratio.

We repeat hierarchical inference of the BBH population, adopting the fiducial model for the primary mass and redshift distribution. At 98.3% credibility, we constrain $\alpha < 0$, indicating that more unequal-mass binaries preferentially possess larger, more positive χ_{eff} . Figure 20 illustrates our constraints on the mean and standard deviation of the χ_{eff} distribution as a function of mass ratio. Each light trace represents a single sample from our hyperposterior, and the solid black lines denote the median values and central 90% bounds on $\mu(q)$ and $\sigma(q)$ at a given value of q . If we adopt these hierarchical results as a new, population-informed prior, Fig. 21 shows the resulting reweighted posteriors for the BBHs among GWTC-3. Each filled contour bounds the central 90% region for a given event in the q - χ_{eff} plane, while black points mark events' one-dimensional median q and χ_{eff} measurements.

VIII. COMPARISON WITH OTHER GW CATALOGS

In this paper, we have presented population inferences based upon events identified by the LIGO Scientific, Virgo and KAGRA Collaborations in data taken by the Advanced LIGO and Advanced Virgo instruments during their first three observing runs [1, 2]. We have imposed a

FAR threshold of $< 0.25 \text{ yr}^{-1}$ across all analyses incorporating NS binaries and a lower threshold $\text{FAR} < 1 \text{ yr}^{-1}$ for BBH analyses. This excludes several events which pass the threshold of $p_{\text{astro}} > 0.5$ for inclusion in GWTC-3. In addition, a number of analyses of the public GW data from O1, O2 and O3a [24, 66–70] have identified additional candidate binary merger events. In the remainder of the paper, we have restricted the primary analysis to events included in GWTC-3. The overriding reason for this is that differences in the analysis methods prevent a detailed evaluation of search sensitivity, as described in Section III, which is critical to interpreting the population. In this Section, we investigate the consistency of the remaining GWTC-3 events and additional non-GWTC events with the population models inferred in this paper.

For concreteness, when referring to results reported by external groups, we include all events identified as GWs in their catalogs. In O1, there is one additional event, GW151216 identified in [67]. The additional events from O2 are GW170304, GW170425, and GW170403 which are identified in [66, 68], and GW170121, GW170202, and GW170727 which were also then independently found in open gravitational-wave catalog two (2-OGC) [70]. In O3, we include 16 additional events. These include GW190916_200658 and GW190926_050336 which were originally identified in 3-OGC [24] and independently identified in Gravitational Wave Transient Catalog 2.1 (GWTC-2.1) [4]; GW190403.051519, GW190426.190642 and GW190514.065416 which are included in GWTC-2.1 but have a FAR below our $< 1 \text{ yr}^{-1}$ threshold; GW191113.071753, GW191126.115259, GW191204.110529, GW191219.163120, GW200208.222617, GW200210.092254, GW200220.061928, GW200220.124850, GW200306.093714, GW200308.173609, GW200322.091133 which are included in GWTC-3 but again have a FAR below our $< 1 \text{ yr}^{-1}$ threshold.

In Figure 22, we show the additional gravitational wave events which were not included in the sample used in this paper. The additional events are broadly consistent with the population presented here although several events lie at the boundaries of the identified population. Specifically, two of the events have aligned spins that lie outside the inferred population. These are GW151216 with a mean $\chi_{\text{eff}} = 0.82$ and GW170403 with a mean $\chi_{\text{eff}} = -0.58$. The analysis in [68] used a prior which is constant in χ_{eff} which is significantly different from the uniform in spin-magnitude prior used in the GWTC papers. A re-analysis of GW151216 and GW170403 [199] leads to inferred χ_{eff} distributions which are more consistent with the population inferred here. Specifically, this gives $\chi_{\text{eff}} = 0.5^{+0.2}_{-0.5}$ for GW151216 and $\chi_{\text{eff}} = -0.2^{+0.4}_{-0.3}$ for GW170403. In addition, the sub-threshold events from GWTC-3 extend the distribution to both higher masses and higher mass ratios. However, only low significance events currently populate these regions. Additional observations in future runs will allow us to deter-

mine whether these low significance events are more likely spurious, or were the first hints of a broader population in the mass space.

With regard to events potentially containing NS, GWTC-3 contains several candidates that do not satisfy our $\text{FAR} < 1/\text{yr}$ threshold but do have m_2 potentially consistent with NS masses, namely GW191219 and GW200210. Both events are inferred to have highly asymmetric masses and could possibly be an indication of additional NSBH sources, or asymmetric BBH similar to GW190814. Further observations in future runs will again allow us to investigate these interesting regions of the binary parameter space in greater detail.

IX. ASTROPHYSICAL INTERPRETATION

A. Implications for binary black hole formation

1. Mass distribution

The statistical distribution of BH source properties such as their mass, spin and redshift can be used to probe the astrophysics of BH binary formation and evolution [94, 108, 110, 172, 180, 184, 196, 197, 200–211]. The analysis performed in Sec. VI has identified structures in the mass distribution of BBHs that go beyond a standard power-law model and can help to shed light on formation processes. These features were previously identified in [45], but we are now more confident that they are statistically significant (see Sec. VI and Appendix E2).

The underlying mass distribution of BBHs inferred in this paper peaks at a primary mass $\sim 10M_{\odot}$, with the majority of BBHs having a primary BHs with a mass lower than this value (e.g., see Fig. 11). Formation in globular clusters has been long recognized as an important channel for merging BBHs [212–220]. In this scenario, BBHs are assembled during three body dynamical interactions in a low metallicity environment. The resulting BH mass distribution is generally predicted to peak at $> 10M_{\odot}$. Three recent studies of globular cluster formation find that the BBH merger rate is severely suppressed where we observe a peak: one study [160] finds that the BBH merger rate is severely suppressed below about $m \simeq 13M_{\odot}$ with a corresponding realistic merger rate at this mass value of $\sim 0.5 \text{ Gpc}^{-3}\text{yr}^{-1}M_{\odot}^{-1}$ (see their Fig. 2); another recent study [221] finds similar results, with a peak in their mass distribution at about $m \simeq 15\text{--}20M_{\odot}$ (see their Fig. 5); a third analysis [220] finds the peak at $m \simeq 20M_{\odot}$. Taking these results at face value, the inferred high merger rate of sources with $\lesssim 10M_{\odot}$ may suggest that globular clusters contribute subdominantly to the detected population. Dynamical formation in young clusters is also disfavored to explain the whole BH population at $m \simeq 10M_{\odot}$ because lighter BHs are ejected by supernova kicks and do not participate to the dynamical evolution of the cluster [222–224].

Galactic nuclei can produce a BBH population with a




FIG. 22. The measured properties of the BBH candidates not included in the population study presented in this paper (shaded regions), compared to the inferred population from the PP model presented in Section VI A (black contours). These include both events which fall below our FAR threshold as well as events identified by other groups. The events are color coded based upon the search which first identified them: catalogs from O1 and O2 [66–68] in red, events from 3-OGC (which incorporates events in O1–O3a) [24] in green, and from GWTC-3 with FAR > 1 yr⁻¹ threshold in blue.

much wider mass spectrum than both young and globular clusters [225–231]. Because of their high metallicities and escape velocities, nuclear star clusters can form and retain a significant number of lighter BHs, which can then pair and merge. BBH formation near an AGN disk can produce a significant population of BBH mergers with a wide mass spectrum [162, 194, 197, 232, 233]. In such scenarios, the observed low mass overdensities without counterparts in spin could be reflections of supernova physics; by contrast, in these hierarchical formation models no evident mechanism can impart them without a corresponding signature in spin. If the BBHs are formed near an AGN disk, this process might select heavier BHs, hardening the BBH mass function and driving the peak of the mass distribution towards values higher than observed [234].

Isolated binary evolution models often predict a peak near $m \simeq 10M_{\odot}$ [172, 235–238]. Recent population models find component masses of merging BBHs that peak at 8–10 M_{\odot} and come from ~ 20 –30 M_{\odot} progenitors [171, 172, 237]. The overall merger rate normalization of the peak remains, however, poorly constrained. Moreover, the peak of the mass distribution can shift significantly depending on the adopted supernova, natal kick, mass transfer, and wind prescriptions, and star formation history of the Universe [32, 164, 171, 222, 235, 239, 240].

The analysis in Section VI suggests two additional peaks in the mass distribution at $m \sim 17M_{\odot}$ and at 35 M_{\odot} . The three most significant mass peaks are therefore separated by roughly a factor of two from each other [45]. Assuming these peaks exist, an explanation consistent with our constraints on BH spins is that they originate either from the initial BH mass function, or that they are produced by different populations formed by separate physical processes or formation channels.

The other feature of the inferred BH mass distribution that was shown in our analysis is the apparent lack of truncation at $m \sim 40M_{\odot}$, which confirms our re-

sults based on GWTC-1 and GWTC-2 [2, 25]. A mass gap between approximately 50 $^{+20}_{-10}M_{\odot}$ and $\sim 120M_{\odot}$ is predicted by stellar evolution models as the result of the pair instability process in the cores of massive stars [144–150]. However, due to our limited knowledge of the evolution of massive stars, the formation of BHs heavier than $\sim 40M_{\odot}$ from stellar collapse cannot be fully excluded [150, 172, 241–244]. The location of the mass gap is sensitive to the uncertain $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ reaction rate, which governs the production of oxygen at the expense of carbon [146, 241, 245]. Moreover, BHs formed from progenitor stars at low metallicities ($Z/Z_{\odot} \lesssim 0.1$) might avoid all together the mass limit imposed by pair-instability [174, 243, 244]. The lack of a sharp truncation at high masses might indicate a dynamical process, such as the hierarchical merger of BHs [108, 151, 195, 196, 198, 221, 226, 230, 246–250] or stars [251–254] in dense clusters or in the gaseous disk surrounding a massive BH [205]. In a hierarchical scenario we would expect the more massive BHs to also have the larger spins [195, 196]. While we do observe such a mass-spin correlation above $m \sim 40M_{\odot}$ (Fig. 19), the binaries with a signature that χ_{eff} is not zero all prefer $\chi_{\text{eff}} > 0$ (see Table IV), while hierarchical formation in dynamical environments would lead to isotropically oriented spins. BHs can also increase their birth mass beyond the pair-instability mass gap through the efficient accretion of gas from a stellar companion or from a surrounding gaseous disk [255–259]. Highly coherent accretion on one of the BHs could also explain the negative correlation between χ_{eff} and q shown in Fig. 20, although accretion in gas-free scenarios should be highly super-Eddington in order to impart significant spin [257, 259]. Alternatively, primordial BHs can have masses above the pair-instability mass threshold, although this most likely requires efficient accretion before the reionization epoch in order not to violate current constraints [260].

2. Redshift distribution

In Section VI we showed that the BBH merger rate increases with redshift, as $(1+z)^\kappa$, with $\kappa \sim 3$. Although error bars are large, current data prefer a model in which the merger rate evolves steeply with redshift and at a rate that is consistent with the growth in star formation. For binary formation in the field, the predicted value of κ is sensitive to the assumed efficiency of common envelope ejection: values between $\kappa = 0.2$ and 2.5 are all possible, although relatively small values $\kappa \sim 1$ are preferred [163, 164, 171, 222, 261]. Delay times in the field are also dependent upon stability of mass transfer, e.g., [262]. Similarly, $\kappa \lesssim 2$ is often found in models of BBHs formation in open and young clusters [224, 263]. Dynamical formation in globular clusters predicts $\kappa \lesssim 2$ [158, 264], e.g., [160] find $\kappa = 1.6_{-0.6}^{+0.4}$, and show that the most important parameter affecting the value of κ in the globular cluster scenario is the initial cluster half-mass density, ρ_h , while uncertainties in other model parameters (e.g., natal kicks, black hole masses, metallicity) have a small effect. Only models in which globular clusters are formed with a high half-mass density, $\rho_h > 10^5 M_\odot \text{pc}^{-3}$, lead to $\kappa \gtrsim 2$ [160]. While uncertainties are large, improved constraints on the merger rate evolution have the potential to unveil important information about the physics of massive binaries [96, 172, 235, 261] and the initial conditions of clusters across cosmic time [158, 211, 264].

3. Spin distribution

We observe evidence that the spin distribution both requires spin-orbit misalignment and also includes events with anti-aligned spins. BBHs with a large spin-orbit misalignment can be formed in dynamical environments such as globular, nuclear, and young star clusters, or active galactic nuclei [179, 198, 226, 229]. In these systems, two single BHs are paired together during a three body interaction and/or undergo a number of subsequent dynamical interactions before merging. Their spins have directions that are therefore uncorrelated with each other and with the orbital angular momentum of the binary, leading to an isotropic spin-orbit alignment [177, 179]. The evolution of BH spins in AGN disks depends on several uncertain factors, such as the importance of accretion and dynamical encounters, the initial spin orientation, and the efficiency of migration [197, 265]. If radial migration of BHs is inefficient, the distribution of χ_{eff} skews toward higher values because scattering encounters that randomize spin directions become less frequent. On the other hand, efficient migration would imply more frequent dynamical encounters, producing a χ_{eff} distribution centered around zero [265]. However, the dispersion of the χ_{eff} distribution also increases characteristically with mass, as with other hierarchical formation scenarios [108, 194, 197].

Formation from field binaries is thought to produce

components with preferentially aligned spins [177, 179, 183]. Such an alignment, however, is not certain. In fact, all population models of isolated binaries customarily start with the stellar progenitor spins initially perfectly aligned with the orbital angular momentum of the binary. This assumption is made due to simplicity and partly because tidal interactions are thought to quickly remove any spin-orbit misalignment prior to BH formation. However, the observational evidence of close massive binaries with highly inclined spin axes suggests that close massive binaries can form with misaligned spins and that tides might not in all cases be able to realign the spins [266–269]. Moreover, a large spin-orbit misalignment can be produced if a binary is the inner component of a triple system [264, 270–273], where the tertiary component can be either a star, another BH, or even a massive BH [228, 274–276]. In this scenario, the secular gravitational interaction of the binary with an external companion can randomize the orbital plane of the binary. The complex precessional dynamics of the BH spins in triple systems also changes the spin orientation and lead to a distribution of χ_{eff} peaked near zero, although with a marginal preference for aligned spins [270, 271]. A spin-orbit misalignment can also be produced in field binaries by a stable episode of mass-transfer prior to the formation of the BHs [277] or by asymmetric mass and neutrino emission during core-collapse [190, 278], although these processes are unlikely to produce a large misalignment for a significant fraction of the population [182, 201, 279]. We conclude that the presence of systems with misaligned spins is not in contradiction with a scenario in which the majority, if not all, BBHs form in the field of galaxies. On the other hand, the fact that the χ_{eff} distribution is not symmetric around zero, if confirmed, can be used to rule out a model in which all BBHs are formed through dynamical encounters in star clusters [191, 206].

Corroborating our previous conclusion based on GWTC-2, we find that the BH population is typically described by small spins. Predictions for BH spin magnitudes vary depending on the assumptions about stellar winds and their metallicity dependence, tides, and are particularly sensitive to the efficiency of angular momentum transport within the progenitor star [185, 186]. If the stellar core remains strongly coupled to the outer envelope during the stellar expansion off the main sequence, then a significant amount of spin can be carried from the core to the envelope. In this case, a BH formed from stellar collapse may be born with nearly zero spins. For formation in isolated binaries, this implies that the first-born BH will essentially be a Schwarzschild black hole. The second-born BH can still form with significant spin as tidal interactions may realign and increase the spin of its stellar progenitor in between the two supernova explosions [280, 281]. If the binary undergoes chemically homogeneous evolution [282, 283], its components may both be tidally spun up to near break-up velocity, and keep this rotation rate throughout main sequence evolution, evolving into BHs with large and aligned spins.

Black holes that form during the QCD phase transition in the early Universe will all have essentially zero natal spins [284, 285]. However, a significant spin can be attained through subsequent gas accretion [260]. Finally, if BBHs are formed or migrate within the accretion disk of a supermassive BH, they can accrete from the surrounding gaseous environment and spin up [265].

We observe evidence for an increase in spin magnitude for systems with higher masses [205, 286, 287] and more unequal mass ratios [193]. Current stellar evolution models suggest that larger BH masses should correlate with smaller spins because larger BHs originate from more massive stars which undergo more extensive mass loss, carrying away most of the angular momentum and producing BHs with small spins [172, 183, 280]. A consequence of this should be either a decrease in spin magnitude (and χ_{eff}) with mass above $\sim 20M_{\odot}$ or no correlation, where the predicted trend depends on the specific stellar evolution models adopted [172, 183]. Predictions remain uncertain and are strongly dependent on modeling assumptions about the angular momentum transport within the star, spin dissipation during the supernova and the treatment of binary interaction prior to BH formation. The observed increase in spin magnitude with mass might suggest a hierarchical formation scenario. However, as mentioned above, this scenario seems currently at odds with the fact that binaries with more unequal mass ratios and massive components exhibit preferentially positive χ_{eff} .

B. Implications for neutron stars

One result from gravitational wave (GW) observations is tension with the strong preference for $1.35M_{\odot}$ mass objects which has been recovered in galactic BNS [40]. Instead, conservatively assuming all objects below the maximum neutron star mass are neutron stars, our unmodeled analysis of the lowest-mass compact objects is consistent with a broad unimodal Gaussian, allowing for highly asymmetric binaries. Our analysis of all individual low-mass (assumed NS) objects suggests a wide NS mass distribution, without the bimodal structure seen in the galactic NS population. The GW-observed population of low-mass mergers is still small. If this tension persists, however, several avenues exist to explain a discrepancy between Galactic and GW observations, including but not limited to additional formation channels for GW systems; strong observational selection effects, like those used to explain GW190425 [21, 288, 289] and the smaller body in GW200105; or the prospect that BHs form below the maximum NS mass.

Our conclusions about the compact object mass spectrum in general and the mass spectrum of NS in particular will have substantial impact on the understanding of the stellar explosions that generate such compact objects [32, 290, 291] and the binary interactions that carry these objects towards merger, assuming a stellar origin

for low-mass binary mergers.

Our analyses show no evidence for or against the presence of a mass distribution feature closely corresponding to the maximum neutron star mass. Rather, the shape of the neutron star mass distribution, the existence of GW190814, and our results for the mass distribution for compact objects between $3M_{\odot}$ and $7M_{\odot}$ may instead suggest a continuous mass spectrum, albeit strongly suppressed above the masses of known NS.

Fortunately, the comparatively high prevalence of objects close to the maximum neutron star mass suggests that we will likely observe several objects near this region in the future, potentially providing several avenues to connect features in the NS mass distribution to fundamental nuclear physics. Our analysis of NS in merging binaries alone suggests the NS mass distribution extends to the maximum NS mass M_{TOV} expected from the EOS.

Our analyses are also consistent with both symmetric ($q \simeq 0.8$) and significantly asymmetric ($q < 0.8$) binaries containing NS in BNS, and modestly ($q \in [0.5, 0.8]$) to highly ($q < 0.5$) asymmetric binaries in NSBH. Compared to equal-mass mergers [292, 293], modestly asymmetric NS mergers (with either NS or BH counterparts) are potentially strong candidates for multimessenger counterparts [294], since an asymmetric merger can eject more mass [295], produce a larger remnant disk [296, 297], and potentially produce significant associated gamma ray burst emission [298–300]. For BNS, our analyses are consistent with a significant fraction of highly asymmetric events. For NSBH, the discovery of GW200105 and GW200115 demonstrate the existence of asymmetric binaries containing NS with a range of mass ratios. Based on these events, our inferences about the low-mass compact object distribution suggests that EM-bright NSBH mergers could occur at a significant fraction of the overall NSBH rate. Generally, a broad mass ratio distribution suggests modestly more favorable prospects for electromagnetic follow-up observations. Conversely, a broad mass ratio distribution complicates simple efforts to interpret existing GW observations which were developed under the assumption that low-mass binary mergers are very frequently of comparable mass [301].

Finally, our analyses here leave GW190814 as an outlier both from BBH systems and from systems that contain a likely NS. Neither component of this binary has exceptional masses; for example, the secondary component could easily be produced from conventional supernova engines [32]. However, based on the merger rates versus mass identified in our study, this system (and the larger sample of high-mass-ratio binaries available at a lower threshold) may require a different formation pathway [229, 302, 303].

X. THE GW BACKGROUND FROM BINARY MERGERS

The observation of binaries with masses in the NSBH range allows us to provisionally complete a census of the different classes of compact binaries that contribute to an astrophysical gravitational-wave background, assuming our existing surveys are sensitive to all relevant sources (i.e., not accounting for frequent subsolar mass mergers). We have previously predicted the contributions of BBH and BNS mergers to the gravitational-wave background, based on the compact binary population observed in GWTC-2 [169]. In Fig. 23, we update this forecast with our latest knowledge of the BBH and BNS population and the newly measured rate of NSBH mergers.

The shaded bands on the left side of Fig. 23 shows estimates of and uncertainties on the dimensionless energy-density spectra

$$\Omega(f) = \frac{1}{\rho_c} \frac{d\rho}{d \ln f} \quad (21)$$

of gravitational waves radiated by each class of compact binary. In Eq. (21), $d\rho$ is the gravitational-wave energy density per logarithmic frequency interval $d \ln f$, while ρ_c is the critical energy density required to close the Universe. We adopt the same model for the merger history of compact binaries used previously [169, 304], assuming that compact binary formation rate traces a metallicity-weighted star formation rate model [305–307] with a $p(t_d) \propto t_d^{-1}$ distribution of time delays t_d between binary formation and merger. Time delays are restricted to $20 \text{ Myr} \leq t_d \leq 13.5 \text{ Gyr}$ for BNS and NSBH mergers and $50 \text{ Myr} \leq t_d \leq 13.5 \text{ Gyr}$ for BBHs [171, 308], with binary formation restricted to redshifts below $z_{\text{max}} = 10$. The birth rate of BBH progenitors is further weighted by the fraction of star formation at metallicities $Z < 0.1 Z_\odot$ [170, 309].

Within Fig. 23, the stochastic energy-density due to BBHs has been marginalized over our uncertainty on both the local merger rate and mass distribution of the BBH population, as measured using the PP mass model. At 25 Hz, we estimate the energy-density due to BBHs to be $\Omega_{\text{BBH}}(25 \text{ Hz}) = 4.9_{-1.5}^{+1.5} \times 10^{-10}$. Contributions due to BNS and NSBH systems, meanwhile, are estimated adopting the BGP rates reported in Table II. For simplicity, we assume a uniform distribution of neutron star masses between 1 and $2.5 M_\odot$, and a logarithmically uniform distribution of black hole masses between 5 and $50 M_\odot$ among NSBH binaries. We find $\Omega_{\text{BNS}}(25 \text{ Hz}) = 0.6_{-0.5}^{+2.4} \times 10^{-10}$ and $\Omega_{\text{NSBH}}(25 \text{ Hz}) = 1.0_{-0.7}^{+2.4} \times 10^{-10}$.

The blue band in the right side of Fig. 23 denotes the our estimate of the total gravitational-wave background due to the superposition of these three source classes; we expect a total energy-density of $\Omega(25 \text{ Hz}) = 6.9_{-2.0}^{+3.5} \times 10^{-10}$. For comparison, the solid black curve marks our present sensitivity to the gravitational-wave background [169, 310]. Although our estimate for the

background amplitude lies well below current limits, it may be accessible with future detectors, such as the planned “A+” LIGO configuration.

XI. CONCLUSIONS

The third LIGO–Virgo gravitational wave transient catalog (GWTC-3) [1] has increased our census of the population of compact mergers by nearly a factor of two, compared to our analysis of the first half of O3 [11]. We simultaneously employ all observations with $\text{FAR} < 0.25 \text{ yr}^{-1}$ to infer the merger rate versus both component masses across the observed mass spectrum. For NS, we find a broad mass distribution, extending up to $2.0_{-0.2}^{+0.3} M_\odot$, in contrast to the unimodal mass distribution observed for Galactic BNS. We find the BBH mass distribution is nonuniform, with overdensities at BH masses of $10 M_\odot$ and $35 M_\odot$. These overdensities may reflect the astrophysics associated with generating coalescing binaries, potentially reflecting properties of stellar physics or astrophysical environments. These features may assist future applications of gravitational wave astronomy. As an example, these sharp features could be redshift-independent and, if so, used as standard candles for cosmology [311, 312]. We find the BH mass distribution exhibits an interval between $2.1 M_\odot$ and $5.8 M_\odot$ where merger rates are suppressed, which could be consistent with past X-ray observations suggesting a mass gap [28–31]. Our analysis lacks sufficient sensitivity to probe the structure of the mass distribution at the highest masses $m_1 > 70 M_\odot$ in detail; however, so far, we find no evidence for or against an upper mass gap.

We find that observed BH spins are typically small (half less than 0.26). We still conclude that at least some of these spins exhibit substantial spin–orbit misalignment. We corroborate a correlation between BBH effective aligned inspiral spins and mass ratio.

Using parametric models to infer the distribution of BBH merger rate with redshift, we find the BBH merger rate likely increases with redshift; we cannot yet assess more complex models where the shape or extent of the mass distribution changes with redshift.

Analyses presented in our previous work [11] and in a companion paper [313] employ coarse-grained models for the BBH population, smoothing over some of the subtle features identified above. We find that these coarse-grained models draw similar conclusions on current data as our previous studies; see Sec. VI A. Applications that focus on large-scale features of the mass distribution (e.g., the stochastic background, as described in Sec. X) only require these coarse-grained results. Nonetheless, the mass distribution remains a critical source of systematic uncertainty in any merger rate integrated over any mass interval, particularly in mass intervals with few observations. We specifically find the BNS and NSBH merger rates exhibit considerable uncertainty in the mass distribution, with relative merger rate errors within (and




FIG. 23. Forecast of astrophysical gravitational-wave background due to binary mergers following O3. (*Left*): The individual contributions expected from BNS, NSBH and BBH mergers. While uncertainties on the energy-density due to BNS and NSBH are due to Poisson uncertainties in their merger rates, our forecast for the stochastic background due to BBHs additionally includes systematic uncertainties associated with their imperfectly known mass distribution. (*Right*): Estimate of the total gravitational-wave background (blue), as well as our experimental current sensitivity (solid black) [169, 310]. For comparison, we additionally show the expected sensitivities of the LIGO-Virgo network at design sensitivity, as well as that of LIGO’s anticipated “A+” configuration.

between) models far in excess of the expected statistical Poisson error associated with the count of these events. These systematics propagate directly into our most conservative estimates for their merger rates.

The next GW survey could have a BNS detection range increased by approximately 15–40% [314]. Even without allowing for increased merger rates at higher redshift, the next survey should identify roughly 3 times more events of each class than used in this study, including several new events from the BNS and BHNS category. We continuously revise our assessment of future observing prospects [314].

ACKNOWLEDGMENTS

This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation. The authors also gratefully acknowledge the support of the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Netherlands Organization for Scientific Research (NWO), for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Council

of Scientific and Industrial Research of India, the Department of Science and Technology, India, the Science & Engineering Research Board (SERB), India, the Ministry of Human Resource Development, India, the Spanish Agencia Estatal de Investigación (AEI), the Spanish Ministerio de Ciencia e Innovación and Ministerio de Universidades, the Conselleria de Fons Europeus, Universitat i Cultura and the Direcció General de Política Universitaria i Recerca del Govern de les Illes Balears, the Conselleria d’Innovació, Universitats, Ciència i Societat Digital de la Generalitat Valenciana and the CERCA Programme Generalitat de Catalunya, Spain, the National Science Centre of Poland and the European Union – European Regional Development Fund; Foundation for Polish Science (FNP), the Swiss National Science Foundation (SNSF), the Russian Foundation for Basic Research, the Russian Science Foundation, the European Commission, the European Social Funds (ESF), the European Regional Development Funds (ERDF), the Royal Society, the Scottish Funding Council, the Scottish Universities Physics Alliance, the Hungarian Scientific Research Fund (OTKA), the French Lyon Institute of Origins (LIO), the Belgian Fonds de la Recherche Scientifique (FRS-FNRS), Actions de Recherche Concertées (ARC) and Fonds Wetenschappelijk Onderzoek – Vlaanderen (FWO), Belgium, the Paris Île-de-France Region, the National Research, Development and Innovation Office Hungary (NKFIH), the National Research Foundation of Korea, the Natural Science and Engineering Research Council Canada, Canadian Foundation for Innovation (CFI), the Brazilian Ministry of Science, Technology, and Innovations, the International Center for Theoretical Physics South American Institute for Fundamental Research (ICTP-SAIFR), the Research Grants Council

of Hong Kong, the National Natural Science Foundation of China (NSFC), the Leverhulme Trust, the Research Corporation, the Ministry of Science and Technology (MOST), Taiwan, the United States Department of Energy, and the Kavli Foundation. The authors gratefully acknowledge the support of the NSF, STFC, INFN and CNRS for provision of computational resources. This work was supported by MEXT, JSPS Leading-edge Research Infrastructure Program, JSPS Grant-in-Aid for Specially Promoted Research 26000005, JSPS Grant-in-Aid for Scientific Research on Innovative Areas 2905: JP17H06358, JP17H06361 and JP17H06364, JSPS Core-to-Core Program A. Advanced Research Networks, JSPS Grant-in-Aid for Scientific Research (S) 17H06133 and 20H05639, JSPS Grant-in-Aid for Transformative Research Areas (A) 20A203: JP20H05854, the joint research program of the Institute for Cosmic Ray Research, University of Tokyo, National Research Foundation (NRF) and Computing Infrastructure Project of KISTI-GSDC in Korea, Academia Sinica (AS), AS Grid Center (ASGC) and the Ministry of Science and Technology (MoST) in Taiwan under grants including AS-CDA-105-M06, Advanced Technology Center (ATC) of NAOJ, Mechanical Engineering Center of KEK. *We would like to thank all of the essential workers who put their health at risk during the COVID-19 pandemic, without whom we would not have been able to complete this work.*

Appendix A: Sensitivity estimation

A key ingredient in Eqs. (3) and (4) is the detection fraction $\xi(\Lambda)$, which estimates the fraction of systems that we expect to successfully detect from some prior volume that extends past our detector’s reach. The detection fraction quantifies selection biases, and so it is critical to accurately characterize. For a population described by parameters Λ , the detection fraction is

$$\xi(\Lambda) = \int P_{\text{det}}(\theta)\pi(\theta|\Lambda)d\theta. \quad (\text{A1})$$

Here, $P_{\text{det}}(\theta)$ is the detection probability: the probability that an event with parameters θ would be detected by a particular search. The detection probability depends on the angular/sky position and orientation of the source binary, and crucially for our purposes, on the masses and redshift of a system, and, to a lesser degree, on the spins.

Given the non-ideal nature of the detector data, the variation in network sensitivity over time, and the complexity of both the signal waveforms and the search pipelines, an accurate estimate of $P_{\text{det}}(\theta)$ and $\xi(\Lambda)$ requires empirical methods, specifically the use of a large suite of simulated signals added to the data: injections. For analyses that focus on the BBH subpopulation in Section VI, we simulate compact binary signals from a reference BBH population and record which ones are successfully detected by the PyCBC, GstLAL or MBTA

search pipelines. We omit the cWB search from our volume estimate, since at present any detection of a binary merger was corroborated by a detection in the remaining pipelines. In addition, we also simulate compact binary signals from reference BNS, NSBH and IMBH populations. These injections include binaries with component masses in the range $1\text{--}600M_{\odot}$, have spins that are isotropic in orientation and are uniform in comoving volume. Spins are drawn from a distribution that is uniform in the dimensionless spin magnitude up to a maximum of $\chi_{\text{max}} = 0.998$ for black holes and $\chi_{\text{max}} = 0.4$ for neutron stars. To control computational costs, the expected network signal to noise ratio (SNR) of each injection is computed using representative detector power spectral densities (PSDs) for O3. Injections with expected SNR below 6 are assumed not to be detected, and are thus removed from the set analyzed by the search pipelines. A thorough description of the injections and their underlying probability distribution is available in [315]. These injections are then combined into a single dataset as a mixture model [1, 315] in order to assess sensitivity across the entire parameter space and subpopulations. Our analyses in Sec. IV make use of these injections to estimate sensitivity.

Unlike previous synthetic simulation sets used in our population analysis following GWTC-2 [11], the injections used here model spins that are isotropically distributed in orientation and hence allow for orbital precession. Further, the maximum spin magnitude we assume for NS components, 0.4, is significantly larger than for previous injection sets used to estimate BNS merger rates [21]. That said, our injections have an effective χ_{eff} distribution that is narrow and centered at 0 while analyses using BNS populations with small NS spins inherently have $\chi_{\text{eff}} \approx 0$. Because the merger rate depends on spins primarily through the system’s χ_{eff} , the specific assumptions made about the spin distribution at low mass have modest impact on the inferred low-mass merger rate.

Following [94, 95, 316, 317], the point estimate for Eq. (A1) is calculated using a Monte Carlo integral over found injections:

$$\hat{\xi}(\Lambda) = \frac{1}{N_{\text{inj}}} \sum_{j=1}^{N_{\text{found}}} \frac{\pi(\theta_j | \Lambda)}{p_{\text{draw}}(\theta_j)}, \quad (\text{A2})$$

where N_{inj} is the total number of injections, N_{found} are the injections that are successfully detected, and p_{draw} is the probability distribution from which the injections are drawn. When using this approach to estimate sensitivity, we marginalize over the uncertainty in $\hat{\xi}(\Lambda)$ and ensure that the effective number of found injections remaining after population re-weighting is sufficiently high ($N_{\text{eff}} > 4N_{\text{det}}$) following [317]. We also compute (and some analyses like MS employ) semi-analytic approximations to the integrated network sensitivity $VT(\theta, \kappa) = \int dt dz dV_c / dz / (1+z) \langle P_{\text{det}}(\theta, z) \rangle (1+z)^\kappa$ for fiducial choices of κ , appropriate to characterize sensitivity to a population with a fixed redshift evolution.

Parameter	Description	Prior
α	Spectral index for the power-law of the primary mass distribution.	$U(-4, 12)$
β_q	Spectral index for the power-law of the mass ratio distribution.	$U(-2, 7)$
m_{\min}	Minimum mass of the power-law component of the primary mass distribution.	$U(2 M_{\odot}, 10 M_{\odot})$
m_{\max}	Maximum mass of the power-law component of the primary mass distribution.	$U(30 M_{\odot}, 100 M_{\odot})$

TABLE V. Summary of TRUNCATED model parameters.

For the O3 observing period, we characterize the found injections as those recovered with a FAR below the corresponding thresholds used in population analyses described in this paper (1 per year and 1 per 4 years) in either PYCBC, GstLAL or MBTA. For the O1 and O2 observing periods, we supplement the O3 pipeline injections with mock injections drawn from the same distribution p_{draw} above. For the mock injections, we calculate $P_{\text{det}}(m_1, m_2, z, \chi_{1,z}, \chi_{2,z})$ according to the semi-analytic approximation used in our analysis of GWTC-2 [25], based on a network signal-to-noise ratio threshold $\rho = 10$ and representative strain noise power spectral densities estimated from data recorded during the O1 and O2 observing runs. We combine O1, O2 and O3 injection sets ensuring a constant rate of injections across the total observing time [318].

Appendix B: Population Model Details

In this section we provide details about the low-dimensional parameterized population models described above in Section III. Each subsection includes a table with a summary of the parameters for that model and the prior distribution used for each parameter. The prior distributions are indicated using abbreviations: for example, $U(0, 1)$ translates to uniform on the interval $(0, 1)$, $LU(10^{-6}, 10^5)$ translates to log-uniform on the interval $10^{-6}, 10^5$, and $N(0, 1)$ translates to a Gaussian distribution with mean 0 and standard deviation of 1.

For the TRUNCATED, POWER LAW + PEAK, POWER LAW + SPLINE, POWER LAW + DIP + BREAK, DEFAULT, GAUSSIAN, and POWER-LAW population models, to ensure convergence of the Monte Carlo integral used to compute Equation 4 we only assign nonzero likelihoods to points in parameter space with an effective sample size of at least the number of observed events in our event list [317]. While this restriction effectively enforces a data-dependent constraint on the prior relative to what is listed here, it prevents the inclusion of biased samples in our inferred posterior distributions which may otherwise contain relics of unconverged Monte Carlo integrals.

1. Details of mass population models

a. TRUNCATED mass model

TRUNCATED mass model serves as the primary component for some of our mass models. The primary mass distribution for this model follows a power-law with spectral index α , and with a sharp cut-off at the lower end m_{\min} and the upper end of the distribution m_{\max} :

$$\pi(m_1 | \alpha, m_{\min}, m_{\max}) \propto \begin{cases} m_1^{-\alpha} & m_{\min} < m_1 < m_{\max} \\ 0 & \text{otherwise,} \end{cases} \quad (\text{B1})$$

Meanwhile, the mass ratio $q \equiv m_2/m_1$ follows a power-law distribution with spectral index β_q

$$\pi(q | \beta_q, m_{\min}, m_1) \propto \begin{cases} q^{\beta_q} & m_{\min} < m_2 < m_1 \\ 0 & \text{otherwise.} \end{cases} \quad (\text{B2})$$

The parameters for this model are summarized in Table V. For this model, as well as further mass models where a prior on the total merger rate is not specified, the rate prior is proportional to $1/R$, or equivalently to $1/N$ in the notation of Eq.3-4.

Parameter	Description	Prior
α	Spectral index for the power-law of the primary mass distribution.	$U(-4, 12)$
β_q	Spectral index for the power-law of the mass ratio distribution.	$U(-2, 7)$
m_{\min}	Minimum mass of the power-law component of the primary mass distribution.	$U(2 M_{\odot}, 10 M_{\odot})$
m_{\max}	Maximum mass of the power-law component of the primary mass distribution.	$U(30 M_{\odot}, 100 M_{\odot})$
λ_{peak}	Fraction of BBH systems in the Gaussian component.	$U(0, 1)$
μ_m	Mean of the Gaussian component in the primary mass distribution.	$U(20 M_{\odot}, 50 M_{\odot})$
σ_m	Width of the Gaussian component in the primary mass distribution.	$U(1 M_{\odot}, 10 M_{\odot})$
δ_m	Range of mass tapering at the lower end of the mass distribution.	$U(0 M_{\odot}, 10 M_{\odot})$

TABLE VI. Summary of POWER LAW + PEAK model parameters.

b. POWER LAW + PEAK *mass model*

This is equivalent to Model C from [25]. The primary mass distribution is a truncated powerlaw, with the addition of tapering at the lower mass end of the distribution and a Gaussian component:

$$\pi(m_1 | \lambda_{\text{peak}}, \alpha, m_{\min}, \delta_m, m_{\max}, \mu_m, \sigma_m) = \left[(1 - \lambda_{\text{peak}}) \mathfrak{P}(m_1 | -\alpha, m_{\max}) + \lambda_{\text{peak}} G(m_1 | \mu_m, \sigma_m) \right] S(m_1 | m_{\min}, \delta_m). \quad (\text{B3})$$

Here, $\mathfrak{P}(m_1 | -\alpha, m_{\max})$ is a normalized power-law distribution with spectral index $-\alpha$ and high-mass cut-off m_{\max} . Meanwhile, $G(m_1 | \mu_m, \sigma_m)$ is a normalized Gaussian distribution with mean μ_m and width σ_m . The parameter λ_{peak} is a mixing fraction determining the relative prevalence of mergers in \mathfrak{P} and G . Finally, $S(m_1, m_{\min}, \delta_m)$ is a smoothing function, which rises from 0 to 1 over the interval $(m_{\min}, m_{\min} + \delta_m)$:

$$S(m | m_{\min}, \delta_m) = \begin{cases} 0 & (m < m_{\min}) \\ [f(m - m_{\min}, \delta_m) + 1]^{-1} & (m_{\min} \leq m < m_{\min} + \delta_m) \\ 1 & (m \geq m_{\min} + \delta_m) \end{cases}, \quad (\text{B4})$$

with

$$f(m', \delta_m) = \exp\left(\frac{\delta_m}{m'} + \frac{\delta_m}{m' - \delta_m}\right). \quad (\text{B5})$$

The conditional mass ratio distribution in this model also includes the smoothing term:

$$\pi(q | \beta, m_1, m_{\min}, \delta_m) \propto q^{\beta q} S(qm_1 | m_{\min}, \delta_m). \quad (\text{B6})$$

The parameters for this model are summarized in Table VI.

c. POWER LAW + SPLINE *mass model*

The POWER LAW + SPLINE mass model explicitly applies a perturbation to a modified version of the fiducial POWER LAW + PEAK model that does not include the Gaussian peak [46]. Let $p(m_1 | \alpha, m_{\min}, m_{\max}, \delta_m)$ be the modified POWER LAW + PEAK model without the Gaussian, then primary mass distribution for the POWER LAW + SPLINE model is given as:

$$p_{\text{PS}}(m_1 | \alpha, m_{\min}, m_{\max}, \delta_m, \{f_i\}) = k p(m_1 | \alpha, m_{\min}, m_{\max}, \delta_m) \exp(f(m_1 | \{f_i\})). \quad (\text{B7})$$

Above, k is a normalization factor found by numerically integrating p_{PS} over the range of allowed primary masses, $f(m_1 | \{f_i\})$ is the perturbation function we model with cubic splines, and $\{f_i\}$ are the heights of the n knots from which f is interpolated. The n knot locations are fixed, spaced linearly in $\log m_1$ space from 2–100 M_{\odot} . We additionally restrict the perturbations to converge to the underlying distribution at the boundary nodes by fixing both f_0 and f_{n-1} to be 0. We chose $n = 20$ to be the optimal number of knots for this analysis following the same procedure in [46], which adds a total of 18 additional parameters describing the perturbations to the underlying model. In

Parameter	Description	Prior
α	Spectral index for the power-law of the primary mass distribution.	$U(-4, 12)$
β_q	Spectral index for the power-law of the mass ratio distribution.	$U(-2, 7)$
m_{\min}	Minimum mass of the power-law component of the primary mass distribution.	$U(2 M_{\odot}, 10 M_{\odot})$
m_{\max}	Maximum mass of the power-law component of the primary mass distribution.	$U(30 M_{\odot}, 100 M_{\odot})$
δ_m	Range of mass tapering at the lower end of the mass distribution.	$U(0 M_{\odot}, 10 M_{\odot})$
$\{f_i\}$	Spline perturbation knot heights.	$N(0, 1)$

TABLE VII. Summary of POWER LAW + SPLINE model parameters.

Parameter	Description	Prior
w_i	Mixing weights.	Dirichlet($\boldsymbol{\alpha}$), $\alpha_{1\dots N} = 1/N$
$\mu_i^{\mathcal{M}}$	Mean of the normal distribution modeling the chirp mass.	$LU(5.2 M_{\odot}, 65 M_{\odot})$
$\sigma_i^{\mathcal{M}}$	Scale of the normal distribution modeling the chirp mass.	$U(0.02 \mu_i^{\mathcal{M}}, 0.18 \mu_i^{\mathcal{M}})$
μ_i^{sz}	Mean of the normal distribution modeling the aligned spin distribution.	$U(-0.5, 0.5)$
σ_i^{sz}	Scale of the normal distribution modeling the aligned spin distribution.	$U(0.05, 0.6)$
q_i^{\min}	Minimum value of the mass ratio.	$U(0.1, 0.95)$
α_i^q	Slope of the power-law.	$U(-7, 2)$
\mathcal{R}	Merger rate.	$LU(1, 100)$

TABLE VIII. Summary of FLEXIBLE MIXTURES model parameters. All rates are in $\text{Gpc}^{-3} \text{yr}^{-1}$.

addition to the primary mass, the conditional mass ratio distribution follows the same form as the POWER LAW + PEAK model defined in Eq. (B6). For each mass distribution inference with the POWER LAW + SPLINE model, we simultaneously fit the spin distribution with the DEFAULT model and the redshift evolution of the merger rate with the POWER LAW evolution model. The parameters and chosen prior distributions for the POWER LAW + SPLINE model are summarized in Table VII.

d. FLEXIBLE MIXTURES model

The FLEXIBLE MIXTURES model, Vamana, predicts the population using a sum of weighted components. Each component is composed of a Gaussian, another Gaussian and a power-law to model the chirp mass, the aligned spins and the mass ratio respectively. The model is defined as

$$p(\mathcal{M}, q, s_{1z}, s_{2z} | \boldsymbol{\lambda}) = \sum_{i=1}^N w_i G(\mathcal{M} | \mu_i^{\mathcal{M}}, \sigma_i^{\mathcal{M}}) G(s_{1z} | \mu_i^{sz}, \sigma_i^{sz}) \sigma_i^{\mathcal{M}} G(s_{2z} | \mu_i^{sz}, \sigma_i^{sz}) \mathcal{P}(q | \alpha_i^q, q_i^{\min}, 1), \quad (\text{B8})$$

where G is the normal distribution and \mathcal{P} is the truncated power-law. For the presented analysis we use $N = 11$ components. This choice maximises the marginal likelihood, however, the predicted population is robust for a wide range of N . For detailed description of this model see [49]. FLEXIBLE MIXTURES model uses a power-law to model the redshift evolution of the merger rate, as described in subsection B 3. The merger rate has a uniform-in-log distributed prior; the prior distributions for parameters in Eq. B8 are summarized in Table VIII.

e. BINNED GAUSSIAN PROCESS model

The BINNED GAUSSIAN PROCESS models the rate densities, $m_1 m_2 \frac{dR^i}{dm_1 dm_2} = n^i$, as a binned Gaussian Process where the index i denotes a particular bin in the two-dimensional $\log m_1 - \log m_2$ parameter space [115, 116]. The bin edges in the analysis presented in the paper are located at $[1, 2, 2.5, 3, 4, 5, 6.5, 8, 10, 15, 20, 30, 40, 50, 60, 70, 80, 100] M_{\odot}$ with the assumption that $m_2 \leq m_1$. The probabilistic model for the logarithm of the rate density in each bin is defined as

$$\log n^i \sim N(\mu, \Sigma), \quad (\text{B9})$$

Parameter	Description	Prior
μ	Mean log (Rate) in each bin.	$N(0, 10)$
σ	Amplitude of the covariance kernel.	$N(0, 10)$
$\log(l)$	log (Length scale) of the covariance kernel.	$N(-0.085, 0.93)$

TABLE IX. Summary of BINNED GAUSSIAN PROCESS model parameters.

where μ is the mean of the Gaussian process and Σ is the covariance matrix that correlates the bins. Each element of the covariance matrix Σ is generated using a squared-exponential kernel $k(x, x')$ which is defined as

$$k(x, x') = \sigma^2 \exp\left(\frac{-(x - x')^2}{2l^2}\right). \quad (\text{B10})$$

For the specific analysis here we take x, x' to be the bin centers in $\log m$. The parameter σ models the amplitude of the covariances while l is a parameter that defines the length scales over which bins are correlated. The prior distribution chosen here for the length scale is a log-normal distribution with a mean that is the average between the minimum bin spacing

$$\Delta_{\min} \equiv \min_{m_1, m_2} \Delta \log m \quad (\text{B11})$$

and the maximum bin spacing

$$\Delta_{\max} \equiv \max_{m_1, m_2} \Delta \log m \quad (\text{B12})$$

with a standard deviation of $\frac{(\Delta_{\max} - \Delta_{\min})}{4}$. This constrains (at “2- σ ” in the prior) the correlation length for the GP to lie between “one bin” and “all the bins.” For our analyses presented in the paper, the mean and standard deviation are -0.085 and 0.93 respectively. The BINNED GAUSSIAN PROCESS model assumes a redshift distribution such that the overall merger rate of compact binaries is uniform-in-comoving volume. The spin distributions for each component are isotropic in direction and uniform in the spin magnitude with a maximum spin of 0.998 for BHs and 0.4 for NSs; the prior distribution for the relevant parameters in Equations B9 and B10 is summarized in Table IX.

f. POWER LAW + DIP + BREAK mass model

The POWER LAW + DIP + BREAK model explicitly searches for separation in masses between two subpopulations by employing a broken power law with a dip at the location of the power law break. As described in [105] and [106], the dip is modeled by a notch filter with depth A , which is fit along with the other model parameters in order to determine the existence and depth of a potential mass gap. No gap corresponds to $A = 0$, whereas $A = 1$ corresponds to precisely zero merger rate over some interval. POWER LAW + DIP + BREAK also employs a low-pass filter at high masses to allow for a tapering of the mass spectrum, which has the effect of a smooth second break to the power law.

The joint mass distribution in this model has the form:

$$p(m_1, m_2) \propto p(m_1)p(m_2)(m_2/m_1)^\beta, \quad (\text{B13})$$

$$p(m) \propto p_{\text{pl}}(m)n(m)\ell(m), \quad (\text{B14})$$

$$n(m) = 1 - \frac{A}{(1 + (M_{\text{low}}^{\text{gap}} m)^{\eta_{\text{low}}})(1 + (M_{\text{high}}^{\text{gap}} m)^{\eta_{\text{high}}})}, \text{ and} \quad (\text{B15})$$

$$\ell(m) = \frac{1}{1 + (m/m_{\text{max}})^\eta}. \quad (\text{B16})$$

where $p_{\text{pl}}(m)$ is a broken power law with exponents α_1 between m_{\min} and $M_{\text{low}}^{\text{gap}}$ and α_2 between $M_{\text{low}}^{\text{gap}}$ and m_{max} . The parameters for this model are summarized in Table X.

g. Neutron star mass models

The mass models adopted for the BNS and NSBH events in Sec. V assume a basic mass distribution that is common to all NSs, with random pairing into compact binaries. The basic mass distribution is taken to be either a power law

Parameter	Description	Prior
α_1	Spectral index for the power-law of the mass distribution at low mass.	$U(-8, 2)$
α_2	Spectral index for the power-law of the mass distribution at high mass.	$U(-3, 2)$
A	Lower mass gap depth.	$U(0, 1)$
$M_{\text{low}}^{\text{gap}}$	Location of lower end of the mass gap.	$U(1.4M_{\odot}, 3M_{\odot})$
$M_{\text{high}}^{\text{gap}}$	Location of upper end of the mass gap	$U(3.4M_{\odot}, 9M_{\odot})$
η_{low}	Parameter controlling how the rate tapers at the low end of the mass gap	50
η_{high}	Parameter controlling how the rate tapers at the low end of the mass gap.	50
η	Parameter controlling tapering the truncated power law at high mass	$U(-4, 12)$
β	Spectral index for the power-law-in-mass-ratio pairing function.	$U(-2, 7)$
m_{min}	Minimum mass of the power-law component of the mass distribution.	$U(1M_{\odot}, 1.4M_{\odot})$
m_{max}	Maximum mass of the power-law component of the mass distribution.	$U(35M_{\odot}, 100M_{\odot})$

TABLE X. Summary of POWER LAW + DIP + BREAK model parameters.

Parameter	Description	Prior
α	Spectral index for the power-law in the POWER NS mass distribution.	$U(-4, 12)$
m_{min}	Minimum mass of the NS mass distribution.	$U(1.0M_{\odot}, 1.5M_{\odot})$
m_{max}	Maximum mass of the NS mass distribution.	$U(1.5M_{\odot}, 3.0M_{\odot})$
μ	Location of the Gaussian peak in the PEAK NS mass distribution.	$U(1.0M_{\odot}, 3.0M_{\odot})$
σ	Width of the Gaussian peak in the PEAK NS mass distribution.	$U(0.01M_{\odot}, 2.00M_{\odot})$

TABLE XI. Summary of POWER and PEAK NS mass model parameters.

or, inspired by the shape of the Galactic BNS mass distribution [39–41], a Gaussian. The BH mass distribution is fixed to be uniform between 3 and 60 M_{\odot} . The NS mass distribution analysis assumes definite source classifications for the events. Thus, the joint mass distribution takes the form

$$p(m_1, m_2) \propto \begin{cases} p(m_1)p(m_2) & \text{if BNS} \\ U(3M_{\odot}, 60M_{\odot})p(m_2) & \text{if NSBH,} \end{cases} \quad (\text{B17})$$

with $p(m)$ either a power law with exponent α , minimum mass m_{min} and maximum mass m_{max} , or a Gaussian with a peak of width σ at μ , plus sharp minimum and maximum mass cutoffs m_{min} , m_{max} . We call these models POWER and PEAK, respectively. Their hyper-parameters, and the choices for their prior distributions, are listed in Table XI. We additionally impose the constraint $m_{\text{min}} \leq \mu \leq m_{\text{max}}$ on the PEAK model. Besides the flat m_{max} prior described in the table, for the analyses excluding GW190814 we use a prior proportional to the cumulative distribution function of $M_{\text{max,TOV}}$, i.e., $p(m_{\text{max}}) \propto \int_{m_{\text{max}}}^{\infty} dM_{\text{max,TOV}} p(M_{\text{max,TOV}})$. This enforces our expectation that the NS masses in the gravitational-wave population should not exceed $M_{\text{max,TOV}}$.

2. Details of spin population models

a. DEFAULT spin model

This model was introduced in [25]. Following [110], the dimensionless spin magnitude distribution is taken to be a Beta distribution,

$$\pi(\chi_{1,2}|\alpha_{\chi}, \beta_{\chi}) = \text{Beta}(\alpha_{\chi}, \beta_{\chi}), \quad (\text{B18})$$

where α_{χ} and β_{χ} are the standard shape parameters that determine the distribution's mean and variance. The Beta distribution is convenient because it is bounded on (0,1). The distributions for χ_1 and χ_2 are assumed to be the same. Following [111], we define $z_i = \cos \theta_i$ as the cosine of the tilt angle between component spin and a binary's orbital angular momentum, and assume that \mathbf{z} is distributed as a mixture of two populations:

$$\pi(\mathbf{z}|\zeta, \sigma_t) = \zeta G_t(\mathbf{z}|\sigma_t) + (1 - \zeta)\mathcal{J}(\mathbf{z}). \quad (\text{B19})$$

Parameter	Description	Prior
μ_χ	Mean of the Beta distribution of spin magnitudes.	U(0,1)
σ_χ^2	Variance of the Beta distribution of spin magnitudes.	U(0.005,0.25)
ζ	Mixing fraction of mergers from truncated Gaussian distribution.	U(0,1)
σ_t	Width of truncated Gaussian, determining typical spin misalignment.	U(0.1,4)

TABLE XII. Summary of DEFAULT spin parameters.

Parameter	Description	Prior
μ_{eff}	Mean of the χ_{eff} distribution.	U(-1, 1)
σ_{eff}	Standard deviation of the χ_{eff} distribution.	U(0.05,1)
μ_p	Mean of the χ_p distribution.	U(0.05, 1)
σ_p	Standard deviation of the χ_p distribution.	U(0.05, 1)
ρ	Degree of correlation between χ_{eff} and χ_p .	U(-0.75, 0.75)
$\chi_{\text{eff,min}}$	Lower truncation bound on χ_{eff} .	U(-1, μ_{eff})
ζ	Non-vanishing mixture fraction in Eq. (18).	U(0, 1)

TABLE XIII. Summary of GAUSSIAN spin parameters. The $\chi_{\text{eff,min}}$ and ζ parameters appear only in variants of the GAUSSIAN model, as discussed below.

Here, $\mathcal{J}(z)$ is an isotropic distribution, while $G_t(z|\sigma_t)$ is a truncated two-dimensional Gaussian, peaking at $\mathbf{z} = 0$ (perfect alignment) with width σ_t . The mixing parameter ζ controls the relative fraction of mergers drawn from the isotropic distribution and Gaussian subpopulations. The isotropic subpopulation is intended to accommodate dynamically assembled binaries, while G_t is a model for field mergers. The parameters for this model and their priors are summarized in Table XII. Additional constraints to the priors on μ_χ and σ_χ^2 are applied by setting $\alpha_\chi, \beta_\chi > 1$.

b. GAUSSIAN spin model

In addition to the distribution of component spin magnitudes and tilts, we explore the distribution of the effective inspiral spin parameter χ_{eff} and the effective precession spin parameter χ_p . In particular, we wish to measure the mean and variance of each parameter, and so model the joint distribution of χ_{eff} and χ_p as a bivariate Gaussian:

$$\pi(\chi_{\text{eff}}, \chi_p | \mu_{\text{eff}}, \sigma_{\text{eff}}, \mu_p, \sigma_p, \rho) \propto G(\chi_{\text{eff}}, \chi_p | \boldsymbol{\mu}, \boldsymbol{\Sigma}). \quad (\text{B20})$$

The mean of this distribution is $\boldsymbol{\mu} = (\mu_{\text{eff}}, \mu_p)$, and its covariance matrix is

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_{\text{eff}}^2 & \rho\sigma_{\text{eff}}\sigma_p \\ \rho\sigma_{\text{eff}}\sigma_p & \sigma_p^2 \end{pmatrix}. \quad (\text{B21})$$

The population parameters governing this model and their corresponding priors are shown in Table XIII. Equation (B20) is truncated to the physically allowed range of each effective spin parameter, with $\chi_{\text{eff}} \in (-1, 1)$ and $\chi_p \in (0, 1)$. All results in the main text using the GAUSSIAN model are obtained while simultaneously fitting for the BBH mass distribution, assuming the POWER LAW + PEAK model, and the evolving redshift distribution model in Appendix B3 below.

Two variants of this model are additionally discussed in Sect. VII. In the first, Eq. (B20) is modified such that the effective inspiral spin parameter is truncated not on the interval $(-1, 1)$, but on $(\chi_{\text{eff,min}}, 1)$, where $\chi_{\text{eff,min}}$ is another parameter to be inferred by the data. The second variant, inspired by [191] and [192] and defined in Eq. (18), alternatively treats the χ_{eff} distribution as a mixture between a bulk component with a variable mean and width and a narrow zero-spin component centered on $\chi_{\text{eff}} = 0$. In this second variant, we measure only the *marginal* χ_{eff} distribution, implicitly assuming that the remaining spin degree of freedom are distributed uniformly and isotropically. As χ_{eff} is the primary spin measurable, we do not expect this implicit prior to have a strong effect.

3. Redshift Evolution Model

The POWER-LAW redshift evolution model parameterizes the merger rate density per comoving volume and source time as [96]

$$\mathcal{R}(z) = \mathcal{R}_0(1+z)^\kappa, \quad (\text{B22})$$

where \mathcal{R}_0 denotes the merger rate density at $z = 0$. This implies that the redshift distribution is

$$\frac{dN}{dz} = \mathcal{C} \frac{dV_c}{dz} (1+z)^{\kappa-1}, \quad (\text{B23})$$

where dV_c/dz is the differential comoving volume, and \mathcal{C} is related to \mathcal{R}_0 by

$$\mathcal{R}_0 = \mathcal{C} \frac{dV_c}{dz} \left[\int_0^{z_{\max}} \frac{dV_c}{dz} (1+z)^{\kappa-1} \right]^{-1}. \quad (\text{B24})$$

We adopt $z_{\max} = 2.3$ as this is a conservative upper bound on the redshift at which we could detect BBH systems during O3, for both detection thresholds used in this work. We employ a uniform prior on κ centered at $\kappa = 0$. We take a sufficiently wide prior so that the likelihood is entirely within the prior range, $\kappa \in (-6, 6)$.

4. Models with multiple independent components

a. MULTI SOURCE model

The MULTI SOURCE model, introduced in [319], extends the MULTISPIN BBH model introduced in [11] to include additional subpopulations for BNS and NSBH systems. Each subpopulation (two for BBH, one for BNS, and one for NSBH) is assumed to have an independent rate parameter.

The BBH subpopulation is itself a mixture of two subpopulations, i. a power law mass distribution $m_1^{-\alpha} q^\beta$ truncated to a range $[m_{\min, \text{BBH}}, m_{\max, \text{BBH}}]$ which is inferred from the data, and ii. a Gaussian in (m_1, m_2) with independent mean and standard deviation parameters $\mu_{m_1, \text{BBH}}, \mu_{m_2, \text{BBH}}, \sigma_{m_1, \text{BBH}}, \sigma_{m_2, \text{BBH}}$. Both subpopulations, and both binary components within them follow independent DEFAULT spin models, with $\zeta \equiv 1$.

Two more bivariate Gaussians in m_1, m_2 are used to model BNS and NSBH. The BH in NSBH follow a Gaussian mass distribution, with free parameters $\mu_{m, \text{NSBH}}, \sigma_{m, \text{NSBH}}$. As with BBH, these BH follow an independent DEFAULT spin model with $\zeta \equiv 1$. All three types of NS (two in BNS and one in NSBH) follow the same Gaussian mass distribution, with free parameters $\mu_{m, \text{NS}}, \sigma_{m, \text{NS}}, m_{\max, \text{NS}}$ (the minimum mass is assumed to be $1 M_\odot$). Each type of NS follows an independent DEFAULT spin model. To stay within astrophysically plausible spins, the magnitude distributions are scaled down to $\chi_{\max} = 0.05$. Since NS spin tilts are not well measured, we set $\zeta \equiv 0$, assuming they are isotropic, which has the effect of not wasting any samples from parameter estimation.

In addition to any mass cutoffs mentioned above, all BHs component masses are assumed to lie on the range $[2, 100] M_\odot$, with those in NSBHs further restricted to $[2, 50] M_\odot$ due to our limited injections.

Priors for all parameters are given in Table XIV.

Appendix C: Validation studies

We employ several methods to validate our calculations, notably including comparing results from multiple independent analyses; reproducing previous work through O3a [11]; assessing the sensitivity of our results to threshold choices (changing from 1 yr^{-1} to 0.2 yr^{-1} for BBH; or from 0.25 yr^{-1} to 1 yr^{-1} for analyses containing NS); and performing posterior predictive checks as in our analysis of GWTC-2 [11]. Though these specific technical checks will not be described here, some of these checks can be reproduced with the data release associated with this paper.

Below, we describe additional validation studies we have performed to assess whether our results for merger rates are sensitive to the choice of threshold; waveform systematics; or updates to our sensitivity model.

1. NS mass distribution including marginal events

If we loosen the FAR threshold to $< 1 \text{ yr}^{-1}$ so as to include the marginal events GW190917 and GW190426, and repeat the analysis of Sec. VB, the inferred NS mass distribution is virtually unchanged. This can be seen in Fig. 24, which compares the posterior population distri-

Parameter	Description	Prior
$R_{\text{BBH,pl}}$	Local merger rate for the BBH power-law subpopulation.	$U(0, 1000)$
$R_{\text{BBH,g}}$	Local merger rate for the BBH Gaussian subpopulation.	$U(0, 1000)$
R_{BNS}	Local merger rate for the BNS subpopulation.	$U(0, 2000)$
R_{NSBH}	Local merger rate for the NSBH subpopulation.	$U(0, 500)$
α	Primary mass spectral index for the BBH power-law subpopulation.	$U(-4, 12)$
β	Mass ratio spectral index for the BBH power-law subpopulation	$U(-4, 10)$
$m_{\text{min,BBH,pl}}$	Minimum mass of the BBH power-law subpopulation.	$U(2, 10)$
$m_{\text{max,BBH,pl}}$	Maximum mass of the BBH power-law subpopulation.	$U(30, 100)$
$\mu_{m_1,\text{BBH,g}} (\mu_{m_2,\text{BBH,g}})$	Centroid of the primary (secondary) mass distribution for the BBH Gaussian subpopulation	$U(20, 50)$
$\sigma_{m_1,\text{BBH,g}} (\sigma_{m_2,\text{BBH,g}})$	Width of the primary (secondary) mass distribution for the BBH Gaussian subpopulation	$U(0.4, 20)$
$\mu_{m,\text{NSBH}}$	Centroid of the BH mass distribution for NSBH	$U(3, 50)$
$\sigma_{m,\text{NSBH}}$	Width of the BH mass distribution for NSBH	$U(0.4, 20)$
$\mu_{m,\text{NS}}$	Centroid of the NS mass distribution	$U(1, 3)$
$\sigma_{m,\text{NS}}$	Width of the NS mass distribution	$U(0.05, 3)$
$m_{\text{max,NS}}$	Maximum mass of all NS.	$U(2, 3)$
$\mu_{\chi_1,\text{BBH,pl}} (\mu_{\chi_2,\text{BBH,pl}})$	Mean of the Beta distribution of primary (secondary) spin magnitudes for the BBH Gaussian sub-population.	$U(0, 1)$
$\sigma_{\chi_1,\text{BBH,pl}}^2 (\sigma_{\chi_2,\text{BBH,pl}}^2)$	Variance of the Beta distribution of primary (secondary) spin magnitudes for the BBH Gaussian sub-population.	$U(0, 0.25)$
$\sigma_{t_1,\text{BBH,pl}} (\sigma_{t_2,\text{BBH,pl}})$	Width of truncated Gaussian, determining typical primary (secondary) spin misalignment for the BBH Gaussian sub-population.	$U(0, 4)$
$\mu_{\chi_1,\text{BBH,g}} (\mu_{\chi_2,\text{BBH,g}})$	Mean of the Beta distribution of primary (secondary) spin magnitudes for the BBH Gaussian sub-population.	$U(0, 1)$
$\sigma_{\chi_1,\text{BBH,g}}^2 (\sigma_{\chi_2,\text{BBH,g}}^2)$	Variance of the Beta distribution of primary (secondary) spin magnitudes for the BBH Gaussian sub-population.	$U(0, 0.25)$
$\sigma_{t_1,\text{BBH,g}} (\sigma_{t_2,\text{BBH,g}})$	Width of truncated Gaussian, determining typical primary (secondary) spin misalignment for the BBH Gaussian sub-population.	$U(0, 4)$
$\mu_{\chi,\text{NSBH}}$	Mean of the Beta distribution of spin magnitudes for BH in the NSBH sub-population.	$U(0, 1)$
$\sigma_{\chi,\text{NSBH}}^2$	Variance of the Beta distribution of spin magnitudes for BH in the NSBH sub-population.	$U(0, 0.25)$
$\sigma_{t,\text{NSBH}}$	Width of truncated Gaussian, determining typical primary (secondary) spin misalignment for BH in the NSBH sub-population.	$U(0, 4)$
$\mu_{\chi,\text{NSBH}}$	Mean of the Beta distribution of spin magnitudes for NS in the NSBH sub-population.	$U(0, 0.05)$
$\sigma_{\chi,\text{NSBH}}^2$	Variance of the Beta distribution of spin magnitudes for NS in the NSBH sub-population.	$U(0, 0.0125)$
$\mu_{\chi_1,\text{BNS}} (\mu_{\chi_2,\text{BNS}})$	Mean of the Beta distribution of primary (secondary) spin magnitudes in the BNS sub-population.	$U(0, 0.05)$
$\sigma_{\chi_1,\text{BNS}}^2 (\sigma_{\chi_2,\text{BNS}}^2)$	Variance of the Beta distribution of primary (secondary) spin magnitudes in the BNS sub-population.	$U(0, 0.0125)$

TABLE XIV. Summary of MULTI SOURCE model parameters. All rates are in $\text{Gpc}^{-3}\text{yr}^{-1}$, and all masses in M_{\odot} . Rate, mass, and spin hyperparameters are separated by horizontal lines.

butions inferred with and without the marginal events. Traces from the posterior population distribution with respect to the original FAR threshold are also shown. This alternative analysis strongly suggests that substantial uncertainties in the merger rate versus mass dominate our error budget; the handful of observations made to date is not sufficient to overcome the strong impact of our highly

uncertain model priors. Moreover, the masses of the NS secondaries in the marginal events are poorly constrained relative to those in GW170817, GW190425, GW200105 and GW200115, such that the $\text{FAR} < 0.25 \text{yr}^{-1}$ events continue to drive the inference.

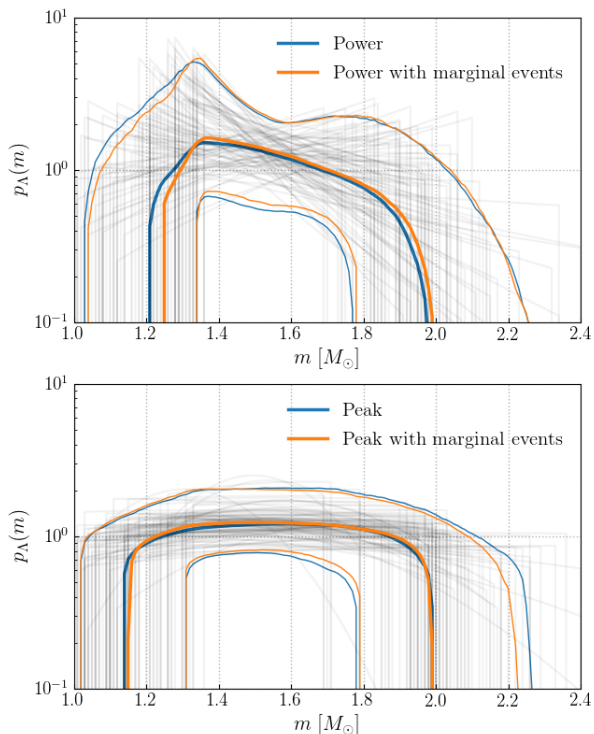


FIG. 24. Inferred neutron star mass distribution with and without the marginal events GW190426 and GW190917. *Top panel:* Median and 90% confidence region of the inferred NS mass distribution for the POWER model, using the event list at a FAR threshold of 0.25 yr^{-1} (blue) and 1 yr^{-1} (orange). Traces from the posterior population distribution with respect to the stricter FAR threshold are plotted in grey. *Bottom panel:* Same as the top panel but for the PEAK model. The inclusion of the marginal events has a negligible impact on the inferred mass distribution.

2. Merger rates including subthreshold triggers

In the main text, our merger rates were calculated after adopting a fixed significance threshold to identify confident events, then fitting population model families to the recovered events’ posteriors. By design, such an approach depends on the threshold. Here we employ an alternative threshold-free method of rate estimation which lacks potential biases from an arbitrary choice of significance threshold [12].

We extend methods from GWTC-2.1 [4], also applied to the discovery of GW200105 and GW200115 [320], to estimate the event rate from the full set of triggers (including subthreshold triggers) from a specific binary merger search: here, GSTLAL [50–52]. In doing so, we allow for population distributions that fit our observations and account for still-considerable uncertainty in the

mass distribution, rather than adopting a fixed population model with fixed model hyper-parameters. Compared to previous publications, the results presented in this section update the BBH merger rates presented in GWTC-2.1 by including O3b events [1]. We also update the NSBH rate quoted in [320] by incorporating all O3 triggers, rather than as previously truncating to the first 9 months of O3.

We use a multi-component mixture model [118] to construct the posterior of astrophysical counts of CBC events by assuming that foreground and background events are independent Poisson processes. We then estimate the space-time volume sensitivity of the pipeline using simulated events which are re-scaled to an astrophysical population model [316]. We then compute the rates as the ratio of the counts to VT . In order to marginalize over population hyper-parameters we compute several VT ’s, each corresponding to a population hyper-parameter sample drawn from the inferred hyper-posterior for the astrophysical population model. Finally, we integrate over the count posterior obtained for each of these samples with an appropriate weight, effectively marginalizing over the population hyper-parameters:

$$\begin{aligned}
 p(R|\vec{x}) &= \int p(R|\vec{\Lambda}, \vec{x}) p(\Lambda|\{d\}) d\vec{\Lambda} \\
 &= \int VT p(N|\vec{x}) p(VT|\vec{\Lambda}) p(\vec{\Lambda}|\{d\}) d(VT) d\vec{\Lambda} \\
 &= \sum_{i,j} VT_{ij} \times p(N_{ij}|\vec{x}), \tag{C1}
 \end{aligned}$$

where \vec{x} is the complete set of triggers (including subthreshold triggers) and $\{d\}$ is the set of data from gravitational-wave detections used in population model inference, as in III B. The astrophysical count posterior is given by $p(N|\vec{x})$, where $N = R \times VT$; we evaluate by sampling via $N_{ij} = R \times VT_{ij}$ where VT_{ij} is the i ’th VT sample drawn from $p(VT|\vec{\Lambda}_j)$ for the j ’th hyper-parameter sample $\vec{\Lambda}_j$ drawn from the inferred hyper-posterior $p(\vec{\Lambda}|\{d\})$.¹ Following [118], we take the distribution $p(VT|\vec{\Lambda}_j)$ to be

$$p(VT|\vec{\Lambda}_j) = \frac{1}{VT \sqrt{2\pi\sigma^2}} \exp \left[-\frac{[\ln VT - \ln \langle VT \rangle(\vec{\Lambda}_j)]^2}{2\sigma^2} \right], \tag{C2}$$

where $\langle VT(\vec{\Lambda}_j) \rangle$ is calculated by re-weighting simulated sources to an astrophysical population with hyper-parameter $\vec{\Lambda}_j$, and σ is the quadrature sum of a calibration error of 10% [91] and Monte-Carlo uncertainty.

Using hyper-parameter samples from the posterior inferred using the PP model with data through the end

¹ The extra VT factor in Eq. (C1) arises from the Jacobian dN/dR .

of O3, as in Section VI and imposing a Jeffreys prior $\propto N^{-1/2}$ on the astrophysical counts, we compute a BBH merger rate of 24.58–67.44 $\text{Gpc}^{-3}\text{yr}^{-1}$. A similar calculation for the BGP model, again with a Jeffreys prior $\propto N^{-1/2}$ imposed on the astrophysical counts, yields an NSBH merger rate of 15.16–199.06 $\text{Gpc}^{-3}\text{yr}^{-1}$, which is consistent with 15.0–180.0 $\text{Gpc}^{-3}\text{yr}^{-1}$, the joint inference for the NSBH merger rate presented in the main text.

3. Effect of Waveform Systematics on Population

All O3b BBH events analyzed in this paper have source properties inferred using two different waveform models: SEOBNRv4PHM [87] and IMRPhenomXPHM [88], both of which include effects of higher-order multipole moments and spin precession. The posterior distribution for each event is then checked for consistency between waveform models before use in our analyses.

The event GW200129 is the highest SNR event exhibiting notable inconsistencies between the source properties inferred with the two waveform models. The event analysis using IMRPhenomXPHM infers much more support for unequal masses and precessing spins relative to the analysis using SEOBNRv4PHM. See [1] for an extended discussion of these systematic differences.

To test if the inferred BBH spin population depends on the waveform model chosen for this event, we repeat our O3 population inference using the PP model for three different choices of waveform model for GW200129: IMRPhenomXPHM, SEOBNRv4PHM, and a mix of the two. As shown in Fig. 25, the inferred spin population is not significantly affected by changes in the waveform model for this event.

4. Impact of Sensitivity on Redshift Evolution Inference

As noted in Sec. II, one change in the sensitivity estimation procedure between this work and our previous study of GWTC-2 [11] is the use of injections that account for the effect of precession and as well as updates to our detection pipelines as detailed in [1]. Since precession was not included in the injections used in [11], the full spin distribution could not be reweighted to calculate the sensitivity via Equation A2, and thus, for the purposes of sensitivity estimation, an approximation was made that $S_{x,y} \in (-0.5, 0.5)$. Since we now use precessing injections, we do the reweighting procedure including the full spin distribution as a function of Λ . To test if this difference in our sensitivity estimation procedure is responsible for the change in the inferred redshift evolution, we repeat the population analysis reported in Sec. VI, using our updated sensitivity model, but only including events analyzed in the GWTC-2 populations study [11]. From this analysis, we infer $\kappa > 0$ at 97.6% credibility, as opposed to the 85% credibility reported in




FIG. 25. Inferred differential merger rate as a function of the cosine of the tilt angle (t_i), where i indexes the body of the binary. We demonstrate that the differences in the posterior distribution for GW200129’s spin parameters have a minimal effect on the inferred spin tilt population. The population is inferred using posterior distributions for GW200129 using the IMRPhenomXPHM waveform model (orange), SEOBNRv4PHM model (blue), and a mixture of both (green). Dashed lines are 90% credible intervals.

[11], indicating a much stronger preference for a merger rate increasing with redshift. We conclude that the differences between our current results for the evolution of the BBH merger rate and those reported in [11] are due to improvements to our sensitivity model rather than the presence of the additional events in GWTC-3.

In Fig. 26 we compare the redshift dependence of our current sensitivity model to that of the sensitivity model used in [11]. To make this comparison, we reweight the injections used in [11] to the same spin distribution assumed in that study, and assuming a fiducial PP and POWERLAW model for the mass and redshift distributions, respectively. We reweight the current injections to this same mass and redshift distribution, but reweight them to the median inferred spin distribution in [11], to mimic an astrophysically-realistic population. Both injection sets only cover the observing times of the O3a observing run. Taking the ratio of the corresponding sensitivities, we find our sensitivity has increased for low redshift events and decreased for high redshift events, relative to the sensitivity model used in [11]. We expect to see an increase in sensitivity between [11] and our current analysis due to updates to the detection pipelines. The relative decrease in sensitivity at higher redshifts indicates a bias in the previous sensitivity estimate, implying that the BBH merger rate at high redshift was underestimated in [11]. Accounting for the shift in sensitivity as a function of redshift causes a relative decrease in local BBH merger rate and a relative increase in high-redshift BBH merger rate, leading to a higher inferred value for κ .

One possible explanation for the shift in sensitivity is that the use of precession in the injections for sensitivity




FIG. 26. Comparison of our current BBH merger sensitivity estimate in the O3a observing run (VT_{new}) to that used in [11] (VT_{old}) as a function of redshift, for events with chirp masses between $20M_{\odot}$ and $50M_{\odot}$. Our current sensitivity model differs from what was used in [11] in two important ways: we use updated detection pipelines relative to those used in [11] and we use injections which include spin precession. Note the relative increase (decrease) in sensitivity at low (high) redshift. Computed by reweighting injections to a fiducial population for each of the two injection sets.

estimation caused a non-trivial change in the inferred sensitive hypervolume, given that we do observe precession in the BBH population. Our current detection pipelines use template banks that include only aligned-spin components; this can result in up to tens of percent reduced sensitivity to a population of BBHs with spin precession, depending on the degree of precession possible [321–323]. The farthest precessing sources, which, due to their distances, correspond to FARs closest to the detection threshold, are therefore the most susceptible to dropping below the detection threshold with our current pipelines, causing us to see a decrease in sensitivity to a population of BBHs with precession relative to a strictly non-precessing population.

Additionally, both the use of population-informed reweighting of the spin distribution to calculate sensitivity to a population and the incorporation of additional detection pipelines may have contributed to a more accurate estimate of our sensitivity across parameter space.

Appendix D: Additional studies of the binary black hole distribution

1. Analyses from GWTC-2

We report updated Bayes factor comparisons for these various models in Table XV, showing that the BROKEN POWERLAW + PEAK model is slightly preferred over our fiducial POWER LAW + PEAK model. We highlight the

Model	GWTC-2	GWTC-3
	$\log_{10} \mathcal{B}$	$\log_{10} \mathcal{B}$
POWER LAW + PEAK	0.0	0.0
BROKEN POWERLAW + PEAK	-0.11	0.36
MULTI PEAK	-0.3	-0.37
BROKEN POWERLAW	-0.92	-2.1

TABLE XV. Bayes factors for each of the previously used phenomenological mass models relative to the model with highest marginal likelihood, POWER LAW + PEAK. The previous results from GWTC-2 are shown in the second column with the updated catalog results in the third column.

key differences between the model priors for GWTC-2 compared to GWTC-3: the prior on β_q was changed from $U(-4, 12)$ to $U(-2, 7)$, and the main population results now include an evolving redshift model where the prior on κ is changed from 0 to $U(-10, 10)$.

In addition to these analyses, we used a variation of the MULTYPEAK model to study the feature in the mass distribution at $\sim 10M_{\odot}$. In GWTC-2 the prior on the mean of the peaks were $U(20, 50)$ and $U(50, 100)$ for the lower and upper mass peaks respectively. We modified these priors to be $U(5, 20)$ and $U(20, 100)$. This updated MULTYPEAK model is most preferred model with a Bayes factor of 1.1 compared to the POWER LAW + PEAK model. This further supports our findings of the peak-like feature at $\sim 10M_{\odot}$ in the mass distribution.

2. Comprehensive BBH merger rates

In Table IV, we evaluate BBH merger rates over targeted mass subsets of the whole BBH space, using models specifically targeted to reproduce new features of the binary black hole mass distribution. For broader context, Table XVI also provides the corresponding merger rates in these intervals from all the models presented in this work.

Appendix E: Population-weighted posteriors

With an increasing number of events, we can use the distribution of the population of compact binaries to inform our priors for parameter estimation. By reweighting the initial analysis of compact binaries with the population distribution we can obtain posterior distribution for the events in GWTC-3 with population-informed priors. Using our population analysis with models POWER LAW + PEAK and FLEXIBLE MIXTURES we provide population-weighted posteriors (Fig. 27) for m_1 , q and χ_{eff} for the BBHs population (69 events).

Some of our analyses will show apparent changes in the inferences about the mass ratio. These seemingly-substantive changes reflect the relatively weak con-

	$m_1 \in [5, 20]M_\odot$	$m_1 \in [20, 50]M_\odot$	$m_1 \in [50, 100]M_\odot$	All BBH
	$m_2 \in [5, 20]M_\odot$	$m_2 \in [5, 50]M_\odot$	$m_2 \in [5, 100]M_\odot$	
PDB (pair)	$17.0^{+10.0}_{-6.0}$	$6.8^{+2.3}_{-1.7}$	$0.68^{+0.42}_{-0.3}$	$25.0^{+10.0}_{-7.0}$
PDB (ind)	$9.3^{+5.7}_{-3.6}$	$11.0^{+3.0}_{-2.0}$	$1.6^{+1.0}_{-0.7}$	$22.0^{+9.0}_{-6.0}$
MS	$34.0^{+86.0}_{-20.0}$	$6.1^{+3.1}_{-2.1}$	$0.93^{+0.87}_{-0.64}$	$42.0^{+88.0}_{-20.0}$
BGP	$20.0^{+10.0}_{-8.0}$	$6.4^{+3.0}_{-2.1}$	$0.74^{+1.2}_{-0.46}$	$33.0^{+16.0}_{-10.0}$
PS	$27^{+12}_{-9.4}$	$3.6^{+1.5}_{-1.1}$	$0.2^{+0.18}_{-0.1}$	$32^{+14}_{-9.6}$
FM	$21.1^{+10.7}_{-8.3}$	$4.1^{+2.0}_{-1.4}$	$0.2^{+0.3}_{-0.1}$	$26.0^{+11.5}_{-8.7}$
PP	$23.4^{+12.9}_{-8.6}$	$4.5^{+1.8}_{-1.3}$	$0.2^{+0.1}_{-0.1}$	$28.1^{+14.8}_{-10.0}$
MERGED	12.8 – 40	0.098 – 6.3	2.5 – 0.5	17.3 – 45
PP (O3a)	$16.0^{+13.0}_{-7.7}$	$6.8^{+2.7}_{-1.9}$	$0.5^{+0.4}_{-0.3}$	$25.3^{+16.1}_{-9.9}$

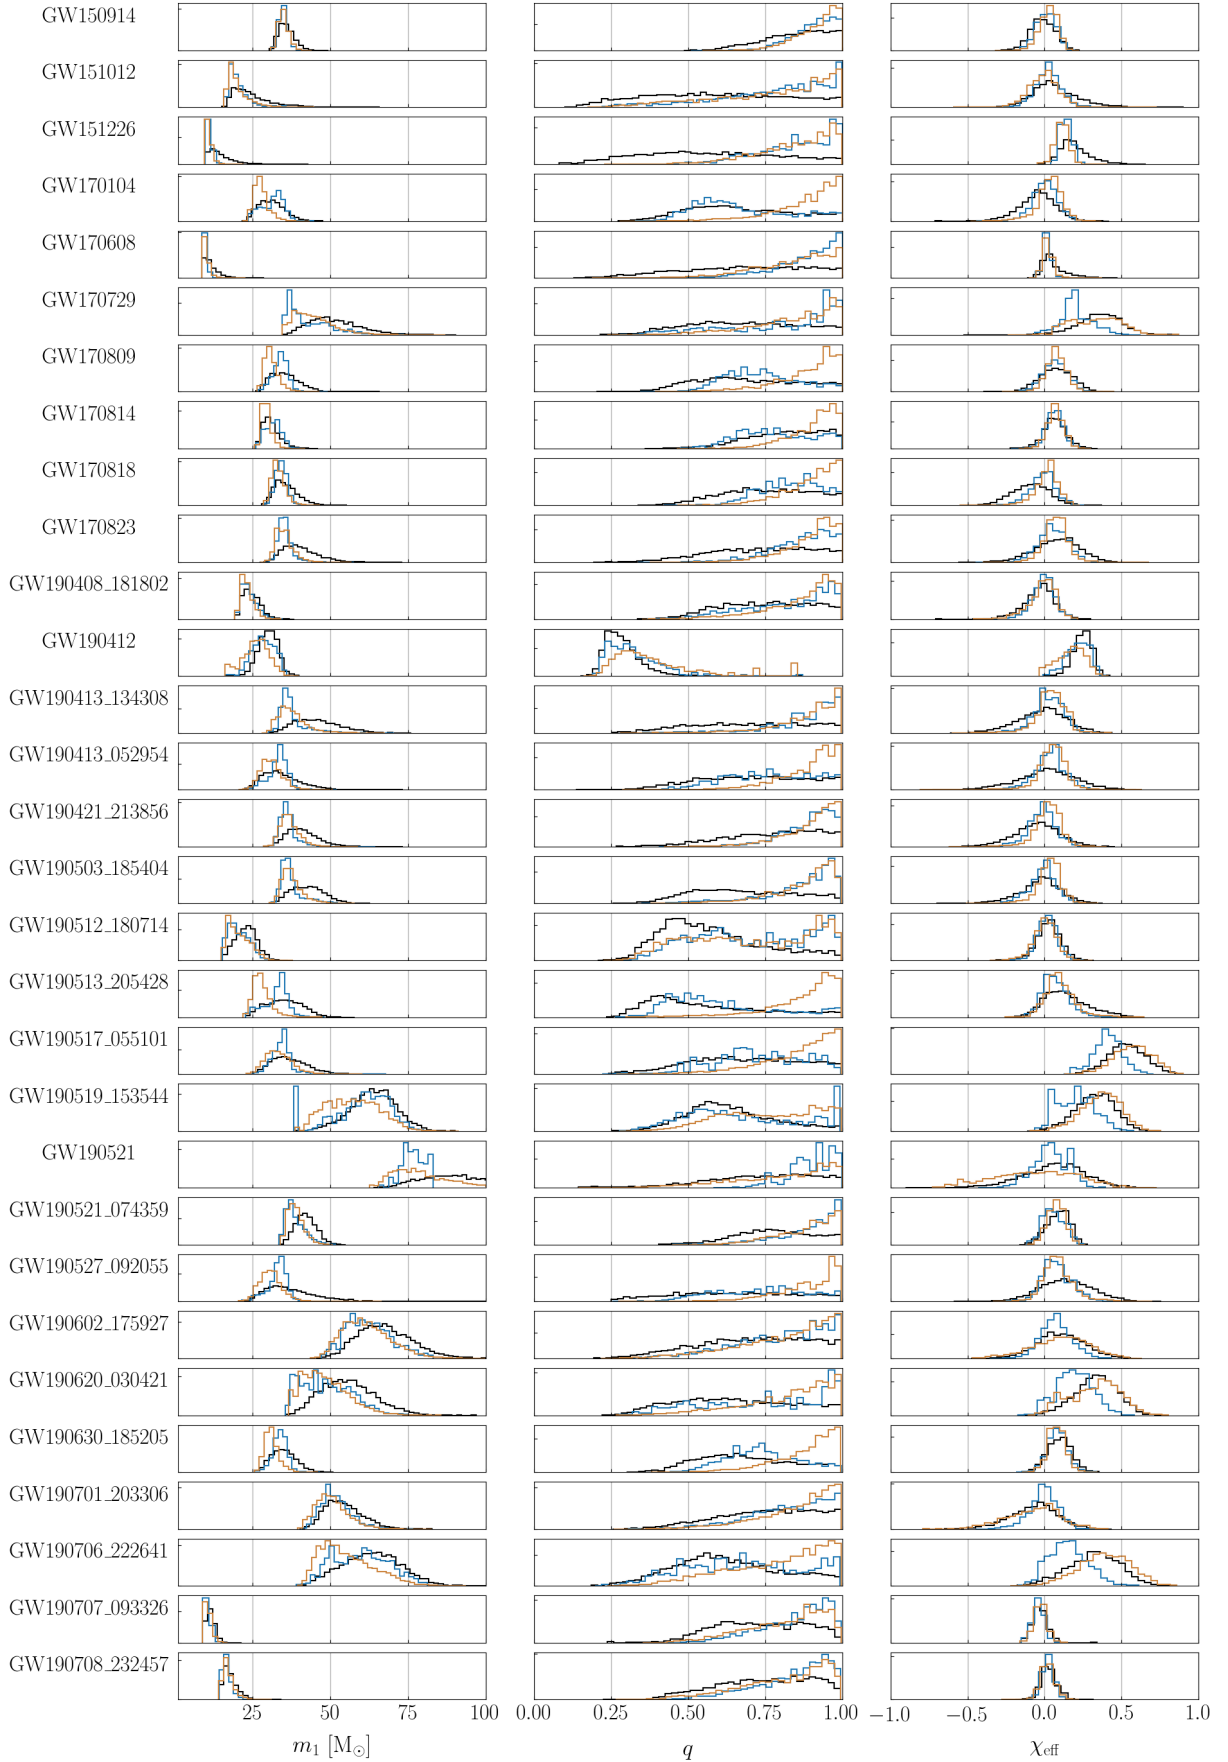
TABLE XVI. Merger rates in $\text{Gpc}^{-3} \text{yr}^{-1}$ for black hole binaries, quoted at the 90% credible interval. Rates are given for three ranges of primary mass, m_1 as well as for the entire population. The PDB, MS, and BGP merger rates are derived assuming the merger rate does not increase with redshift, using a threshold $\text{FAR} < 0.25 \text{yr}^{-1}$ (Sec. IV). For FM, PS, and PP, merger rates are reported at $z = 0.2$, estimated using a threshold $\text{FAR} < 1 \text{yr}^{-1}$ (Sec. VI). The merged rates reported in the MERGED row are the union of the preceding three rows, which all account for distance-dependent merger rate and adopt a consistent threshold. The final row shows merger rates deduced from our analysis of GWTC-2 [11], which assumed a redshift-independent merger rate. Compare to Table IV.

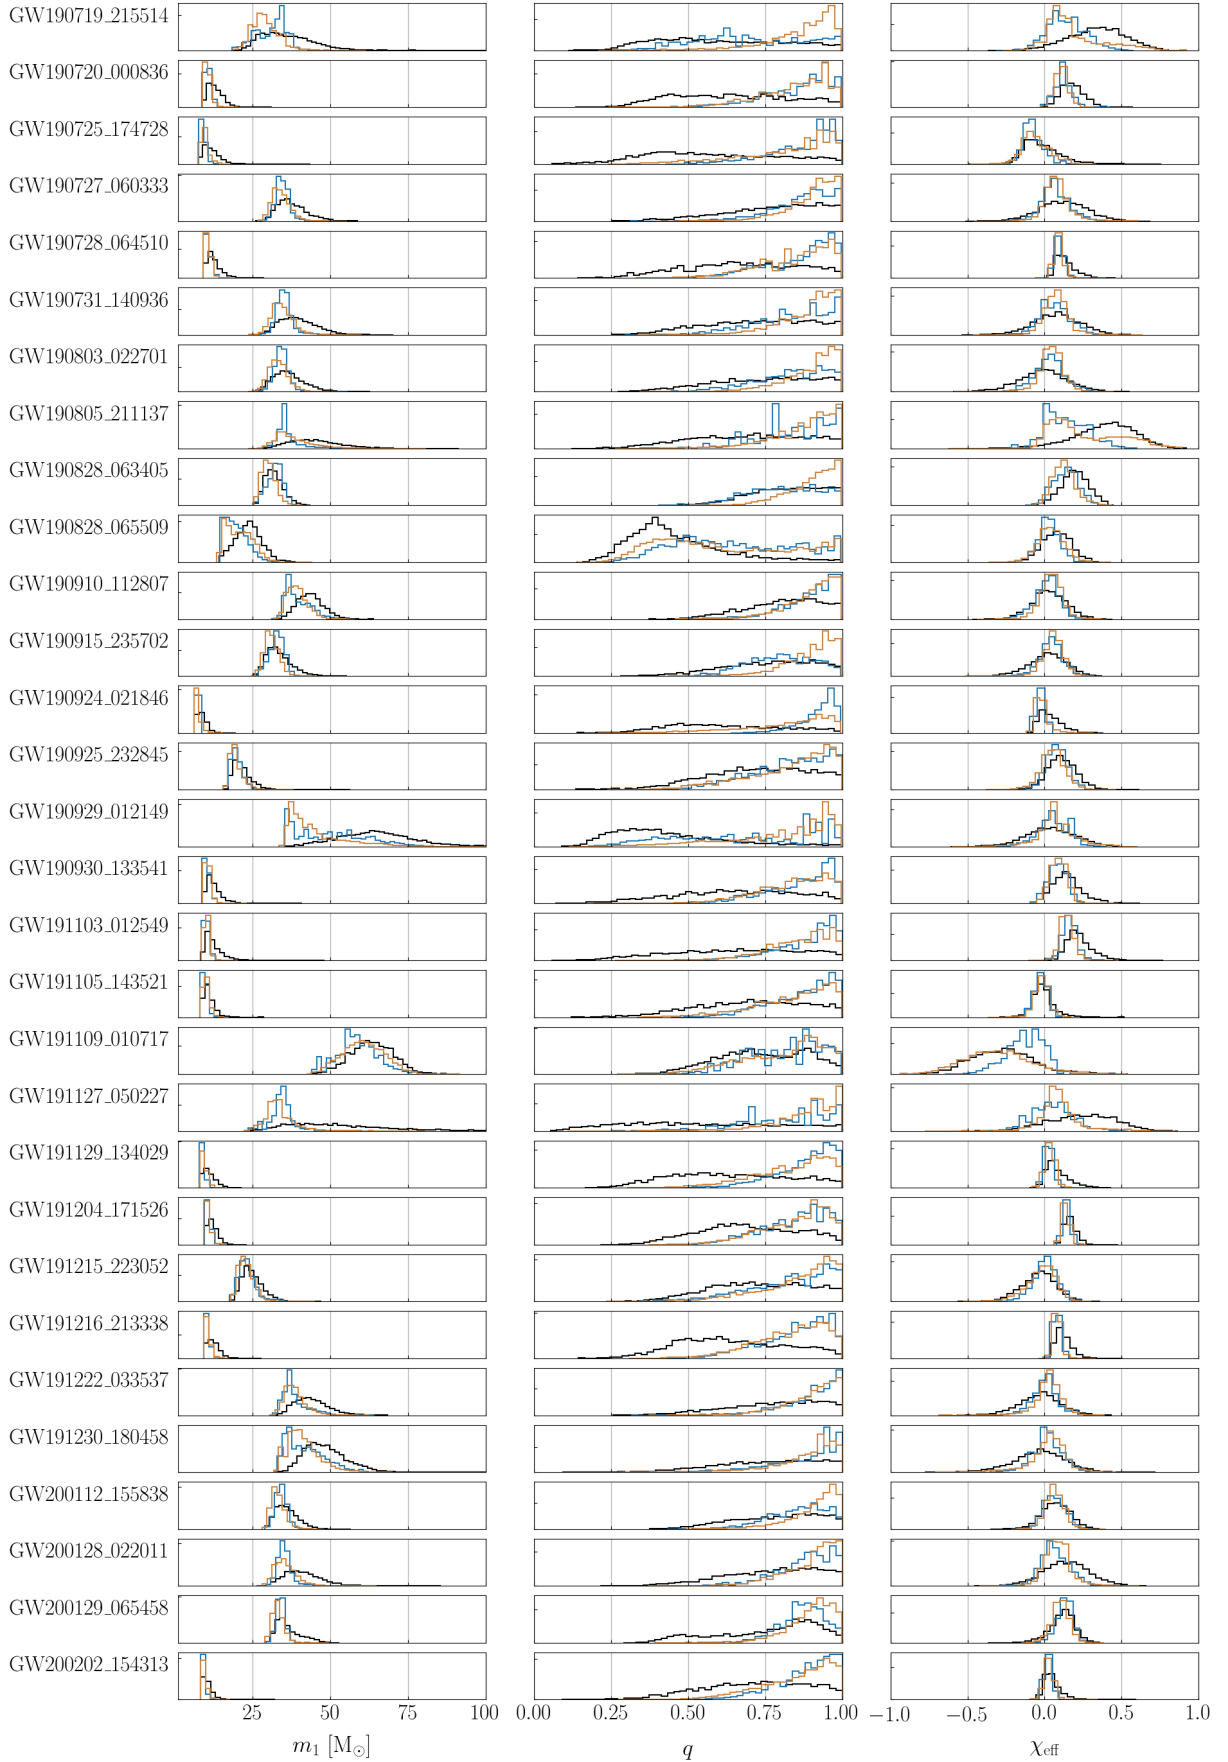
straints provided by the fiducial parameter inferences used as input and shown in black. Specifically, several low-amplitude or low-mass events have extremely weak constraints on mass ratios, with posterior support extending to $q < 0.4$. This extended feature reflects the prior distribution on component masses, conditioned on modest constraints on chirp mass. To be concrete, using the corresponding prior distributions for these events, conditioned on a suitable chirp mass interval, we often find a posterior distribution with comparable support for $q = 1$ (i.e., the Savage-Dickey estimate of the Bayes factor for unequal mass would be nearly unity).

Examining Fig. 27 in light of this caveat, we find that our population models and the fiducial model agree. For these events the population reweight-

ings, as expected, strongly favor symmetric component masses (e.g., GW190503_185404, GW190720_000836, GW191127_050227). For a few binaries, however, the two population reweightings disagree. The most notable example is GW190513_205428, where the FLEXIBLE MIXTURES model pulls the posterior distribution to more symmetric component masses and a lower primary mass. Both population models also pull the majority of the posteriors closer to $\chi_{\text{eff}} \sim 0$. However, given the FLEXIBLE MIXTURES models spins as dependent on chirp mass the events with higher mass and higher spin do not drawn to $\chi_{\text{eff}} \sim 0$ as strongly as the POWER LAW + PEAK model (e.g., GW191109_010717) and in some cases the FLEXIBLE MIXTURES reweighted posterior move to higher χ_{eff} values (e.g., GW190706_222641).

-
- [1] R. Abbott, T. Abbott, S. Abraham, *et al.*, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Half of the Third Observing Run, [Available as LIGO-P2000318 \(2021\)](#).
- [2] B. Abbott, R. Abbott, T. Abbott, *et al.*, GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs, *Phys. Rev. X* **9**, 031040 (2019).
- [3] R. Abbott, T. Abbott, S. Abraham, *et al.*, GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run, *Physical Review X* **11**, 021053 (2021), [arXiv:2010.14527 \[gr-qc\]](#).
- [4] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, and *et al.*, GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run, Available as [arxiv:2108.01045](#), [arXiv:2108.01045 \(2021\)](#), [arXiv:2108.01045 \[gr-qc\]](#).
- [5] J. Aasi *et al.*, Advanced LIGO, *Class. Quant. Grav.* **32**, 074001 (2015).
- [6] F. Acernese *et al.*, Advanced Virgo: a second-generation interferometric gravitational wave detector, *Class. Quant. Grav.* **32**, 024001 (2015).
- [7] R. Essick and P. Landry, Discriminating between Neutron Stars and Black Holes with Imperfect Knowledge of the Maximum Neutron Star Mass, *Astrophys. J.* **904**, 80 (2020), [arXiv:2007.01372 \[astro-ph.HE\]](#).
- [8] P. Landry, R. Essick, and K. Chatziioannou, Non-parametric constraints on neutron star matter with existing and upcoming gravitational wave and pulsar observations, *Phys. Rev. D* **101**, 123007 (2020), [arXiv:2003.04880 \[astro-ph.HE\]](#).
- [9] I. Legred, K. Chatziioannou, R. Essick, S. Han, and P. Landry, Impact of the PSR J 0740 +6620 radius constraint on the properties of high-density matter, *Phys.*








FIG. 27. Posterior distributions (black) for binary black hole events weighted by the population results from POWER LAW + PEAK (blue) and FLEXIBLE MIXTURES (orange)

- Rev. D **104**, 063003 (2021), [arXiv:2106.05313 \[astro-ph.HE\]](#).
- [10] R. Abbott, T. Abbott, S. Abraham, *et al.*, Observation of Gravitational Waves from Two Neutron Star-Black Hole Coalescences, *Astrophys. J. Lett.* **915**, L5 (2021), [arXiv:2106.15163 \[astro-ph.HE\]](#).
- [11] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog, *Astrophys. J. Lett.* **913**, L7 (2021), [arXiv:2010.14533 \[astro-ph.HE\]](#).
- [12] W. M. Farr, J. R. Gair, I. Mandel, and C. Cutler, Counting and confusion: Bayesian rate estimation with multiple populations, *Phys. Rev. D* **91**, 023005 (2015).
- [13] R. Abbott, T. Abbott, S. Abraham, *et al.*, Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* **116**, 061102 (2016), [arXiv:1602.03837 \[gr-qc\]](#).
- [14] B. Abbott, R. Abbott, T. Abbott, *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Binary black hole mergers in the first advanced ligo observing run, *Phys. Rev. X* **6**, 041015 (2016).
- [15] B. P. Abbott *et al.* (LIGO Scientific, Virgo), GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence, *Phys. Rev. Lett.* **116**, 241103 (2016), [arXiv:1606.04855 \[gr-qc\]](#).
- [16] B. Abbott, R. Abbott, T. Abbott, *et al.*, GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2, *Phys. Rev. Lett.* **118**, 221101 (2017), [arXiv:1706.01812 \[gr-qc\]](#).
- [17] B. Abbott, R. Abbott, T. Abbott, *et al.*, GW170608: Observation of a 19 Solar-mass Binary Black Hole Coalescence, *Astrophys. J. Lett.* **851**, L35 (2017), [arXiv:1711.05578 \[astro-ph.HE\]](#).
- [18] B. Abbott, R. Abbott, T. Abbott, *et al.*, GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence, *Phys. Rev. Lett.* **119**, 141101 (2017), [arXiv:1709.09660 \[gr-qc\]](#).
- [19] B. Abbott, R. Abbott, T. Abbott, *et al.* (LIGO Scientific, Virgo), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Phys. Rev. Lett.* **119**, 161101 (2017).
- [20] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), GW190412: Observation of a binary-black-hole coalescence with asymmetric masses, *Phys. Rev. D* **102**, 043015 (2020).
- [21] B. Abbott, R. Abbott, T. Abbott, *et al.*, GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4 M_{\odot}$, *Astrophys. J. Lett.* **892**, L3 (2020).
- [22] R. Abbott *et al.* (LIGO Scientific, Virgo), GW190521: A Binary Black Hole Merger with a Total Mass of $150 M_{\odot}$, *Phys. Rev. Lett.* **125**, 101102 (2020), [arXiv:2009.01075 \[gr-qc\]](#).
- [23] R. Abbott, T. Abbott, S. Abraham, *et al.*, GW190814: Gravitational waves from the coalescence of a 23 solar mass black hole with a 2.6 solar mass compact object, *Astrophys. J. Lett.* **896**, L44 (2020).
- [24] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kasta, M. Schäfer, R. Dhurkunde, and M. Cabero, 3-OGC: Catalog of gravitational waves from compact-binary mergers, Available as [arxiv:2105.09151](#) (2021), [arXiv:2105.09151 \[astro-ph.HE\]](#).
- [25] B. Abbott, R. Abbott, T. Abbott, *et al.*, Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo, *Astrophys. J. Lett.* **882**, L24 (2019).
- [26] P. Ajith *et al.*, Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins, *Phys. Rev. Lett.* **106**, 241101 (2011), [arXiv:0909.2867 \[gr-qc\]](#).
- [27] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and *et al.*, Planck 2015 results. XIII. Cosmological parameters, *Astron. Astrophys.* **594**, A13 (2016), [arXiv:1502.01589 \[astro-ph.CO\]](#).
- [28] C. D. Bailyn, R. K. Jain, P. Coppi, and J. A. Orosz,

- The Mass distribution of stellar black holes, *Astrophys. J.* **499**, 367 (1998), [arXiv:astro-ph/9708032](#).
- [29] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock, The Black Hole Mass Distribution in the Galaxy, *Astrophys. J.* **725**, 1918 (2010), [arXiv:1006.2834 \[astro-ph.GA\]](#).
- [30] W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, I. Mandel, and V. Kalogera, The Mass Distribution of Stellar-Mass Black Holes, *Astrophys. J.* **741**, 103 (2011), [arXiv:1011.1459 \[astro-ph.GA\]](#).
- [31] L. Kreidberg, C. D. Bailyn, W. M. Farr, and V. Kalogera, Mass Measurements of Black Holes in X-Ray Transients: Is There a Mass Gap?, *Astrophys. J.* **757**, 36 (2012), [arXiv:1205.1805 \[astro-ph.HE\]](#).
- [32] C. L. Fryer, K. Belczynski, G. Wiktorowicz, M. Dominik, V. Kalogera, and D. E. Holz, Compact Remnant Mass Function: Dependence on the Explosion Mechanism and Metallicity, *Astrophys. J.* **749**, 91 (2012), [arXiv:1110.1726 \[astro-ph.SR\]](#).
- [33] I. Mandel and B. Müller, Simple recipes for compact remnant masses and natal kicks, *Mon. Not. Roy. Astron. Soc.* **499**, 3214 (2020), [arXiv:2006.08360 \[astro-ph.HE\]](#).
- [34] M. Zevin, M. Spera, C. P. L. Berry, and V. Kalogera, Exploring the Lower Mass Gap and Unequal Mass Regime in Compact Binary Evolution, *Astrophys. J. Lett.* **899**, L1 (2020), [arXiv:2006.14573 \[astro-ph.HE\]](#).
- [35] T. Liu, Y.-F. Wei, L. Xue, and M.-Y. Sun, Final Compact Remnants in Core-collapse Supernovae from 20 to 40 M_{\odot} : The Lower Mass Gap, *Astrophys. J.* **908**, 106 (2021), [arXiv:2011.14361 \[astro-ph.HE\]](#).
- [36] R. A. Patton, T. Sukhbold, and J. J. Eldridge, Comparing Compact Object Distributions from Mass- and Pre-supernova Core Structure-based Prescriptions, Available as [arxiv:2106.05987](#) (2021), [arXiv:2106.05978 \[astro-ph.HE\]](#).
- [37] T. A. Thompson, C. S. Kochanek, K. Z. Stanek, C. Badenes, R. S. Post, T. Jayasinghe, D. W. Latham, A. Bieryla, G. A. Esquerdo, P. Berlind, M. L. Calkins, J. Tayar, L. Lindgren, J. A. Johnson, T. W. S. Holoién, K. Auchettl, and K. Covey, A noninteracting low-mass black hole-giant star binary system, *Science* **366**, 637 (2019), [arXiv:1806.02751 \[astro-ph.HE\]](#).
- [38] T. Jayasinghe, K. Z. Stanek, T. A. Thompson, C. S. Kochanek, D. M. Rowan, P. J. Vallely, K. G. Strassmeier, M. Weber, J. T. Hinkle, F. J. Hambach, D. V. Martin, J. L. Prieto, T. Pessi, D. Huber, K. Auchettl, L. A. Lopez, I. Ilyin, C. Badenes, A. W. Howard, H. Isaacson, and S. J. Murphy, A unicorn in monoceros: the 3 M_{\odot} dark companion to the bright, nearby red giant V723 Mon is a non-interacting, mass-gap black hole candidate, *Mon. Not. R. Ast. Soc.* **504**, 2577 (2021), [arXiv:2101.02212 \[astro-ph.SR\]](#).
- [39] B. Kiziltan, A. Kottas, M. De Yoreo, and S. E. Thorsett, The Neutron Star Mass Distribution, [arXiv e-prints](#), [arXiv:1309.6635](#) (2013), [arXiv:1309.6635 \[astro-ph.SR\]](#).
- [40] F. Özel and P. Freire, Masses, Radii, and the Equation of State of Neutron Stars, *Ann. Rev. Astron. Astrophys.* **54**, 401 (2016), [arXiv:1603.02698 \[astro-ph.HE\]](#).
- [41] N. Farrow, X.-J. Zhu, and E. Thrane, The mass distribution of galactic double neutron stars, *Astrophys. J.* **876**, 18 (2019).
- [42] B. Wang and X. Wang, Bandwidth selection for weighted kernel density estimation (2011), [arXiv:0709.1616 \[stat.ME\]](#).
- [43] T. Menne, awkde code, <https://github.com/mennthor/awkde> (2020), [Online; accessed Dec-2020].
- [44] J. Sadiq, T. Dent, and D. Wysocki, Flexible and Fast Estimation of Binary Merger Population Distributions with Adaptive KDE, Available as LIGO-T2100447 (2021), <https://dcc.ligo.org/LIGO-T2100447/public>.
- [45] V. Tiwari and S. Fairhurst, The Emergence of Structure in the Binary Black Hole Mass Distribution, *Astrophys. J. Lett.* **913**, L19 (2021), [arXiv:2011.04502 \[astro-ph.HE\]](#).
- [46] B. Edelman, Z. Doctor, J. Godfrey, and B. Farr, Ain't No Mountain High Enough: Semi-Parametric Modeling of LIGO-Virgo Binary Black Hole Mass Distribution, Available as [arxiv:2109.06137](#), [arXiv:2109.06137 \(2021\)](#), [arXiv:2109.06137 \[astro-ph.HE\]](#).
- [47] Y.-J. Li, Y.-Z. Wang, M.-Z. Han, S.-P. Tang, Q. Yuan, Y.-Z. Fan, and D.-M. Wei, A Flexible Gaussian Process Reconstruction Method and the Mass Function of the Coalescing Binary Black Hole Systems, *Astrophys. J.* **917**, 33 (2021), [arXiv:2104.02969 \[astro-ph.HE\]](#).
- [48] D. Veske, I. Bartos, Z. Márka, and S. Márka, Characterizing the observation bias in gravitational-wave detections and finding structured population properties, Available as [arxiv:2105.13983](#), [arXiv:2105.13983 \(2021\)](#), [arXiv:2105.13983 \[gr-qc\]](#).
- [49] V. Tiwari, VAMANA: modeling binary black hole population with minimal assumptions, *Class. Quant. Grav.* **38**, 155007 (2021), [arXiv:2006.15047 \[astro-ph.HE\]](#).
- [50] S. Sachdev *et al.*, The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO's Second and Advanced Virgo's First Observing Runs, [arXiv eprints](#) (2019), [arXiv:1901.08580 \[gr-qc\]](#).
- [51] C. Hanna *et al.*, Fast evaluation of multi-detector consistency for real-time gravitational wave searches, *Phys. Rev. D* **101**, 022003 (2020).
- [52] C. Messick *et al.*, Analysis Framework for the Prompt Discovery of Compact Binary Mergers in Gravitational-wave Data, *Phys. Rev. D* **95**, 042001 (2017).
- [53] A. Nitz *et al.*, (2019), [gwastro/pycbc: Pycbc release v1.15.0](#).
- [54] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, FIND-CHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries, *Phys. Rev. D* **85**, "122006" (2012).
- [55] B. Allen, χ^2 time-frequency discriminator for gravitational wave detection, *Phys. Rev. D* **71**, 062001 (2005).
- [56] T. D. Canton *et al.*, Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors, *Phys. Rev. D* **90**, 082004 (2014).
- [57] S. A. Usman *et al.*, The PyCBC search for gravitational waves from compact binary coalescence, *Class. Quant. Grav.* **33**, 215004 (2016).
- [58] A. H. Nitz, T. Dent, T. D. Canton, S. Fairhurst, and D. A. Brown, Detecting binary compact-object mergers with gravitational waves: Understanding and Improving the sensitivity of the PyCBC search, *Astrophys. J.* **849**, 118 (2017).
- [59] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang, Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave

- detector era, *Class. Quant. Grav.* **33**, 175012 (2016).
- [60] S. Klimenko and G. Mitselmakher, A wavelet method for detection of gravitational wave bursts, *8th Gravitational Wave Data Analysis Workshop (GWDAAW 2003) Milwaukee, Wisconsin, December 17-20, 2003*, *Class. Quant. Grav.* **21**, S1819 (2004).
- [61] S. Klimenko *et al.*, Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors, *Phys. Rev. D* **93**, 042004 (2016).
- [62] S. M. Gaebel, J. Veitch, T. Dent, and W. M. Farr, Digging the population of compact binary mergers out of the noise, *Mon. Not. R. Ast. Soc.* **484**, 4008 (2019), [arXiv:1809.03815 \[astro-ph.IM\]](#).
- [63] S. Galaudage, C. Talbot, and E. Thrane, Gravitational-wave inference in the catalog era: Evolving priors and marginal events, *Phys. Rev. D* **102**, 083026 (2020), [arXiv:1912.09708 \[astro-ph.HE\]](#).
- [64] J. Roulet, T. Venumadhav, B. Zackay, L. Dai, and M. Zaldarriaga, Binary black hole mergers from LIGO/Virgo O1 and O2: Population inference combining confident and marginal events, *Phys. Rev. D* **102**, 123022 (2020), [arXiv:2008.07014 \[astro-ph.HE\]](#).
- [65] B. Zackay, L. Dai, T. Venumadhav, J. Roulet, and M. Zaldarriaga, Detecting Gravitational Waves With Disparate Detector Responses: Two New Binary Black Hole Mergers, *Phys. Rev. D* **101**, 083030 (2020).
- [66] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New Binary Black Hole Mergers in the Second Observing Run of Advanced LIGO and Advanced Virgo, *Phys. Rev. D* **101**, 083030 (2020).
- [67] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New search pipeline for compact binary mergers: Results for binary black holes in the first observing run of Advanced LIGO, *Phys. Rev. D* **100**, 023011 (2019).
- [68] B. Zackay, T. Venumadhav, L. Dai, J. Roulet, and M. Zaldarriaga, Highly spinning and aligned binary black hole merger in the Advanced LIGO first observing run, *Phys. Rev. D* **100**, 023007 (2019).
- [69] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, 1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO data, *Astrophys. J.* **872**, 195 (2019).
- [70] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, 2-OGC: Open Gravitational-wave Catalog of binary mergers from analysis of public Advanced LIGO and Virgo data, *Astrophys. J.* **891**, 123 (2020).
- [71] R. Magee *et al.*, Sub-threshold Binary Neutron Star Search in Advanced LIGO's First Observing Run, *Astrophys. J.* **878**, L17 (2019).
- [72] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, *SoftwareX* **13**, 100658 (2021), [arXiv:1912.11716 \[gr-qc\]](#).
- [73] A. Trovato *et al.* (Ligo Scientific, The Virgo), GWOSC: Gravitational Wave Open Science Center, *Proceedings, The New Era of Multi-Messenger Astrophysics (ASTERICS 2019): Groningen, Netherlands, March 25-29, 2019*, PoS *Asterics2019*, 082 (2020).
- [74] R. Abbott, T. Abbott, S. Abraham, *et al.*, Search for intermediate mass black hole binaries in the third observing run of Advanced LIGO and Advanced Virgo, Available as [arxiv:2105.15120](#) (2021).
- [75] G. Ashton and E. Thrane, The astrophysical odds of GW151216, *Mon. Not. Roy. Astron. Soc.* **498**, 1905 (2020), [arXiv:2006.05039 \[astro-ph.HE\]](#).
- [76] G. Pratten and A. Vecchio, Assessing gravitational-wave binary black hole candidates with Bayesian odds, *arXiv eprints* (2020), [arXiv:2008.00509 \[gr-qc\]](#).
- [77] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin, W. Del Pozzo, F. Feroz, J. Gair, C. J. Haster, V. Kalogera, T. Littenberg, I. Mandel, R. O'Shaughnessy, M. Pitkin, C. Rodriguez, C. Röver, T. Sidery, R. Smith, M. Van Der Sluys, A. Vecchio, W. Vousden, and L. Wade, Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library, *Phys. Rev. D* **91**, 042003 (2015).
- [78] J. Lange, R. O'Shaughnessy, and M. Rizzo, Rapid and accurate parameter inference for coalescing, precessing compact binaries, (2018), [arxiv/1805.10457](#).
- [79] D. Wysocki, R. O'Shaughnessy, J. Lange, and Y.-L. L. Fang, Accelerating parameter inference with graphics processing units, *Phys. Rev. D* **99** (2019).
- [80] G. Ashton, M. Hübner, P. D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu, Q. Chu, A. Divakarla, P. J. Easter, B. Goncharov, F. Hernandez Vivanco, J. Harms, M. E. Lower, G. D. Meadors, D. Melchor, E. Payne, M. D. Pitkin, J. Powell, N. Sarin, R. J. E. Smith, and E. Thrane, BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy, *Astrophys. J. Supp.* **241**, 27 (2019).
- [81] I. M. Romero-Shaw, C. Talbot, S. Biscoveanu, V. D'Emilio, G. Ashton, C. P. L. Berry, S. Coughlin, S. Galaudage, C. Hoy, M. Hübner, K. S. Phukon, M. Pitkin, M. Rizzo, N. Sarin, R. Smith, S. Stevenson, A. Vajpeyi, M. Arène, K. Athar, S. Banagiri, N. Bose, M. Carney, K. Chatziioannou, J. A. Clark, M. Colleoni, R. Cotesta, B. Edelman, H. Estellés, C. García-Quirós, A. Ghosh, R. Green, C. J. Haster, S. Husa, D. Keitel, A. X. Kim, F. Hernandez-Vivanco, I. Magaña Hernandez, C. Karathanasis, P. D. Lasky, N. De Lillo, M. E. Lower, D. Macleod, M. Mateu-Lucena, A. Miller, M. Millhouse, S. Morisaki, S. H. Oh, S. Ossokine, E. Payne, J. Powell, G. Pratten, M. Pürrer, A. Ramos-Buades, V. Raymond, E. Thrane, J. Veitch, D. Williams, M. J. Williams, and L. Xiao, Bayesian inference for compact binary coalescences with BILBY: validation and application to the first LIGO-Virgo gravitational-wave transient catalogue, *Mon. Not. R. Ast. Soc.* **499**, 3295 (2020), [arXiv:2006.00714 \[astro-ph.IM\]](#).
- [82] B. Abbott, R. Abbott, T. Abbott, *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Properties of the Binary Black Hole Merger GW150914, *Phys. Rev. Lett.* **116**, 241102 (2016).
- [83] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, N. W. Taylor, and A. Zenginoglu, Effective-one-body model for black-hole binaries with generic mass ratios and spins, *Phys. Rev. D* **89**, 061502 (2014), [arXiv:1311.2544 \[gr-qc\]](#).
- [84] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel,

- and B. Szilágyi, Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism, *Phys. Rev. D* **89**, 084006 (2014), [arXiv:1307.6232 \[gr-qc\]](#).
- [85] M. Hannam *et al.*, Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [86] A. Bohé *et al.*, Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors, *Phys. Rev. D* **95**, 044028 (2017).
- [87] S. Ossokine, A. Buonanno, S. Marsat, R. Cotesta, S. Babak, T. Dietrich, R. Haas, I. Hinder, H. P. Pfeiffer, M. Pürrer, C. J. Woodford, M. Boyle, L. E. Kidder, M. A. Scheel, and B. Szilágyi, Multipolar effective-one-body waveforms for precessing binary black holes: Construction and validation, *Phys. Rev. D* **102**, 044055 (2020), [arXiv:2004.09442 \[gr-qc\]](#).
- [88] G. Pratten, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, M. Mateu-Lucena, R. Jaume, M. Haney, D. Keitel, J. E. Thompson, and S. Husa, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, *Phys. Rev. D* **103**, 104056 (2021), [arXiv:2004.06503 \[gr-qc\]](#).
- [89] L. Sun, E. Goetz, J. S. Kissel, J. Betzwieser, S. Karki, D. Bhattacharjee, P. B. Covas, L. E. H. Datrier, S. Kandhasamy, Y. K. Lecoecuche, G. Mendell, T. Mistry, E. Payne, R. L. Savage, A. Viets, M. Wade, A. J. Weinstein, S. Aston, C. Cahillane, J. C. Driggers, S. E. Dwyer, and A. Urban, Characterization of systematic error in Advanced LIGO calibration in the second half of O3, Available as [arxiv:2107.00129](#), [arXiv:2107.00129](#) (2021), [arXiv:2107.00129 \[astro-ph.IM\]](#).
- [90] D. Estevez, P. Lagabbe, A. Masserot, L. Rolland, M. Seglar-Arroyo, and D. Verkindt, The Advanced Virgo photon calibrators, *Class. Quant. Grav.* **38**, 075007 (2021).
- [91] L. Sun, E. Goetz, J. S. Kissel, J. Betzwieser, S. Karki, A. Viets, M. Wade, D. Bhattacharjee, V. Bossilkov, P. B. Covas, and *et al.*, Characterization of systematic error in Advanced LIGO calibration, *Class. Quant. Grav.* **37**, 225008 (2020), [arXiv:2005.02531 \[astro-ph.IM\]](#).
- [92] E. Thrane and C. Talbot, An introduction to Bayesian inference in gravitational-wave astronomy: parameter estimation, model selection, and hierarchical models, *PASA* **36**, E010 (2019).
- [93] I. Mandel, W. M. Farr, and J. R. Gair, Extracting distribution parameters from multiple uncertain observations with selection biases, *Mon. Not. R. Ast. Soc.* **486**, 1086 (2019).
- [94] S. Vitale, One, No One, and One Hundred Thousand – Inferring the properties of a population in presence of selection effects, Available as [arxiv:2007.05579](#) (2020), [arXiv:2007.05579 \[astro-ph.IM\]](#).
- [95] T. J. Loredo, Accounting for Source Uncertainties in Analyses of Astronomical Survey Data, in *American Institute of Physics Conference Series*, Vol. 735 (2004) pp. 195–206.
- [96] M. Fishbach, D. E. Holz, and W. M. Farr, Does the Black Hole Merger Rate Evolve with Redshift?, *Astrophys. J. Lett.* **863**, L41 (2018), [arXiv:1805.10270 \[astro-ph.HE\]](#).
- [97] C. Talbot, R. Smith, E. Thrane, and G. B. Poole, Parallelized inference for gravitational-wave astronomy, *Phys. Rev. D* **100**, 043030 (2019).
- [98] D. Wysocki and R. O’Shaughnessy, *Bayesian parametric population models* (2017), [bayesian-parametric-population-models.readthedocs.io](#).
- [99] P. Landry and J. S. Read, The mass distribution of neutron stars in gravitational-wave binaries, Available as [arxiv:2107.04559](#), [arXiv:2107.04559](#) (2021), [arXiv:2107.04559 \[astro-ph.HE\]](#).
- [100] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, emcee: The MCMC Hammer, *Publ. Astron. Soc. Pac* **125**, 306 (2013).
- [101] J. S. Speagle, dynesty: A Dynamic Nested Sampling Package for Estimating Bayesian Posteriors and Evidences, *Mon. Not. R. Ast. Soc.* **493**, 3132 (2020).
- [102] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell, Stan: A probabilistic programming language, *Journal of Statistical Software* **76**, 10.18637/jss.v076.i01 (2017).
- [103] A. Riddell, A. Hartikainen, D. Lee, Riddell-Stan, *et al.*, *Stan-dev/pystan: V2.18.0.0* (2018).
- [104] C. Talbot and E. Thrane, Measuring the binary black hole mass spectrum with an astrophysically motivated parameterization, *Astrophys. J.* **856**, 173 (2018).
- [105] M. Fishbach, R. Essick, and D. E. Holz, Does Matter Matter? Using the Mass Distribution to Distinguish Neutron Stars and Black Holes, *Astrophys. J. Lett.* **899**, L8 (2020), [arXiv:2006.13178 \[astro-ph.HE\]](#).
- [106] A. M. Farah, M. Fishbach, R. Essick, and D. E. Holz, Bridging the Gap: Categorizing Gravitational-Wave Events at the Transition Between Neutron Stars and Black Holes, Available as [LIGO-P2100403](#) (2021).
- [107] M. Fishbach and D. E. Holz, Picky Partners: The Pairing of Component Masses in Binary Black Hole Mergers, *Astrophys. J. Lett.* **891**, L27 (2020), [arXiv:1905.12669 \[astro-ph.HE\]](#).
- [108] Z. Doctor, D. Wysocki, R. O’Shaughnessy, D. E. Holz, and B. Farr, Black hole coagulation: Modeling hierarchical mergers in black hole populations, *Astrophys. J.* **893**, 35 (2019).
- [109] B. Farr, E. Ochsner, W. M. Farr, and R. O’Shaughnessy, A more effective coordinate system for parameter estimation of precessing compact binaries from gravitational waves, *Phys. Rev. D* **90**, 024018 (2014), [arXiv:1404.7070 \[gr-qc\]](#).
- [110] D. Wysocki, J. Lange, and R. O’Shaughnessy, Reconstructing phenomenological distributions of compact binaries via gravitational wave observations, *Phys. Rev. D* **100**, 043012 (2019), [arXiv:1805.06442 \[gr-qc\]](#).
- [111] C. Talbot and E. Thrane, Determining the population properties of spinning black holes, *Phys. Rev. D* **96**, 023012 (2017).
- [112] S. Miller, T. A. Callister, and W. Farr, The Low Effective Spin of Binary Black Holes and Implications for Individual Gravitational-Wave Events, *Astrophys. J.* **895**, 128 (2020).
- [113] J. Roulet and M. Zaldarriaga, Constraints on binary black hole populations from LIGO-Virgo detections, *Mon. Not. R. Ast. Soc.* **484**, 4216 (2019), [arXiv:1806.10610 \[astro-ph.HE\]](#).
- [114] M. Fishbach and D. E. Holz, Where Are LIGO’s Big Black Holes?, *Astrophys. J. Lett.* **851**, L25 (2017).

- [115] D. Foreman-Mackey, D. W. Hogg, and T. D. Morton, Exoplanet population inference and the abundance of earth analogs from noisy, incomplete catalogs, *The Astrophysical Journal* **795**, 64 (2014).
- [116] I. Mandel, W. M. Farr, A. Colonna, S. Stevenson, P. Tiño, and J. Veitch, Model-independent inference on compact-binary observations, *Mon. Not. R. Ast. Soc.* **465**, 3254 (2017), arXiv:1608.08223 [astro-ph.HE].
- [117] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, Probabilistic programming in python using PyMC3, *PeerJ Computer Science* **2**, e55 (2016).
- [118] S. J. Kapadia *et al.*, A self-consistent method to estimate the rate of compact binary coalescences with a Poisson mixture model, *Class. Quant. Grav.* **37**, 045007 (2020).
- [119] Y. Lim and J. W. Holt, Bayesian modeling of the nuclear equation of state for neutron star tidal deformabilities and GW170817, *European Physical Journal A* **55**, 209 (2019), arXiv:1902.05502 [nucl-th].
- [120] T. Dietrich, M. W. Coughlin, P. T. H. Pang, M. Bulla, J. Heinzel, L. Issa, I. Tews, and S. Antier, Multimessenger constraints on the neutron-star equation of state and the Hubble constant, *Science* **370**, 1450 (2020), arXiv:2002.11355 [astro-ph.HE].
- [121] J.-L. Jiang, S.-P. Tang, Y.-Z. Wang, Y.-Z. Fan, and D.-M. Wei, PSR J0030+0451, GW170817, and the Nuclear Data: Joint Constraints on Equation of State and Bulk Properties of Neutron Stars, *Astrophys. J.* **892**, 55 (2020), arXiv:1912.07467 [astro-ph.HE].
- [122] B. Margalit and B. D. Metzger, Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817, *Astrophys. J. Lett.* **850**, L19 (2017), arXiv:1710.05938 [astro-ph.HE].
- [123] L. Rezzolla, E. R. Most, and L. R. Weih, Using Gravitational-wave Observations and Quasi-universal Relations to Constrain the Maximum Mass of Neutron Stars, *Astrophys. J. Lett.* **852**, L25 (2018), arXiv:1711.00314 [astro-ph.HE].
- [124] M. Ruiz, S. L. Shapiro, and A. Tsokaros, GW170817, general relativistic magnetohydrodynamic simulations, and the neutron star maximum mass, *Phys. Rev. D* **97**, 021501 (2018), arXiv:1711.00473 [astro-ph.HE].
- [125] M. Shibata, E. Zhou, K. Kiuchi, and S. Fujibayashi, Constraint on the maximum mass of neutron stars using GW170817 event, *Phys. Rev. D* **100**, 023015 (2019), arXiv:1905.03656 [astro-ph.HE].
- [126] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, C. Afeldt, M. Agathos, and *et al.*, Model comparison from LIGO-Virgo data on GW170817's binary components and consequences for the merger remnant, *Class. Quant. Grav.* **37**, 045006 (2020), arXiv:1908.01012 [gr-qc].
- [127] A. Nathanail, E. R. Most, and L. Rezzolla, GW170817 and GW190814: Tension on the Maximum Mass, *Astrophys. J. Lett.* **908**, L28 (2021), arXiv:2101.01735 [astro-ph.HE].
- [128] S. Galaudage, C. Adamcewicz, X.-J. Zhu, S. Stevenson, and E. Thrane, Heavy Double Neutron Stars: Birth, Midlife, and Death, *Astrophys. J. Lett.* **909**, L19 (2021), arXiv:2011.01495 [astro-ph.HE].
- [129] É. É. Flanagan and T. Hinderer, Constraining neutron-star tidal Love numbers with gravitational-wave detectors, *Phys. Rev. D* **77**, 021502 (2008), arXiv:0709.1915 [astro-ph].
- [130] J. Antoniadis, P. C. C. Freire, N. Wex, T. M. Tauris, R. S. Lynch, M. H. van Kerkwijk, M. Kramer, C. Bassa, V. S. Dhillon, T. Driebe, and *et al.*, A Massive Pulsar in a Compact Relativistic Binary, *Science* **340**, 448 (2013), arXiv:1304.6875 [astro-ph.HE].
- [131] H. T. Cromartie, E. Fonseca, S. M. Ransom, P. B. Demorest, Z. Arzoumanian, H. Blumer, P. R. Brook, M. E. DeCesar, T. Dolch, J. A. Ellis, and *et al.*, Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar, *Nature Astronomy* **4**, 72 (2020), arXiv:1904.06759 [astro-ph.HE].
- [132] I. Bombaci, The maximum mass of a neutron star., *Astron. Astrophys.* **305**, 871 (1996).
- [133] V. Kalogera and G. Baym, The Maximum Mass of a Neutron Star, *Astrophys. J. Lett.* **470**, L61 (1996), arXiv:astro-ph/9608059 [astro-ph].
- [134] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, and *et al.*, GW170817: Measurements of Neutron Star Radii and Equation of State, *Phys. Rev. Lett.* **121**, 161101 (2018), arXiv:1805.11581 [gr-qc].
- [135] G. B. Cook, S. L. Shapiro, and S. A. Teukolsky, Rapidly Rotating Neutron Stars in General Relativity: Realistic Equations of State, *Astrophys. J.* **424**, 823 (1994).
- [136] C. Breu and L. Rezzolla, Maximum mass, moment of inertia and compactness of relativistic stars, *Mon. Not. R. Ast. Soc.* **459**, 646 (2016), arXiv:1601.06083 [gr-qc].
- [137] F. Özel, D. Psaltis, R. Narayan, and A. Santos Villareal, On the Mass Distribution and Birth Masses of Neutron Stars, *Astrophys. J.* **757**, 55 (2012), arXiv:1201.1006 [astro-ph.HE].
- [138] J. Alsing, H. O. Silva, and E. Berti, Evidence for a maximum mass cut-off in the neutron star mass distribution and constraints on the equation of state, *Mon. Not. R. Ast. Soc.* **478**, 1377 (2018), arXiv:1709.07889 [astro-ph.HE].
- [139] W. M. Farr and K. Chatzioannou, A Population-Informed Mass Estimate for Pulsar J0740+6620, *Research Notes of the American Astronomical Society* **4**, 65 (2020), arXiv:2005.00032 [astro-ph.GA].
- [140] D.-S. Shao, S.-P. Tang, J.-L. Jiang, and Y.-Z. Fan, Maximum mass cutoff in the neutron star mass distribution and the prospect of forming supramassive objects in the double neutron star mergers, *Phys. Rev. D* **102**, 063006 (2020), arXiv:2009.04275 [astro-ph.HE].
- [141] K. Chatzioannou and W. M. Farr, Inferring the maximum and minimum mass of merging neutron stars with gravitational waves, *Phys. Rev. D* **102**, 064063 (2020), arXiv:2005.00482 [astro-ph.HE].
- [142] R. Essick, A. Farah, S. Galaudage, C. Talbot, M. Fishbach, E. Thrane, and D. E. Holz, Probing Extremal Gravitational-Wave Events with Coarse-Grained Likelihoods, arXiv e-prints, arXiv:2109.00418 (2021), arXiv:2109.00418 [astro-ph.HE].
- [143] E. Wilson, Probable inference, the law of succession, and statistical inference, *Journal of the American Statistical Association*, 209 (1927).
- [144] W. A. Fowler and F. Hoyle, Neutrino Processes and Pair Formation in Massive Stars and Supernovae., *Astrophys. J. Supp.* **9**, 201 (1964).
- [145] K. Belczynski, A. Heger, W. Gladysz, A. J. Ruiter,

- S. Woosley, G. Wiktorowicz, H. Y. Chen, T. Bulik, R. O’Shaughnessy, D. E. Holz, C. L. Fryer, and E. Berti, The effect of pair-instability mass loss on black-hole mergers, *Astron. Astrophys.* **594**, A97 (2016), [arXiv:1607.03116 \[astro-ph.HE\]](#).
- [146] S. E. Woosley, Pulsational pair-instability supernovae, *Astrophys. J.* **836**, 244 (2017).
- [147] M. Spera and M. Mapelli, Very massive stars, pair-instability supernovae and intermediate-mass black holes with the sevn code, *Mon. Not. R. Ast. Soc.* **470**, 4739 (2017).
- [148] P. Marchant, M. Renzo, R. Farmer, K. M. W. Pappas, R. E. Taam, S. E. de Mink, and V. Kalogera, Pulsational Pair-instability Supernovae in Very Close Binaries, *Astrophys. J.* **882**, 36 (2019), [arXiv:1810.13412 \[astro-ph.HE\]](#).
- [149] S. Stevenson, M. Sampson, J. Powell, A. Vigna-Gómez, C. J. Neijssel, D. Szécsi, and I. Mandel, The Impact of Pair-instability Mass Loss on the Binary Black Hole Mass Distribution, *Astrophys. J.* **882**, 121 (2019), [arXiv:1904.02821 \[astro-ph.HE\]](#).
- [150] S. E. Woosley and A. Heger, The Pair-instability Mass Gap for Black Holes, *Astrophys. J. Lett.* **912**, L31 (2021), [arXiv:2103.07933 \[astro-ph.SR\]](#).
- [151] R. Abbott *et al.* (LIGO Scientific, Virgo), Properties and Astrophysical Implications of the 150 M_{\odot} Binary Black Hole Merger GW190521, *Astrophys. J. Lett.* **900**, L13 (2020), [arXiv:2009.01190 \[astro-ph.HE\]](#).
- [152] A. H. Nitz and C. D. Capano, GW190521 May Be an Intermediate-mass Ratio Inspiral, *Astrophys. J. Lett.* **907**, L9 (2021), [arXiv:2010.12558 \[astro-ph.HE\]](#).
- [153] M. Fishbach and D. E. Holz, Minding the Gap: GW190521 as a Straddling Binary, *Astrophys. J. Lett.* **904**, L26 (2020), [arXiv:2009.05472 \[astro-ph.HE\]](#).
- [154] B. Edelman, Z. Doctor, and B. Farr, Poking Holes: Looking for Gaps in LIGO/Virgo’s Black Hole Population, *Astrophys. J. Lett.* **913**, L23 (2021), [arXiv:2104.07783 \[astro-ph.HE\]](#).
- [155] P. Madau and M. Dickinson, Cosmic Star-Formation History, *Annual Review of Astronomy and Astrophysics* **52**, 415 (2014).
- [156] R. O’Shaughnessy, V. Kalogera, and K. Belczynski, Binary Compact Object Coalescence Rates: The Role of Elliptical Galaxies, *Astrophys. J.* **716**, 615 (2010), [arXiv:0908.3635 \[astro-ph.CO\]](#).
- [157] M. Mapelli and N. Giacobbo, The cosmic merger rate of neutron stars and black holes, *Mon. Not. R. Ast. Soc.* **479**, 4391 (2018), [arXiv:1806.04866 \[astro-ph.HE\]](#).
- [158] C. L. Rodriguez and A. Loeb, Redshift Evolution of the Black Hole Merger Rate from Globular Clusters, *Astrophys. J. Lett.* **866**, L5 (2018), [arXiv:1809.01152 \[astro-ph.HE\]](#).
- [159] F. Santoliquido, M. Mapelli, Y. Bouffanais, N. Giacobbo, U. N. Di Carlo, S. Rastello, M. C. Artale, and A. Ballone, The Cosmic Merger Rate Density Evolution of Compact Binaries Formed in Young Star Clusters and in Isolated Binaries, *Astrophys. J.* **898**, 152 (2020), [arXiv:2004.09533 \[astro-ph.HE\]](#).
- [160] F. Antonini and M. Gieles, Merger rate of black hole binaries from globular clusters: Theoretical error bars and comparison to gravitational wave data from GWTC-2, *Phys. Rev. D* **102**, 123016 (2020), [arXiv:2009.01861 \[astro-ph.HE\]](#).
- [161] M. A. S. Martinez, G. Fragione, K. Kremer, S. Chatterjee, C. L. Rodriguez, J. Samsing, C. S. Ye, N. C. Weatherford, M. Zevin, S. Naoz, and F. A. Rasio, Black Hole Mergers from Hierarchical Triples in Dense Star Clusters, *Astrophys. J.* **903**, 67 (2020), [arXiv:2009.08468 \[astro-ph.GA\]](#).
- [162] Y. Yang, I. Bartos, Z. Haiman, B. Kocsis, S. Márka, and H. Tagawa, Cosmic Evolution of Stellar-mass Black Hole Merger Rate in Active Galactic Nuclei, *Astrophys. J.* **896**, 138 (2020), [arXiv:2003.08564 \[astro-ph.HE\]](#).
- [163] M. Fishbach and V. Kalogera, The Time Delay Distribution and Formation Metallicity of LIGO-Virgos Binary Black Holes, *Astrophys. J. Lett.* **914**, L30 (2021), [arXiv:2105.06491 \[astro-ph.HE\]](#).
- [164] L. A. C. van Son, S. E. de Mink, T. Callister, S. Justham, M. Renzo, T. Wagg, F. S. Broekgaarden, F. Kummer, R. Pakmor, and I. Mandel, The redshift evolution of the binary black hole merger rate: a weighty matter, Available as [arxiv:2110.01634](#), [arXiv:2110.01634](#) (2021), [arXiv:2110.01634 \[astro-ph.HE\]](#).
- [165] T. Regimbau, The astrophysical gravitational wave stochastic background, *Research in Astronomy and Astrophysics* **11**, 369 (2011), [arXiv:1101.2762 \[astro-ph.CO\]](#).
- [166] J. D. Romano and N. J. Cornish, Detection methods for stochastic gravitational-wave backgrounds: a unified treatment, *Living Reviews in Relativity* **20**, 2 (2017), [arXiv:1608.06889 \[gr-qc\]](#).
- [167] N. Christensen, Stochastic gravitational wave backgrounds, *Reports on Progress in Physics* **82**, 016903 (2019), [arXiv:1811.08797 \[gr-qc\]](#).
- [168] T. Callister, M. Fishbach, D. E. Holz, and W. M. Farr, Shouts and Murmurs: Combining Individual Gravitational-wave Sources with the Stochastic Background to Measure the History of Binary Black Hole Mergers, *Astrophys. J. Lett.* **896**, L32 (2020), [arXiv:2003.12152 \[astro-ph.HE\]](#).
- [169] R. Abbott, T. Abbott, S. Abraham, *et al.*, Upper limits on the isotropic gravitational-wave background from Advanced LIGO and Advanced Virgo’s third observing run, *Phys. Rev. D* **104**, 022004 (2021), [arXiv:2101.12130 \[gr-qc\]](#).
- [170] M. Mapelli, N. Giacobbo, F. Santoliquido, and M. C. Artale, The properties of merging black holes and neutron stars across cosmic time, *Mon. Not. R. Ast. Soc.* **487**, 2 (2019), [arXiv:1902.01419 \[astro-ph.HE\]](#).
- [171] C. J. Neijssel, A. Vigna-Gómez, S. Stevenson, J. W. Barrett, S. M. Gaebel, F. S. Broekgaarden, S. E. de Mink, D. Szécsi, S. Vinciguerra, and I. Mandel, The effect of the metallicity-specific star formation history on double compact object mergers, *Mon. Not. R. Ast. Soc.* **490**, 3740 (2019).
- [172] K. Belczynski, J. Klencki, C. E. Fields, A. Olejak, E. Berti, G. Meynet, C. L. Fryer, D. E. Holz, R. O’Shaughnessy, D. A. Brown, T. Bulik, S. C. Leung, K. Nomoto, P. Madau, R. Hirschi, E. Kaiser, S. Jones, S. Mondal, M. Chruslinska, P. Drozda, D. Gerosa, Z. Doctor, M. Giersz, S. Ekstrom, C. Georgy, A. Askar, V. Baibhav, D. Wysocki, T. Natan, W. M. Farr, G. Wiktorowicz, M. Coleman Miller, B. Farr, and J. P. Lasota, Evolutionary roads leading to low effective spins, high black hole masses, and O1/O2 rates for LIGO/Virgo binary black holes, *Astron. Astrophys.* **636**, A104 (2020), [arXiv:1706.07053 \[astro-ph.HE\]](#).

- [173] R. Farmer, M. Renzo, S. E. de Mink, P. Marchant, and S. Justham, Mind the Gap: The Location of the Lower Edge of the Pair-instability Supernova Black Hole Mass Gap, *Astrophys. J.* **887**, 53 (2019), [arXiv:1910.12874 \[astro-ph.SR\]](#).
- [174] J. S. Vink, E. R. Higgins, A. A. C. Sander, and G. N. Sabahit, Maximum black hole mass across cosmic time, *Mon. Not. R. Ast. Soc.* **504**, 146 (2021), [arXiv:2010.11730 \[astro-ph.HE\]](#).
- [175] T. Kinugawa, T. Nakamura, and H. Nakano, Formation of binary black holes similar to GW190521 with a total mass of $\sim 150 M_{\odot}$ from Population III binary star evolution, *Mon. Not. R. Ast. Soc.* **501**, L49 (2021), [arXiv:2009.06922 \[astro-ph.HE\]](#).
- [176] M. Fishbach, Z. Doctor, T. Callister, B. Edelman, J. Ye, R. Essick, W. M. Farr, B. Farr, and D. E. Holz, When are LIGO/Virgo's Big Black-Hole Mergers?, arXiv e-prints, [arXiv:2101.07699 \(2021\)](#), [arXiv:2101.07699 \[astro-ph.HE\]](#).
- [177] I. Mandel and R. O'Shaughnessy, Compact binary coalescences in the band of ground-based gravitational-wave detectors, *Class. Quant. Grav.* **27**, 114007 (2010).
- [178] D. Kushnir, M. Zaldarriaga, J. A. Kollmeier, and R. Waldman, GW150914: spin-based constraints on the merger time of the progenitor system, *Mon. Not. R. Ast. Soc.* **462**, 844 (2016), [arXiv:1605.03839 \[astro-ph.HE\]](#).
- [179] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, Illuminating Black Hole Binary Formation Channels with Spins in Advanced LIGO, *Astrophys. J. Lett.* **832**, L2 (2016), [arXiv:1609.05916 \[astro-ph.HE\]](#).
- [180] W. M. Farr, S. Stevenson, M. Coleman Miller, I. Mandel, B. Farr, and A. Vecchio, Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves, *Nature* **548**, 426 (2017).
- [181] S. Vitale, D. Gerosa, C.-J. Haster, K. Chatziioannou, and A. Zimmerman, Impact of bayesian priors on the characterization of binary black hole coalescences, *Phys. Rev. Lett.* **119**, 251103 (2017).
- [182] S. Stevenson, C. P. L. Berry, and I. Mandel, Hierarchical analysis of gravitational-wave measurements of binary black hole spin-orbit misalignments, *Mon. Not. R. Ast. Soc.* **471**, 2801 (2017), [arXiv:1703.06873 \[astro-ph.HE\]](#).
- [183] D. Gerosa, E. Berti, R. O'Shaughnessy, K. Belczynski, M. Kesden, D. Wysocki, and W. Gladysz, Spin orientations of merging black holes formed from the evolution of stellar binaries, *Phys. Rev. D* **98**, 084036 (2018).
- [184] R. O'Shaughnessy, D. Gerosa, and D. Wysocki, Inferences about Supernova Physics from Gravitational-Wave Measurements: GW151226 Spin Misalignment as an Indicator of Strong Black-Hole Natal Kicks, *Phys. Rev. Lett.* **119**, 011101 (2017).
- [185] J. Fuller, A. L. Piro, and A. S. Jermyn, Slowing the spins of stellar cores, *Mon. Not. R. Ast. Soc.* **485**, 3661 (2019).
- [186] J. Fuller and L. Ma, Most Black Holes Are Born Very Slowly Rotating, *Astrophys. J. Lett.* **881**, L1 (2019).
- [187] Y. Qin, P. Marchant, T. Fragos, G. Meynet, and V. Kalogera, On the Origin of Black Hole Spin in High-mass X-Ray Binaries, *Astrophys. J. Lett.* **870**, L18 (2019), [arXiv:1810.13016 \[astro-ph.SR\]](#).
- [188] S. S. Bavera, T. Fragos, Y. Qin, E. Zapartas, C. J. Neijssel, I. Mandel, A. Batta, S. M. Gaebel, C. Kimball, and S. Stevenson, The origin of spin in binary black holes. Predicting the distributions of the main observables of Advanced LIGO, *Astron. Astrophys.* **635**, A97 (2020), [arXiv:1906.12257 \[astro-ph.HE\]](#).
- [189] S. Biscoveanu, M. Isi, S. Vitale, and V. Varma, A new spin on LIGO-Virgo binary black holes, Available as [arxiv:2007.09156 \(2020\)](#), [arXiv:2007.09156 \[astro-ph.HE\]](#).
- [190] V. Kalogera, Spin-Orbit Misalignment in Close Binaries with Two Compact Objects, *Astrophys. J.* **541**, 319 (2000), [arXiv:astro-ph/9911417 \[astro-ph\]](#).
- [191] J. Roulet, H. S. Chia, S. Olsen, L. Dai, T. Venumadhav, B. Zackay, and M. Zaldarriaga, Distribution of effective spins and masses of binary black holes from the LIGO and Virgo O1-O3a observing runs, *Phys. Rev. D* **104**, 083010 (2021), [arXiv:2105.10580 \[astro-ph.HE\]](#).
- [192] S. Galadage, C. Talbot, T. Nagar, D. Jain, E. Thrane, and I. Mandel, Building better spin models for merging binary black holes: Evidence for non-spinning and rapidly spinning nearly aligned sub-populations, arXiv e-prints, [arXiv:2109.02424 \(2021\)](#), [arXiv:2109.02424 \[gr-qc\]](#).
- [193] T. A. Callister, C.-J. Haster, K. K. Y. Ng, S. Vitale, and W. M. Farr, Who Ordered That? Unequal-Mass Binary Black Hole Mergers Have Larger Effective Spins, arXiv e-prints, [arXiv:2106.00521 \(2021\)](#), [arXiv:2106.00521 \[astro-ph.HE\]](#).
- [194] D. Gerosa and M. Fishbach, Hierarchical mergers of stellar-mass black holes and their gravitational-wave signatures, *Nature Astronomy* **5**, 749 (2021), [arXiv:2105.03439 \[astro-ph.HE\]](#).
- [195] D. Gerosa and E. Berti, Are merging black holes born from stellar collapse or previous mergers?, *Phys. Rev. D* **95**, 124046 (2017), [arXiv:1703.06223 \[gr-qc\]](#).
- [196] M. Fishbach, D. E. Holz, and B. Farr, Are LIGO's Black Holes Made From Smaller Black Holes?, *Astrophys. J. Lett.* **840**, L24 (2017), [arXiv:1703.06869 \[astro-ph.HE\]](#).
- [197] B. McKernan, K. E. S. Ford, R. O'Shaughnessy, and D. Wysocki, Monte Carlo simulations of black hole mergers in AGN discs: Low χ_{eff} mergers and predictions for LIGO, *Mon. Not. R. Ast. Soc.* **494**, 1203 (2020), [arXiv:1907.04356 \[astro-ph.HE\]](#).
- [198] V. Baibhav, D. Gerosa, E. Berti, K. W. K. Wong, T. Helfer, and M. Mould, The mass gap, the spin gap, and the origin of merging binary black holes, *Phys. Rev. D* **102**, 043002 (2020), [arXiv:2004.00650 \[astro-ph.HE\]](#).
- [199] Y. Huang, C.-J. Haster, J. Roulet, S. Vitale, A. Zimmerman, T. Venumadhav, B. Zackay, L. Dai, and M. Zaldarriaga, Source properties of the lowest signal-to-noise-ratio binary black hole detections, *Phys. Rev. D* **102**, 103024 (2020), [arXiv:2003.04513 \[gr-qc\]](#).
- [200] S. Stevenson, F. Ohme, and S. Fairhurst, Distinguishing Compact Binary Population Synthesis Models Using Gravitational Wave Observations of Coalescing Binary Black Holes, *Astrophys. J.* **810**, 58 (2015), [arXiv:1504.07802 \[astro-ph.HE\]](#).
- [201] D. Wysocki, D. Gerosa, R. O'Shaughnessy, K. Belczynski, W. Gladysz, E. Berti, M. Kesden, and D. E. Holz, Explaining LIGO's observations via isolated binary evolution with natal kicks, *Phys. Rev. D* **97**, 043014 (2018), [arXiv:1709.01943 \[astro-ph.HE\]](#).
- [202] D. Gerosa, E. Berti, R. O'Shaughnessy, K. Belczynski, M. Kesden, D. Wysocki, and W. Gladysz, Spin orientations of merging black holes formed from the evolution of stellar binaries, *Phys. Rev. D* **98**, 084036 (2018), [arXiv:1808.02491 \[astro-ph.HE\]](#).

- [203] V. Tiwari, S. Fairhurst, and M. Hannam, Constraining black-hole spins with gravitational wave observations, *Astrophys. J.* **868**, 140 (2018), [arXiv:1809.01401 \[gr-qc\]](#).
- [204] J. W. Barrett, S. M. Gaebel, C. J. Neijssel, A. Vigna-Gómez, S. Stevenson, C. P. L. Berry, W. M. Farr, and I. Mandel, Accuracy of inference on the physics of binary evolution from gravitational-wave observations, *Mon. Not. R. Ast. Soc.* **477**, 4685 (2018), [arXiv:1711.06287 \[astro-ph.HE\]](#).
- [205] Y. Yang, I. Bartos, V. Gayathri, K. E. S. Ford, Z. Haiman, S. Klimentko, B. Kocsis, S. Márka, Z. Márka, B. McKernan, and R. O’Shaughnessy, Hierarchical Black Hole Mergers in Active Galactic Nuclei, *Phys. Rev. Lett.* **123**, 181101 (2019), [arXiv:1906.09281 \[astro-ph.HE\]](#).
- [206] B. Farr, D. E. Holz, and W. M. Farr, Using Spin to Understand the Formation of LIGO and Virgo’s Black Holes, *Astrophys. J. Lett.* **854**, L9 (2018).
- [207] M. Fishbach, W. M. Farr, and D. E. Holz, The Most Massive Binary Black Hole Detections and the Identification of Population Outliers, *Astrophys. J. Lett.* **891**, L31 (2020), [arXiv:1911.05882 \[astro-ph.HE\]](#).
- [208] I. Mandel, B. Müller, J. Riley, S. E. de Mink, A. Vigna-Gómez, and D. Chattopadhyay, Binary population synthesis with probabilistic remnant mass and kick prescriptions, *Mon. Not. R. Ast. Soc.* **500**, 1380 (2021), [arXiv:2007.03890 \[astro-ph.HE\]](#).
- [209] K. W. K. Wong, K. Breivik, K. Kremer, and T. Callister, Joint constraints on the field-cluster mixing fraction, common envelope efficiency, and globular cluster radii from a population of binary hole mergers via deep learning, *Phys. Rev. D* **103**, 083021 (2021), [arXiv:2011.03564 \[astro-ph.HE\]](#).
- [210] M. Zevin, S. S. Bavera, C. P. L. Berry, V. Kalogera, T. Fragos, P. Marchant, C. L. Rodriguez, F. Antonini, D. E. Holz, and C. Pankow, One Channel to Rule Them All? Constraining the Origins of Binary Black Holes Using Multiple Formation Pathways, *Astrophys. J.* **910**, 152 (2021), [arXiv:2011.10057 \[astro-ph.HE\]](#).
- [211] I. M. Romero-Shaw, K. Kremer, P. D. Lasky, E. Thrane, and J. Samsing, Gravitational waves as a probe of globular cluster formation and evolution, *Mon. Not. R. Ast. Soc.* **506**, 2362 (2021), [arXiv:2011.14541 \[astro-ph.HE\]](#).
- [212] S. R. Kulkarni, P. Hut, and S. McMillan, Stellar black holes in globular clusters, *Nature* **364**, 421 (1993).
- [213] S. Sigurdsson and L. Hernquist, Primordial black holes in globular clusters, *Nature* **364**, 423 (1993).
- [214] S. F. Portegies Zwart and S. L. W. McMillan, Black Hole Mergers in the Universe, *Astrophys. J. Lett.* **528**, L17 (2000), [arXiv:astro-ph/9910061 \[astro-ph\]](#).
- [215] R. M. O’Leary, R. O’Shaughnessy, and F. A. Rasio, Dynamical interactions and the black-hole merger rate of the Universe, *Phys. Rev. D* **76**, 061504 (2007), [arXiv:astro-ph/0701887 \[astro-ph\]](#).
- [216] S. Banerjee, H. Baumgardt, and P. Kroupa, Stellar-mass black holes in star clusters: implications for gravitational wave radiation, *Mon. Not. R. Ast. Soc.* **402**, 371 (2010), [arXiv:0910.3954 \[astro-ph.SR\]](#).
- [217] C. L. Rodriguez, M. Morscher, B. Pattabiraman, S. Chatterjee, C.-J. Haster, and F. A. Rasio, Binary Black Hole Mergers from Globular Clusters: Implications for Advanced LIGO, *Phys. Rev. Lett.* **115**, 051101 (2015), [arXiv:1505.00792 \[astro-ph.HE\]](#).
- [218] C. L. Rodriguez, S. Chatterjee, and F. A. Rasio, Binary black hole mergers from globular clusters: Masses, merger rates, and the impact of stellar evolution, *Phys. Rev. D* **93**, 084029 (2016), [arXiv:1602.02444 \[astro-ph.HE\]](#).
- [219] A. Askar, M. Arca Sedda, and M. Giersz, MOCCA-SURVEY Database I: Galactic globular clusters harbouring a black hole subsystem, *Mon. Not. R. Ast. Soc.* **478**, 1844 (2018), [arXiv:1802.05284 \[astro-ph.GA\]](#).
- [220] J. Hong, E. Vesperini, A. Askar, M. Giersz, M. Szkudlarek, and T. Bulik, Binary black hole mergers from globular clusters: the impact of globular cluster properties, *Mon. Not. R. Ast. Soc.* **480**, 5645 (2018), [arXiv:1808.04514 \[astro-ph.HE\]](#).
- [221] C. L. Rodriguez, M. Zevin, P. Amaro-Seoane, S. Chatterjee, K. Kremer, F. A. Rasio, and C. S. Ye, Black holes: The next generation—repeated mergers in dense star clusters and their gravitational-wave properties, *Phys. Rev. D* **100**, 043027 (2019), [arXiv:1906.10260 \[astro-ph.HE\]](#).
- [222] F. Santoliquido, M. Mapelli, N. Giacobbo, Y. Bouffanais, and M. C. Artale, The cosmic merger rate density of compact objects: impact of star formation, metallicity, initial mass function, and binary evolution, *Mon. Not. R. Ast. Soc.* **502**, 4877 (2021), [arXiv:2009.03911 \[astro-ph.HE\]](#).
- [223] S. Banerjee, Stellar-mass black holes in young massive and open stellar clusters - V. comparisons with LIGO-Virgo merger rate densities, *Mon. Not. R. Ast. Soc.* **503**, 3371 (2021), [arXiv:2011.07000 \[astro-ph.HE\]](#).
- [224] S. Banerjee, Merger rate density of stellar-mass binary black holes from young massive clusters, open clusters, and isolated binaries: comparisons with LIGO-Virgo-KAGRA results, *arXiv e-prints*, [arXiv:2108.04250](#) (2021), [arXiv:2108.04250 \[astro-ph.HE\]](#).
- [225] F. Antonini and H. B. Perets, Secular Evolution of Compact Binaries near Massive Black Holes: Gravitational Wave Sources and Other Exotica, *ApJ* **757**, 27 (2012), [arXiv:1203.2938 \[astro-ph.GA\]](#).
- [226] F. Antonini and F. A. Rasio, Merging black hole binaries in galactic nuclei: Implications for advanced-ligo detections, *Astrophys. J.* **831**, 187 (2016).
- [227] B.-M. Hoang, S. Naoz, B. Kocsis, F. A. Rasio, and F. Dosopoulou, Black Hole Mergers in Galactic Nuclei Induced by the Eccentric Kozai-Lidov Effect, *ApJ* **856**, 140 (2018), [arXiv:1706.09896 \[astro-ph.HE\]](#).
- [228] G. Fragione, E. Grishin, N. W. C. Leigh, H. B. Perets, and R. Perna, Black hole and neutron star mergers in galactic nuclei, *Mon. Not. R. Ast. Soc.* **488**, 47 (2019), [arXiv:1811.10627 \[astro-ph.GA\]](#).
- [229] M. Arca Sedda, M. Mapelli, M. Spera, M. Benacquista, and N. Giacobbo, Fingerprints of Binary Black Hole Formation Channels Encoded in the Mass and Spin of Merger Remnants, *Astrophys. J.* **894**, 133 (2020), [arXiv:2003.07409 \[astro-ph.GA\]](#).
- [230] M. Mapelli, M. Dall’Amico, Y. Bouffanais, N. Giacobbo, M. Arca Sedda, M. C. Artale, A. Ballone, U. N. Di Carlo, G. Iorio, F. Santoliquido, and S. Torniamenti, Hierarchical black hole mergers in young, globular and nuclear star clusters: the effect of metallicity, spin and cluster properties, *Mon. Not. R. Ast. Soc.* **505**, 339 (2021), [arXiv:2103.05016 \[astro-ph.HE\]](#).
- [231] L. Gondán and B. Kocsis, High eccentricities and high masses characterize gravitational-wave captures in galactic nuclei as seen by Earth-based detectors, *MN-*

- RAS **506**, 1665 (2021), [arXiv:2011.02507 \[astro-ph.HE\]](#).
- [232] K. E. S. Ford and B. McKernan, Binary Black Hole Merger Rates in AGN Disks versus Nuclear Star Clusters: Loud beats Quiet, [arXiv e-prints](#), [arXiv:2109.03212](#) (2021), [arXiv:2109.03212 \[astro-ph.HE\]](#).
- [233] H. Tagawa, B. Kocsis, Z. Haiman, I. Bartos, K. Omukai, and J. Samsing, Mass-gap Mergers in Active Galactic Nuclei, *Astrophys. J.* **908**, 194 (2021), [arXiv:2012.00011 \[astro-ph.HE\]](#).
- [234] Y. Yang, I. Bartos, Z. Haiman, B. Kocsis, Z. Márka, N. C. Stone, and S. Márka, AGN Disks Harden the Mass Distribution of Stellar-mass Binary Black Hole Mergers, *Astrophys. J.* **876**, 122 (2019), [arXiv:1903.01405 \[astro-ph.HE\]](#).
- [235] M. Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, T. Bulik, I. Mandel, and R. O’Shaughnessy, Double Compact Objects. II. Cosmological Merger Rates, *Astrophys. J.* **779**, 72 (2013).
- [236] K. Belczynski, D. E. Holz, T. Bulik, and R. O’Shaughnessy, The first gravitational-wave source from the isolated evolution of two stars in the 40-100 solar mass range, *Nature* **534**, 512 (2016), [arXiv:1602.04531 \[astro-ph.HE\]](#).
- [237] N. Giacobbo and M. Mapelli, The progenitors of compact-object binaries: impact of metallicity, common envelope and natal kicks, *Mon. Not. R. Ast. Soc.* **480**, 2011 (2018), [arXiv:1806.00001 \[astro-ph.HE\]](#).
- [238] G. Wiktorowicz, L. Wyrzykowski, M. Chruslinska, J. Klencki, K. A. Rybicki, and K. Belczynski, Populations of Stellar-mass Black Holes from Binary Systems, *Astrophys. J.* **885**, 1 (2019), [arXiv:1907.11431 \[astro-ph.HE\]](#).
- [239] S. E. de Mink and K. Belczynski, Merger Rates of Double Neutron Stars and Stellar Origin Black Holes: The Impact of Initial Conditions on Binary Evolution Predictions, *Astrophys. J.* **814**, 58 (2015), [arXiv:1506.03573 \[astro-ph.HE\]](#).
- [240] A. Olejak, K. Belczynski, and N. Ivanova, Impact of common envelope development criteria on the formation of LIGO/Virgo sources, *Astron. Astrophys.* **651**, A100 (2021), [arXiv:2102.05649 \[astro-ph.HE\]](#).
- [241] R. Farmer, M. Renzo, S. E. de Mink, M. Fishbach, and S. Justham, Constraints from Gravitational-wave Detections of Binary Black Hole Mergers on the $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ Rate, *Astrophys. J. Lett.* **902**, L36 (2020), [arXiv:2006.06678 \[astro-ph.HE\]](#).
- [242] D. Croon, S. D. McDermott, and J. Sakstein, New physics and the black hole mass gap, *Phys. Rev. D* **102**, 115024 (2020).
- [243] G. Costa, A. Bressan, M. Mapelli, P. Marigo, G. Iorio, and M. Spera, Formation of GW190521 from stellar evolution: the impact of the hydrogen-rich envelope, dredge-up, and $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ rate on the pair-instability black hole mass gap, *Mon. Not. R. Ast. Soc.* **501**, 4514 (2021), [arXiv:2010.02242 \[astro-ph.SR\]](#).
- [244] E. Farrell, J. H. Groh, R. Hirschi, L. Murphy, E. Kaiser, S. Ekström, C. Georgy, and G. Meynet, Is GW190521 the merger of black holes from the first stellar generations?, *Mon. Not. R. Ast. Soc.* **502**, L40 (2021), [arXiv:2009.06585 \[astro-ph.SR\]](#).
- [245] K. Takahashi, The Low Detection Rate of Pair-instability Supernovae and the Effect of the Core Carbon Fraction, *Astrophys. J.* **863**, 153 (2018), [arXiv:1807.05373 \[astro-ph.HE\]](#).
- [246] F. Antonini, M. Gieles, and A. Gualandris, Black hole growth through hierarchical black hole mergers in dense star clusters: implications for gravitational wave detections, *Mon. Not. R. Ast. Soc.* **486**, 5008 (2019), [arXiv:1811.03640 \[astro-ph.HE\]](#).
- [247] C. Kimball, C. Talbot, C. P. L. Berry, M. Carney, M. Zevin, E. Thrane, and V. Kalogera, Black Hole Genealogy: Identifying Hierarchical Mergers with Gravitational Waves, *Astrophys. J.* **900**, 177 (2020), [arXiv:2005.00023 \[astro-ph.HE\]](#).
- [248] G. Fragione, A. Loeb, and F. A. Rasio, On the Origin of GW190521-like Events from Repeated Black Hole Mergers in Star Clusters, *Astrophys. J. Lett.* **902**, L26 (2020), [arXiv:2009.05065 \[astro-ph.GA\]](#).
- [249] A. Palmese and C. J. Conselice, GW190521 from the Merger of Ultradwarf Galaxies, *Phys. Rev. Lett.* **126**, 181103 (2021), [arXiv:2009.10688 \[astro-ph.GA\]](#).
- [250] H. Tagawa, Z. Haiman, I. Bartos, B. Kocsis, and K. Omukai, Signatures of hierarchical mergers in black hole spin and mass distribution, *Mon. Not. R. Ast. Soc.* **10.1093/mnras/stab2315** (2021), [arXiv:2104.09510 \[astro-ph.HE\]](#).
- [251] U. D. Carlo *et al.*, Merging black holes in young star clusters, *Mon. Not. R. Ast. Soc.* **487**, 2947 (2019).
- [252] U. N. Di Carlo, M. Mapelli, Y. Bouffanais, N. Giacobbo, F. Santoliquido, A. Bressan, M. Spera, and F. Haardt, Binary black holes in the pair instability mass gap, *Mon. Not. R. Ast. Soc.* **497**, 1043 (2020), [arXiv:1911.01434 \[astro-ph.HE\]](#).
- [253] K. Kremer, M. Spera, D. Becker, S. Chatterjee, U. N. Di Carlo, G. Fragione, C. L. Rodriguez, C. S. Ye, and F. A. Rasio, Populating the Upper Black Hole Mass Gap through Stellar Collisions in Young Star Clusters, *Astrophys. J.* **903**, 45 (2020), [arXiv:2006.10771 \[astro-ph.HE\]](#).
- [254] M. Renzo, M. Cantiello, B. D. Metzger, and Y. F. Jiang, The Stellar Merger Scenario for Black Holes in the Pair-instability Gap, *Astrophys. J. Lett.* **904**, L13 (2020), [arXiv:2010.00705 \[astro-ph.SR\]](#).
- [255] K. Inayoshi, Z. Haiman, and J. P. Ostriker, Hyper-Eddington accretion flows on to massive black holes, *Mon. Not. R. Ast. Soc.* **459**, 3738 (2016), [arXiv:1511.02116 \[astro-ph.HE\]](#).
- [256] Z. Roupas and D. Kazanas, Generation of massive stellar black holes by rapid gas accretion in primordial dense clusters, *Astron. Astrophys.* **632**, L8 (2019), [arXiv:1911.03915 \[astro-ph.GA\]](#).
- [257] L. A. C. van Son, S. E. De Mink, F. S. Broekgaarden, M. Renzo, S. Justham, E. Laplace, J. Morán-Fraile, D. D. Hendriks, and R. Farmer, Polluting the Pair-instability Mass Gap for Binary Black Holes through Super-Eddington Accretion in Isolated Binaries, *Astrophys. J.* **897**, 100 (2020), [arXiv:2004.05187 \[astro-ph.HE\]](#).
- [258] M. Safarzadeh and Z. Haiman, Formation of GW190521 via Gas Accretion onto Population III Stellar Black Hole Remnants Born in High-redshift Minihalos, *Astrophys. J. Lett.* **903**, L21 (2020), [arXiv:2009.09320 \[astro-ph.HE\]](#).
- [259] S. S. Bavera, T. Fragos, M. Zevin, C. P. L. Berry, P. Marchant, J. J. Andrews, S. Coughlin, A. Dotter, K. Kovalakas, D. Misra, J. G. Serra-Perez, Y. Qin, K. A. Rocha, J. Román-Garza, N. H. Tran, and E. Zapartas,

- The impact of mass-transfer physics on the observable properties of field binary black hole populations, *Astron. Astrophys.* **647**, A153 (2021), [arXiv:2010.16333 \[astro-ph.HE\]](#).
- [260] V. De Luca, V. Desjacques, G. Franciolini, P. Pani, and A. Riotta, GW190521 Mass Gap Event and the Primordial Black Hole Scenario, *Phys. Rev. Lett.* **126**, 051101 (2021), [arXiv:2009.01728 \[astro-ph.CO\]](#).
- [261] V. Baibhav, E. Berti, D. Gerosa, M. Mapelli, N. Giacobbo, Y. Bouffanais, and U. N. Di Carlo, Gravitational-wave detection rates for compact binaries formed in isolation: LIGO/Virgo O3 and beyond, *Phys. Rev. D* **100**, 064060 (2019).
- [262] M. Gallegos-Garcia, C. P. L. Berry, P. Marchant, and V. Kalogera, Binary Black Hole Formation with Detailed Modeling: Stable Mass Transfer Leads to Lower Merger Rates, *arXiv e-prints*, [arXiv:2107.05702 \(2021\)](#), [arXiv:2107.05702 \[astro-ph.HE\]](#).
- [263] F. Santoliquido, M. Mapelli, Y. Bouffanais, N. Giacobbo, U. N. Di Carlo, S. Rastello, M. C. Artale, and A. Ballone, The cosmic merger rate density evolution of compact binaries formed in young star clusters and in isolated binaries, *arXiv e-prints*, [arXiv:2004.09533 \(2020\)](#).
- [264] G. Fragione and B. Kocsis, Black Hole Mergers from an Evolving Population of Globular Clusters, *Phys. Rev. Lett.* **121**, 161103 (2018), [arXiv:1806.02351 \[astro-ph.GA\]](#).
- [265] H. Tagawa, Z. Haiman, I. Bartos, and B. Kocsis, Spin Evolution of Stellar-mass Black Hole Binaries in Active Galactic Nuclei, *Astrophys. J.* **899**, 26 (2020), [arXiv:2004.11914 \[astro-ph.HE\]](#).
- [266] S. Albrecht, S. Reffert, I. A. G. Snellen, and J. N. Winn, Misaligned spin and orbital axes cause the anomalous precession of DI Herculis, *Nature* **461**, 373 (2009), [arXiv:0909.2861 \[astro-ph.SR\]](#).
- [267] S. Albrecht, J. N. Winn, J. A. Carter, I. A. G. Snellen, and E. J. W. de Mooij, The Banana Project. III. Spin-Orbit Alignment in the Long-period Eclipsing Binary NY Cephei, *Astrophys. J.* **726**, 68 (2011), [arXiv:1011.0425 \[astro-ph.SR\]](#).
- [268] S. Albrecht, J. Setiawan, G. Torres, D. C. Fabrycky, and J. N. Winn, The BANANA Project. IV. Two Aligned Stellar Rotation Axes in the Young Eccentric Binary System EP Crucis: Primordial Orientation and Tidal Alignment, *Astrophys. J.* **767**, 32 (2013), [arXiv:1211.7065 \[astro-ph.SR\]](#).
- [269] S. Albrecht, J. N. Winn, G. Torres, D. C. Fabrycky, J. Setiawan, M. Gillon, E. Jehin, A. Triaud, D. Queloz, I. Snellen, and P. Eggleton, The BANANA Project. V. Misaligned and Precessing Stellar Rotation Axes in CV Velorum, *Astrophys. J.* **785**, 83 (2014), [arXiv:1403.0583 \[astro-ph.SR\]](#).
- [270] F. Antonini, C. L. Rodriguez, C. Petrovich, and C. L. Fischer, Precessional dynamics of black hole triples: binary mergers with near-zero effective spin, *Mon. Not. R. Ast. Soc. Lett.* **480**, L58 (2018).
- [271] C. L. Rodriguez and F. Antonini, A Triple Origin for the Heavy and Low-spin Binary Black Holes Detected by LIGO/VIRGO, *Astrophys. J.* **863**, 7 (2018).
- [272] B. Liu, D. Lai, and Y.-H. Wang, Black Hole and Neutron Star Binary Mergers in Triple Systems. II. Merger Eccentricity and Spin-Orbit Misalignment, *Astrophys. J.* **881**, 41 (2019).
- [273] A. S. Hamers, G. Fragione, P. Neunteufel, and B. Kocsis, First- and second-generation black hole and neutron star mergers in 2+2 quadruples: population statistics, *MNRAS* **506**, 5345 (2021), [arXiv:2103.03782 \[astro-ph.HE\]](#).
- [274] C. Petrovich and F. Antonini, Greatly Enhanced Merger Rates of Compact-object Binaries in Non-spherical Nuclear Star Clusters, *Astrophys. J.* **846**, 146 (2017), [arXiv:1705.05848 \[astro-ph.HE\]](#).
- [275] H. Yu, S. Ma, M. Giesler, and Y. Chen, Spin and eccentricity evolution in triple systems: From the Lidov-Kozai interaction to the final merger of the inner binary, *Phys. Rev. D* **102**, 123009 (2020), [arXiv:2007.12978 \[gr-qc\]](#).
- [276] Y. Su, D. Lai, and B. Liu, Spin-orbit misalignments in tertiary-induced binary black-hole mergers: Theoretical analysis, *Phys. Rev. D* **103**, 063040 (2021), [arXiv:2010.11951 \[gr-qc\]](#).
- [277] J. Stegmann and F. Antonini, Flipping spins in mass transferring binaries and origin of spin-orbit misalignment in binary black holes, *Phys. Rev. D* **103**, 063007 (2021), [arXiv:2012.06329 \[astro-ph.HE\]](#).
- [278] C. Chan, B. Müller, and A. Heger, The impact of fallback on the compact remnants and chemical yields of core-collapse supernovae, *Mon. Not. R. Ast. Soc.* **495**, 3751 (2020), [arXiv:2003.04320 \[astro-ph.SR\]](#).
- [279] W. M. Farr, K. Kremer, M. Lyutikov, and V. Kalogera, Spin Tilts in the Double Pulsar Reveal Supernova Spin Angular-momentum Production, *Astrophys. J.* **742**, 81 (2011), [arXiv:1104.5001 \[astro-ph.HE\]](#).
- [280] Y. Qin, T. Fragos, G. Meynet, J. Andrews, M. Sørensen, and H. F. Song, The spin of the second-born black hole in coalescing binary black holes, *Astron. Astrophys.* **616**, A28 (2018), [arXiv:1802.05738 \[astro-ph.SR\]](#).
- [281] I. Mandel and T. Fragos, An alternative interpretation of GW190412 as a binary black hole merger with a rapidly spinning secondary, *Astrophys. J. Lett.* **895**, L28 (2020), [arXiv:2004.09288 \[astro-ph.HE\]](#).
- [282] P. Marchant, N. Langer, P. Podsiadlowski, T. M. Tauris, and T. J. Moriya, A new route towards merging massive black holes, *Astron. Astrophys.* **588**, A50 (2016).
- [283] I. Mandel and S. E. de Mink, Merging binary black holes formed through chemically homogeneous evolution in short-period stellar binaries, *Mon. Not. R. Ast. Soc.* **458**, 2634 (2016).
- [284] V. De Luca, V. Desjacques, G. Franciolini, A. Malhotra, and A. Riotta, The initial spin probability distribution of primordial black holes, *J. Cosmology Astropart. Phys.* **2019**, 018 (2019), [arXiv:1903.01179 \[astro-ph.CO\]](#).
- [285] T. Chiba and S. Yokoyama, Spin distribution of primordial black holes, *Progress of Theoretical and Experimental Physics* **2017**, 083E01 (2017), [arXiv:1704.06573 \[gr-qc\]](#).
- [286] M. Safarzadeh, W. M. Farr, and E. Ramirez-Ruiz, A Trend in the Effective Spin Distribution of LIGO Binary Black Holes with Mass, *Astrophys. J.* **894**, 129 (2020), [arXiv:2001.06490 \[gr-qc\]](#).
- [287] C. Hoy, S. Fairhurst, M. Hannam, and V. Tiwari, Understanding how fast black holes spin by analysing data from the second gravitational-wave catalogue, *arXiv e-prints* (2021), [arXiv:2110.13542 \[gr-qc\]](#).
- [288] M. Safarzadeh, E. Ramirez-Ruiz, and E. Berger, Does GW190425 Require an Alternative Formation Pathway

- than a Fast-merging Channel?, *Astrophys. J.* **900**, 13 (2020), [arXiv:2001.04502 \[astro-ph.HE\]](#).
- [289] I. M. Romero-Shaw, N. Farrow, S. Stevenson, E. Thrane, and X.-J. Zhu, On the origin of GW190425, *Mon. Not. R. Ast. Soc.* **496**, L64 (2020), [arXiv:2001.06492 \[astro-ph.HE\]](#).
- [290] S. E. Woosley, A. Heger, and T. A. Weaver, The evolution and explosion of massive stars, *Rev. of Mod. Phys.* **74**, 1015 (2002).
- [291] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann, How Massive Single Stars End their Life, *Astrophys. J.* **591**, 288 (2003).
- [292] D. Radice, A. Perego, K. Hotokezaka, S. A. Fromm, S. Bernuzzi, and L. F. Roberts, Binary Neutron Star Mergers: Mass Ejection, Electromagnetic Counterparts, and Nucleosynthesis, *Astrophys. J.* **869**, 130 (2018), [arXiv:1809.11161 \[astro-ph.HE\]](#).
- [293] M. W. Coughlin, T. Dietrich, B. Margalit, and B. D. Metzger, Multimessenger Bayesian parameter inference of a binary neutron star merger, *Mon. Not. R. Ast. Soc.* **489**, L91 (2019), [arXiv:1812.04803 \[astro-ph.HE\]](#).
- [294] J. M. Lattimer and D. N. Schramm, Black-Hole-Neutron-Star Collisions, *Astrophys. J. Lett.* **192**, L145 (1974).
- [295] K. Kiuchi, K. Kyutoku, M. Shibata, and K. Taniguchi, Revisiting the Lower Bound on Tidal Deformability Derived by AT 2017gfo, *Astrophys. J. Lett.* **876**, L31 (2019), [arXiv:1903.01466 \[astro-ph.HE\]](#).
- [296] F. Foucart, T. Hinderer, and S. Nissanke, Remnant baryon mass in neutron star-black hole mergers: Predictions for binary neutron star mimickers and rapidly spinning black holes, *Phys. Rev. D* **98**, 081501 (2018), [arXiv:1807.00011 \[astro-ph.HE\]](#).
- [297] C. J. Krüger and F. Foucart, Estimates for disk and ejecta masses produced in compact binary mergers, *Phys. Rev. D* **101**, 103002 (2020), [arXiv:2002.07728 \[astro-ph.HE\]](#).
- [298] R. Mochkovitch, M. Hernanz, J. Isern, and X. Martin, Gamma-ray bursts as collimated jets from neutron star/black hole mergers, *Nature* **361**, 236 (1993).
- [299] H. T. Janka, T. Eberl, M. Ruffert, and C. L. Fryer, Black Hole-Neutron Star Mergers as Central Engines of Gamma-Ray Bursts, *Astrophys. J. Lett.* **527**, L39 (1999), [arXiv:astro-ph/9908290 \[astro-ph\]](#).
- [300] S. E. Woosley, Gamma-Ray Bursts from Stellar Mass Accretion Disks around Black Holes, *Astrophys. J.* **405**, 273 (1993).
- [301] B. Margalit and B. D. Metzger, The Multi-messenger Matrix: The Future of Neutron Star Merger Constraints on the Nuclear Equation of State, *Astrophys. J. Lett.* **880**, L15 (2019), [arXiv:1904.11995 \[astro-ph.HE\]](#).
- [302] B. Liu and D. Lai, Hierarchical black hole mergers in multiple systems: constrain the formation of GW190412-, GW190814-, and GW190521-like events, *Mon. Not. R. Ast. Soc.* **502**, 2049 (2021), [arXiv:2009.10068 \[astro-ph.HE\]](#).
- [303] W. Lu, P. Beniamini, and C. Bonnerot, On the formation of GW190814, *Mon. Not. R. Ast. Soc.* **500**, 1817 (2021), [arXiv:2009.10082 \[astro-ph.HE\]](#).
- [304] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, *et al.*, Search for the isotropic stochastic background using data from Advanced LIGO's second observing run, *prd* **100**, 061101 (2019), [arXiv:1903.02886 \[gr-qc\]](#).
- [305] E. Vangioni, K. A. Olive, T. Prestegard, J. Silk, P. Petitjean, and V. Mandic, The impact of star formation and gamma-ray burst rates at high redshift on cosmic chemical evolution and reionization, *Mon. Not. R. Ast. Soc.* **447**, 2575 (2015), [arXiv:1409.2462 \[astro-ph.GA\]](#).
- [306] E. Vangioni, S. Goriely, F. Daigne, P. François, and K. Belczynski, Cosmic neutron-star merger rate and gravitational waves constrained by the r-process nucleosynthesis, *Mon. Not. R. Ast. Soc.* **455**, 17 (2016), [arXiv:1501.01115 \[astro-ph.HE\]](#).
- [307] E. Vangioni and K. A. Olive, The cosmic evolution of magnesium isotopes, *Mon. Not. R. Ast. Soc.* **484**, 3561 (2019), [arXiv:1809.10514 \[astro-ph.GA\]](#).
- [308] M. Dominik, K. Belczynski, C. Fryer, D. E. Holz, E. Berti, T. Bulik, I. Mandel, and R. O'Shaughnessy, Double Compact Objects. I. The Significance of the Common Envelope on Merger Rates, *apj* **759**, 52 (2012), [arXiv:1202.4901 \[astro-ph.HE\]](#).
- [309] M. Chruslinska, G. Nelemans, and K. Belczynski, The influence of the distribution of cosmic star formation at different metallicities on the properties of merging double compact objects, *Mon. Not. R. Ast. Soc.* **482**, 5012 (2019), [arXiv:1811.03565 \[astro-ph.HE\]](#).
- [310] E. Thrane and J. D. Romano, Sensitivity curves for searches for gravitational-wave backgrounds, *prd* **88**, 124032 (2013), [arXiv:1310.5300 \[astro-ph.IM\]](#).
- [311] S. R. Taylor, J. R. Gair, and I. Mandel, Cosmology using advanced gravitational-wave detectors alone, *Phys. Rev. D* **85**, 023535 (2012), [arXiv:1108.5161 \[gr-qc\]](#).
- [312] W. M. Farr, M. Fishbach, J. Ye, and D. E. Holz, A Future Percent-level Measurement of the Hubble Expansion at Redshift 0.8 with Advanced LIGO, *Astrophys. J. Lett.* **883**, L42 (2019), [arXiv:1908.09084 \[astro-ph.CO\]](#).
- [313] R. Abbott, T. Abbott, S. Abraham, *et al.*, Constraints on the cosmic expansion history from the third LIGO-Virgo-KAGRA Gravitational-Wave Transient Catalog, Available as LIGO-P2100185 (2021).
- [314] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, *Living Reviews in Relativity* **23**, 3 (2020).
- [315] R. Abbott, T. Abbott, S. Abraham, *et al.*, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run - O3 search sensitivity estimates, Zenodo, available as <https://zenodo.org/deposit/5546676> 10.5281/zenodo.5546675 (2021).
- [316] V. Tiwari, Estimation of the sensitive volume for gravitational-wave source populations using weighted Monte Carlo integration, *Class. Quant. Grav.* **35**, 145009 (2018), [arXiv:1712.00482 \[astro-ph.HE\]](#).
- [317] W. M. Farr, Accuracy Requirements for Empirically Measured Selection Functions, Research Notes of the American Astronomical Society **3**, 66 (2019).
- [318] R. Abbott, T. Abbott, S. Abraham, *et al.*, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run - O1+O2+O3 search sensitivity estimates, Zenodo, available as <https://zenodo.org/deposit/5636816> 10.5281/zenodo.5636816 (2021).
- [319] D. Wysocki and R. O'Shaughnessy, Popmodels o3a aps april 2021 presentation, <https://gitlab.com/dwysocki/pop-models-o3a-aps-april-2021> (2021).

- [320] R. Abbott, T. D. Abbott, S. Abraham, et al, Ligo Scientific Collaboration, VIRGO Collaboration, and KAGRA Collaboration, Observation of Gravitational Waves from Two Neutron Star-Black Hole Coalescences, *Astrophys. J. Lett.* **915**, L5 (2021), [arXiv:2106.15163 \[astro-ph.HE\]](#).
- [321] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, N. Mazumder, and A. J. Weinstein, Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins, *Phys. Rev. D* **89**, 084041 (2014).
- [322] P. Ajith, Addressing the spin question in gravitational-wave searches: Waveform templates for inspiralling compact binaries with nonprecessing spins, *Phys. Rev. D* **84**, 084037 (2011), [arXiv:1107.1267 \[gr-qc\]](#).
- [323] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, Searching for gravitational waves from compact binaries with precessing spins, *Physical Review D* **94**, 10.1103/physrevd.94.024012 (2016).