

PAPER

Too small to fail: characterizing sub-solar mass black hole mergers with gravitational waves

To cite this article: Noah E. Wolfe *et al* JCAP11(2023)039

View the [article online](#) for updates and enhancements.

You may also like

- [CONSTRAINING STELLAR FEEDBACK: SHOCK-IONIZED GAS IN NEARBY STARBURST GALAXIES](#)
Sungryong Hong, Daniela Calzetti, John S. Gallagher et al.
- [Near- and sub-solar-mass naked singularities and black holes from transmutation of white dwarfs](#)
Chandrachur Chakraborty and Sudip Bhattacharyya
- [Delayed Detonation Thermonuclear Supernovae with an Extended Dark Matter Component](#)
Ho-Sang Chan, Ming-chung Chu, Shing-Chi Leung et al.



**JCAP ANNIVERSARY
SPECIAL ISSUE**

Too small to fail: characterizing sub-solar mass black hole mergers with gravitational waves

Noah E. Wolfe, Salvatore Vitale and Colm Talbot

*LIGO Laboratory, Massachusetts Institute of Technology,
185 Albany St, Cambridge, MA 02139, U.S.A.*

*Department of Physics and Kavli Institute for Astrophysics and Space Research,
Massachusetts Institute of Technology,
77 Massachusetts Ave, Cambridge, MA 02139, U.S.A.*

E-mail: noah.wolfe@ligo.org, salvo@mit.edu, colmt@mit.edu

ABSTRACT: The detection of a sub-solar mass black hole could yield dramatic new insights into the nature of dark matter and early-Universe physics, as such objects lack a traditional astrophysical formation mechanism. Gravitational waves allow for the direct measurement of compact object masses during binary mergers, and we expect the gravitational-wave signal from a low-mass coalescence to remain within the LIGO frequency band for thousands of seconds. However, it is unclear whether one can confidently measure the properties of a sub-solar mass compact object and distinguish between a sub-solar mass black hole or other exotic objects. To this end, we perform Bayesian parameter estimation on simulated gravitational-wave signals from sub-solar mass black hole mergers to explore the measurability of their source properties. We find that the LIGO/Virgo detectors during the O4 observing run would be able to confidently identify sub-solar component masses at the threshold of detectability; these events would also be well-localized on the sky and may reveal some information on their binary spin geometry. Further, next-generation detectors such as Cosmic Explorer and the Einstein Telescope will allow for precision measurement of the properties of sub-solar mass mergers and tighter constraints on their compact-object nature.

KEYWORDS: Bayesian reasoning, dark matter simulations, gravitational waves / sources, primordial black holes

ARXIV EPRINT: [2305.19907](https://arxiv.org/abs/2305.19907)

Contents

1	Introduction	1
2	Methods	2
2.1	Gravitational-wave parameter estimation	2
2.2	Simulated signals	3
3	Results	5
3.1	Component masses	5
3.1.1	Current-generation detectors	5
3.1.2	Next-generation detectors	8
3.2	Spins	10
3.2.1	Effective spins in current-generation detectors	10
3.2.2	Component spin magnitudes and tilts in next-generation detectors	13
3.3	Sky localization	15
4	Discussion	16
A	Parameter estimation	18
A.1	Waveform model and heterodyned likelihood	18
A.2	Priors	19
A.3	Nested sampling in bilby	19
B	Marginal posteriors on chirp mass and mass ratio	20
C	Effective spin constraints on spin magnitude	21
C.1	Spin-aligned systems	21
C.2	Spin anti-aligned systems	22
D	Features in correlated mass ratio-effective spin posteriors	24
E	Additional marginal posteriors on spin geometry	27
F	Biased tilt posteriors at high signal-to-noise ratios	28
G	Correlation between spin precession and binary geometry	29

1 Introduction

If some fraction of dark matter is composed of black holes, or gravitationally collapses to form black holes, gravitational waves (GWs) may offer the opportunity to directly probe the nature of dark matter. Recent work has proposed that previous GW signals consistent with stellar-mass black holes can be sourced from black holes with a primordial origin [1–4], however, there is currently no preference for those formation models over astrophysical channels [2, 3, 5–11]. For current- and next-generation ground-based gravitational-wave detectors like Advanced LIGO [12], Advanced Virgo [13], KAGRA [14], Cosmic Explorer [15], and the Einstein Telescope [16], cleaner targets for dark matter searches may be compact object mergers involving a black hole with mass $\lesssim 1 M_\odot$. The signals emitted by these events lie firmly in the frequency range accessible to ground-based detectors¹ and are immediately distinguished from traditional astrophysical formation channels by their mass alone.² While microlensing surveys have placed constraints on the fraction of dark matter composed by $\mathcal{O}(1 M_\odot)$ black holes, the constraints may depend on the assumed mass distribution of the objects as well as the distribution of mass in the halo of the Milky Way [18]. Thus, it remains possible for some dark matter to be found in sub-solar mass black holes. There are two categories of hypothesized sub-solar mass black holes: primordial black holes and dark matter black holes.

Primordial black holes (PBHs) could have formed from the gravitational collapse of over-densities in the early Universe [19, 20]. Such objects have been proposed as a population of cold, collisionless dark matter [21]. The existence and mass distribution of primordial black holes depends strongly on their formation mechanism and underlying density power spectrum. Based on horizon scale considerations, PBHs with a sharply peaked mass distribution near $\mathcal{O}(1 M_\odot)$ could have formed during the radiation-dominated era near the quark-hadron phase transition; the abundance of PBH masses may depend on the background cosmology [22–24], equation of state of the early Universe [25], and the (non-)Gaussianity of the density fluctuations [26–28] (for a review of PBH formation in the radiation-dominated era, see [18]). PBH formation could have been driven by other physics during or beyond the radiation-dominated era (for a review of such scenarios, see [21]). Fluctuations generated by inflation could seed primordial black holes [29–33], which would be directly sensitive to the dynamics of inflation. Other structures, like cosmic loops [34], bubbles of broken symmetry [35], and domain walls [36] could also collide or collapse to form primordial black holes. Additionally, the spectrum of PBH masses would be sensitive to accretion physics [18], and their formation into binaries depends on their clustering dynamics [37–42] and their merger rate which are theoretically uncertain [4, 43].

Alternatively, black holes could form directly from the collapse of particle dark matter in certain dissipative dark matter models [44–52]. The precise mass distribution of such “dark matter black holes” (DBHs) is strongly dependent on the microphysical details of the dark matter particles such as the possible dark matter species, their masses, and interaction cross

¹Roughly ~ 10 – 2000 Hz, although the exact frequency band considered while analyzing data depends on the particular detector and its noise properties at the time of detection. Lower-mass compact object systems inspiral more slowly, and systems with a total mass approaching $1 M_\odot$ may spend thousands of seconds radiating at a few tens or hundreds of Hertz before merger.

²As their mass lies below the Chandrasekhar limit of $\sim 1.4 M_\odot$ [17].

sections. The merger rate of DBHs would also depend on the larger-scale dynamics of dark matter halos. Both the microphysical details and galactic-scale dynamics of dark matter are open questions [53]. Ref. [45] studied a set of atomic dark matter scenarios that lower the Chandrasekhar mass with a heavier dark-proton analog, forming black holes $\lesssim 1 M_\odot$. There, the dark matter microphysics is encoded in the dark matter cooling rates, and halo dynamics are encoded in the fraction of dark matter available to dynamically cool.

Thus, the discovery of a sub-solar mass black hole would probe the nature of dark matter and provide critical constraints on a rich theoretical landscape of physics in the early- and modern-Universe. If such objects can form binaries and merge in the Hubble time, gravitational waves are the most promising method of detection. Even the non-detection of sub-solar mass compact objects is already providing unique constraints on the parameter space of dark matter [54–57]. However, fully realizing this promise hinges on positively identifying a compact object involved in a merger as (1) $\lesssim 1 M_\odot$ in mass and (2) a black hole. This first question has previously been studied for astrophysical, super-solar mass black holes in next-generation detectors [58, for example] as well as sub-solar mass neutron stars [59]. Given a gravitational wave signal from the coalescence of compact objects, the identification of a sub-solar mass black hole could be complicated by “mimickers” such as sub-solar mass neutron stars or boson stars [60]. While such objects themselves would be astrophysically exotic [61–63] and potentially sourced from dark matter [64], it is not yet clear how well current methods of gravitational-wave data analysis will distinguish between sub-solar mass black holes and these alternatives. For example, when analyzing a low-significance sub-solar mass trigger, ref. [65] could not exclude a neutron star origin for the sub-solar mass object.³

In this work, we estimate parameters for a set of simulated signals from sub-solar mass black hole mergers in current- and next-generation gravitational wave detectors to understand the feasibility of identifying sub-solar mass black holes with gravitational-wave signals across a range of binary black hole parameters. First, we inspect the constraints we can achieve on the component black hole masses, to determine if we can confidently identify that a compact object is sub-solar mass in nature at all. Then, we inspect two additional parameter sets for these signals — the spins of the compact objects and the sky location of the binary — both of which may rule out neutron stars in such signals. We conclude with a discussion of our main findings.

2 Methods

2.1 Gravitational-wave parameter estimation

To measure the source properties of gravitational-wave signals, we perform Bayesian parameter estimation. Bayes’ Theorem states that

$$p(\theta|d) = \frac{\mathcal{L}(d|\theta)\pi(\theta)}{\mathcal{Z}(d)} \quad (2.1)$$

where $p(\theta|d)$ is the posterior probability that the data d contains a signal described by source parameters θ , $\mathcal{L}$ is the likelihood of observing the signal given some source parameters, $\pi(\theta)$

³See ref. [66] for an analysis focusing on distinguishing light super-solar mass black holes from neutron stars.

is the prior probability of θ , and $\mathcal{Z}$ is a normalization commonly referred to as the evidence. In the case of gravitational-wave parameter estimation for emission from a quasi-circular black hole binary, θ includes 8 intrinsic parameters of the component black holes (their masses and spin vectors) and 7 extrinsic parameters of the binary (including its luminosity distance and location on the sky), for 15 parameters total. We use the Whittle likelihood approximation in the frequency domain for the residual of the data minus the astrophysical signal.⁴ Here, the goal of our analysis will be to compute the posterior probability of θ for a series of simulated signals. In this work, we use the nested sampling algorithm [68, 69] implemented by `dynesty` [70] to estimate $p(\theta|d)$.

As the nested sampling algorithm iterates, we must evaluate $\mathcal{L}(d|\theta)$ which requires evaluating the waveform model at some proposal θ . Signals from merging binaries with relatively low total mass or mass ratio will be relatively long; for the component masses listed in table 1, signals will be $\mathcal{O}(10^3)$ seconds in length compared to $\mathcal{O}(10^2)$ seconds for the larger total mass, near equal-mass binary neutron star GW170817 [71]. This dramatically increases the number of frequencies at which we need to evaluate proposal waveforms for a given θ , in evaluating $\mathcal{L}$, which is computationally expensive. Instead, we use a “heterodyned” likelihood [72], an approximation also known as “relative binning” [73, 74], which well-approximates $\mathcal{L}$ at far fewer frequencies by expanding it around its value at some fiducial parameters. For simulated signals, we choose these parameters to be the true source parameters. We also forego effects due to the rotation of the Earth, which would cause the antenna pattern to vary in time and increase the computational expense of $\mathcal{L}$. We carry out parameter estimation using `bilby` [75, 76] which implements a heterodyned likelihood as in ref. [77]; for details on the heterodyned likelihood, our priors, and sampler settings, see appendix A.

2.2 Simulated signals

We consider a set of simulated gravitational-wave signals from binary black hole mergers involving both sub-solar and super-solar mass components. In table 1, we show the pairs of detector-frame component masses ($\tilde{m}_1, \tilde{m}_2$) of the signals studied in this work; here, $\tilde{m}_1 \geq \tilde{m}_2$. For each of these pairs, we consider mergers where the more massive component has a dimensionless spin magnitude a_1 of 0.6, 0.8, and 0.9. These spins are chosen to be greater than the observed maximum spin for a neutron star of $a_{\text{NS}} = 0.4$ [78]. The spin of the less massive component, a_2 , is always chosen as zero. At each of these masses and spins, we also consider different orientations of the spin vector of the more massive component, $\vec{S}_1$, characterized by the “tilt” zenith angle θ_1 between $\vec{S}_1$ and the orbital angular momentum $\vec{L}$. When the black hole spins are misaligned with $\vec{L}$ (e.g. $\theta_1 > 0$), all angular momenta in the system precess, driven by inertial frame dragging [79]. Here, we consider each of two cases, $\theta_1 = 0$ and $\theta_1 = \pi/2$. While our simulated black hole binaries feature spinning components, and some precess, in order to test our ability to measure these dynamics in their mergers, we emphasize that the distribution of sub-solar mass black hole spins and orientations is

⁴In other words, we treat the noise as stationary and Gaussian-distributed. Note this may not necessarily be true in next-generation detectors which will have many overlapping signals, but in principle, our results hold as e.g. one could subtract out this astrophysical “noise” [67].

$\tilde{m}_1 [M_\odot]$	$\tilde{m}_2 [M_\odot]$	q	O4 SNRs	XG SNRs	Redshift	$m_1 [M_\odot]$	$m_2 [M_\odot]$
0.5	0.5	1.00	10.0	314.7	0.009	0.495	0.495
			29.8	934.1	0.028	0.486	0.486
0.9	0.9	1.00	7.5	235.5	0.028	0.875	0.875
			16.0	502.4	0.059	0.850	0.850
0.9	0.5	0.56	12.5	392.0	0.009	0.891	0.495
			37.1	1163.4	0.028	0.875	0.486
1.4	0.5	0.36	21.2	665.5	0.009	1.387	0.495
			42.4	1331.6	0.019	1.374	0.491
1.0	0.1	0.10	14.3	447.5	0.009	0.991	0.099
1.4	0.1	0.07	13.7	429.0	0.009	1.387	0.099

Table 1. Detector-frame component masses $\tilde{m}_1, \tilde{m}_2$ and the mass ratio q of the simulated signals studied in this work. At each row $(\tilde{m}_1, \tilde{m}_2)$ listed here, we simulated a signal for each of a_1 chosen from $\{0.6, 0.8, 0.9\}$ and the tilt angle θ_1 chosen as 0 (no precession) or $\pi/2$ (precessing). For each set of intrinsic parameters $(\tilde{m}_1, \tilde{m}_2, a_1, \theta_1)$, we simulate a source in an O4 and XG detector network with each of the SNRs shown by varying the distance. We include representative redshifts and source-frame component masses m_1, m_2 for each combination of detector-frame masses and SNR; the redshifts may change by as much ± 0.001 to keep the SNR constant as the spin magnitude and tilt are varied.

unknown, depending on how such objects are formed, pair into binaries, and evolve through cosmic time.⁵ For each set of intrinsic parameters $(\tilde{m}_1, \tilde{m}_2, a_1, \theta_1)$, we vary the luminosity distance of the source, d_L , to achieve one of the network signal-to-noise ratios (SNRs) listed for each $(\tilde{m}_1, \tilde{m}_2)$ in table 1. We also note that the black hole masses in the detector frame are larger than the source-frame component masses m_1, m_2 due to cosmological redshift by at most 6%. This results in a total of 60 different sources. Our nearest source is at $d_L = 42$ Mpc (a redshift $z \sim 0.009$ assuming the Planck15 cosmology [82]), and our furthest is at $d_L = 270.7$ Mpc ($z \sim 0.059$). Additional source parameters are listed in table 3. We simulate each signal in a current-generation network of LIGO-Hanford, LIGO-Livingston, and Virgo at design sensitivities for their fourth observing run (O4),⁶ and a next-generation (XG) network of one Cosmic Explorer at the current site of LIGO-Hanford and the Einstein Telescope at the site of Virgo.⁷ We use zero-noise realizations of the detector sensitivities; posteriors estimated in zero-noise will be equivalent to the average posterior estimated in many Gaussian-noise realizations [87–90]. For all signals considered, we model the gravitational waveform with the phenomenological spin-precessing frequency-domain model IMRPhenomXP [91]. We forego higher-order gravitational-wave modes in our analysis for computational expediency, which are expected to be measurable during the merger and ringdown of systems with large total mass

⁵For example, in [80] the authors argue that primordial black holes will be mostly non-spinning, while the authors of [81] find that they may mostly have maximal spins.

⁶Using sensitivity curves from [83, 84].

⁷Using sensitivity curves from [85, 86].

or extreme mass ratios.⁸ Thus, our constraints without higher-order modes are conservative upper limits. Our choices of component masses complement the parameter space studied by recent LIGO-Virgo-KAGRA (LVK) sub-solar mass compact object searches as well as that of ref. [55]. The most recent LVK search considered source-frame masses $m_1 \in [0.2, 10] M_\odot$ and $m_2 \in [0.2, 1] M_\odot$ [54]; ref. [55] considered detector-frame masses $\tilde{m}_1 \in [0.1, 100] M_\odot$ plus $\tilde{m}_2 \in [0.01, 1] M_\odot$ for $\tilde{m}_1 > 20 M_\odot$ and $\tilde{m}_2 \in [0.1, 1] M_\odot$ for $\tilde{m}_1 < 20 M_\odot$. As IMRPhenomXP is only valid for mass ratios $q = m_2/m_1 > 1/20$ [93, 94], we choose our minimum $\tilde{m}_2 = 0.1 M_\odot$ consistent with ref. [55] and below the LVK search boundary.

3 Results

3.1 Component masses

3.1.1 Current-generation detectors

Our aim is to determine when we can identify that one or both of the component compact objects has mass $< 1 M_\odot$. In general, one expects the estimation of mass parameters to get better and better as the total mass of a binary decreases and the number of inspiral cycles increase; this is why the chirp mass of binary neutron stars is measured much more precisely than that of binary black holes [95–97] (See also table IV of ref. [98]). In figure 1, we show marginal posteriors on the source-frame mass m_1 for the simulated signals described in section 2.2, sorted by mass ratio and whether or not the system is precessing. Posteriors are also colored by network SNR; as expected, as SNR goes from lower values (cooler tones) to higher values (warmer tones), the width of the posteriors decreases. In all cases, we recover the simulated value (thin black lines) and do so more confidently with decreasing q . This is simply due to the greater distinguishability between m_1 and m_2 at lower mass ratios. We achieve the best measurement of m_1 for the precessing, $q = 0.07$ source with $a_1 = 0.9$ with a 90% credible interval of $1.7 \times 10^{-2} M_\odot$. The worst measurement occurs with the $q = 1$, $a_1 = 0.9$, non-precessing source with an SNR of 7.5, with a 90% credible interval of $0.84 M_\odot$. In general, the network SNR drives the measurement of m_1 , over a_1 , however, we do observe that the width of the credible interval weakly depends on a_1 at the $\sim 10\%$ level for non-precessing sources. In the most extreme example, the 90% credible interval for the non-precessing $q = 1$ source at an SNR of 7.5 decreases from $0.84 M_\odot$ to $0.73 M_\odot$ as a_1 decreases from 0.9 to 0.6 (whereas for the precessing system with the same mass parameters, the credible interval only shrinks from $0.68 M_\odot$ to $0.61 M_\odot$).

We can also quantify the “efficiency” of measuring m_1 , which we evaluate with the ratio between the width of the 90% credible interval and the SNR; we call this ratio α . As for the width of the credible interval, we find that this ratio typically changes at most at the $\sim 10\%$ level as we vary a_1 and keep SNR the same. It varies noticeably as the SNR changes, indicating a nonlinear improvement in our measurement of m_1 at increasing SNR. For example, for non-precessing, $q = 0.36$ sources, α improves from $\sim 1.5 \times 10^{-2} M_\odot$ to $\sim 1.8 \times 10^{-3} M_\odot$ as the SNR roughly doubles from 21.2 to 42.4.

⁸We note that [92] recently showed that higher-order modes can improve the measurement of distance and inclination angle for binary neutron star mergers, comparable to the sub-solar mass black hole mergers studied in this work.




Figure 1. Marginal posterior distributions on the source-frame mass of the heavier black hole, m_1 for the simulated signals injected into an O4-design sensitivity network of LIGO-Hanford, LIGO-Livingston, and Virgo. Results for non-precessing ($\theta_1 = 0$) sources are shown on the left and those for precessing sources ($\theta_1 = \pi/2$) on the right. The posteriors are colored by the network SNR of the signal. The linestyle of the posterior reflects the dimensionless spin magnitude of the more massive black hole, $a_1 = 0.6$ (dotted), 0.8 (dashed), or 0.9 (solid). Posteriors are organized by increasing mass ratio of the source, from top to bottom. Thin black lines rising from the m_1 -axis indicate the true value of m_1 , and a grey line is included at the fiducial mass scale of $1 M_\odot$. We note that these posteriors are not normalized so that they may be visualized together.

Comparing each posterior to the vertical grey line at $1 M_\odot$, we are not able to confidently exclude $m_1 \geq 1 M_\odot$ for sources with $m_1 < 1 M_\odot$ (bottom three rows). Among sources with $q = 0.56$, the weakest exclusion of a super-solar mass object occurs for the non-precessing source with $a_1 = 0.9$, at an SNR of 12.5, where $1 M_\odot$ occurs at the 60.9% percentile of the marginal posterior distribution. At $q = 1$, with $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$ this exclusion is weaker, with $1 M_\odot$ occurring between the 23.3% ($a_1 = 0.9$, non-precessing source with an SNR of 7.5) and 47.0% ($a_1 = 0.9$, precessing source with an SNR of 16) percentiles.

Turning to figure 2, we show marginal posteriors on the source-frame mass m_2 . This figure is organized in the same manner as figure 1. The relationships between the width of the credible intervals, q , and SNR are qualitatively much the same as in m_1 . However, these posteriors are noticeably more narrow than those for m_1 (so much so that for the lowest mass ratios, we provide insets to resolve detail in the histograms). In particular, our best (worst)




Figure 2. Marginal posterior distributions on the source-frame mass of the lighter black hole, m_2 for the simulated signals injected into an O4-design sensitivity network of LIGO-Hanford, LIGO-Livingston, and Virgo. Results for non-precessing ($\theta_1 = 0$) sources are shown on the left and those for precessing ($\theta_1 = \pi/2$) on the right. The posteriors are colored by the network SNR of the signal in an O4-design sensitivity network of LIGO-Hanford, LIGO-Livingston, and Virgo. The linestyle of the posterior reflects the dimensionless spin magnitude of the more massive black hole, $a_1 = 0.6$ (dotted), 0.8 (dashed), or 0.9 (solid). Posteriors are organized by increasing mass ratio of the source, from top to bottom. Thin black lines rising from the m_2 -axis indicate the true value of m_2 , and a grey line is included at the fiducial mass scale of $1 M_\odot$. We note that these posteriors are not normalized so that they may be visualized together. We observe that $m_2 \geq 1 M_\odot$ is excluded for every simulated signal.

measurement occurs for the precessing (non-precessing) source with $q = 0.07, a_1 = 0.8$ ($q = 1, \tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot, a_1 = 0.8$, and network SNR of 7.5), achieving a 90% credible interval of $9.1 \times 10^{-4} M_\odot$ ($0.36 M_\odot$). This improvement follows from the definition of the mass ratio; since q is linear in m_2 but $\propto 1/m_1$, at low mass ratios q is more sensitive to m_2 than m_1 . Importantly, we observe that *every* posterior excludes $m_2 \geq 1 M_\odot$. We also stress that the hard edge observed towards $1 M_\odot$ in the $q = 1$ results are the posteriors railing on the prior constraint that $m_1 \geq m_2$. So, at least for systems equivalent to the injections studied in this work, we could confidently report the discovery of a sub-solar mass compact object at SNRs as low as 7.5. For completeness, we show marginal posteriors on the source-frame chirp mass $\mathcal{M}$ and mass ratio q in the O4 network in appendix B.




Figure 3. Marginal posterior distributions on the source-frame mass of the heavier black hole, m_1 for the simulated signals injected into a network of Cosmic Explorer and the Einstein Telescope. Results for non-precessing ($\theta_1 = 0$) sources are shown on the left and those for precessing sources ($\theta_1 = \pi/2$) on the right. The posteriors are colored by the network SNR of the signal. The linestyle of the posterior reflects the dimensionless spin magnitude of the more massive black hole, $a_1 = 0.6$ (dotted), 0.8 (dashed), or 0.9 (solid). Posteriors are organized by increasing mass ratio of the source, from top to bottom. Thin black lines rising from the m_1 -axis indicate the true value of m_1 , and a grey line is included at the fiducial mass scale of $1 M_\odot$. We note that these posteriors are not normalized so that they may be visualized together.

3.1.2 Next-generation detectors

In figures 3 and 4, we show marginal posteriors on the source-frame component masses m_1 and m_2 , respectively, for signals simulated in the XG network of one Cosmic Explorer and Einstein Telescope. These figures are organized in the same manner as figures 1 and 2 in the previous subsection. For the more massive component, m_1 , we correctly characterize the compact object as solar, sub-, or super-solar in mass for all the signals studied here, except in the case of non-precessing sources with detector-frame masses $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$. There, $1 M_\odot$ occurs at worst at the $\gtrsim 99.9\%$ percentile, for the $a_1 = 0.9$ source.

For the less massive component, m_2 , we correctly characterize the compact object as sub-solar in nature in all cases, like in current-generation detectors. Except for signals with $q = 1$, we are also able to constrain both m_1 and m_2 away from the prior, in particular, without observing any railing on the mass ratio prior. Again, we stress that the hard edge on the $q = 1$ posteriors is a result of the definition $m_1 \geq m_2$. We achieve our best (worst) measurement of




Figure 4. Marginal posterior distributions on the source-frame mass of the lighter black hole, m_2 for the simulated signals injected into a network of Cosmic Explorer and the Einstein Telescope. Results for non-precessing ($\theta_1 = 0$) sources are shown on the left and those for precessing ($\theta_1 = \pi/2$) on the right. The posteriors are colored by the network SNR of the signal. The linestyle of the posterior reflects the dimensionless spin magnitude of the more massive black hole, $a_1 = 0.6$ (dotted), 0.8 (dashed), or 0.9 (solid). Posteriors are organized by increasing mass ratio of the source, from top to bottom. Thin black lines rising from the m_2 -axis indicate the true value of m_2 , and a grey line is included at the fiducial mass scale of $1 M_\odot$. We note that these posteriors are not normalized so that they may be visualized together.

m_1 for the precessing (non-precessing) source with $q = 0.07$, $a_1 = 0.8$, and a network SNR of 476.0 ($q = 1$, $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$, $a_1 = 0.9$, and SNR of 235.5), that has a 90% credible of $3.5 \times 10^{-4} M_\odot$ ($0.11 M_\odot$). For m_2 , we achieve our best and worst measurements with the same sources which have credible intervals of $1.7 \times 10^{-5} M_\odot$ ($q = 0.07$, $a_1 = 0.8$, and SNR of 476.0) and $0.10 M_\odot$ ($q = 1$, $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$, $a_1 = 0.9$, and SNR of 235.5), respectively. This improvement may be driven by the same nonlinear improvement in the measurement efficiency α seen in O4 detectors, which continues in the XG network. Previously, we saw that for the non-precessing, $q = 0.36$ source, α for m_1 is quartered as the SNR doubles; for the same source in a XG network, α decreases from $\sim 1.7 \times 10^{-5} M_\odot$ to $\sim 5.8 \times 10^{-6} M_\odot$ as the SNR increases from ~ 666 to ~ 1333 . We see a similar pattern in the measurement efficiency of m_2 . Overall, we observe an order-of-magnitude improvement in our ability to measure the component masses in XG detectors compared to the O4 network. We show marginal posteriors on the source-frame chirp mass and mass ratio in the XG network in appendix B.

3.2 Spins

3.2.1 Effective spins in current-generation detectors

The spin of a compact object may distinguish it from a neutron star. The maximum observed neutron star spin is that of a pulsar in a binary, with $a_{\text{NS}} = 0.4$ [78]. Whether this value is the theoretical maximal neutron star spin depends on the unknown nuclear equation of state; however, we adopt it as a fiducial value, above which a compact object is inconsistent with a neutron star description. Even at SNRs as high as ~ 40 , we do not expect to measure the component spin magnitudes or spin tilt angles well on their own [96, 99, 100]. However, the leading-order contributions of spins to the gravitational inspiral are measurable and can be parameterized with the effective spin χ_{eff} and effective spin precession χ_p .

The effective spin is the mass-weighted projection of the component black hole spins along the orbital angular momentum of the system [101–103],

$$\chi_{\text{eff}} = \frac{m_1 a_1 \cos \theta_1 + m_2 a_2 \cos \theta_2}{m_1 + m_2} = \frac{a_1 \cos \theta_1 + q a_2 \cos \theta_2}{1 + q}. \quad (3.1)$$

The effective spin χ_{eff} is usually measured better than either component spins, even though it is known to be partially degenerate with the mass ratio for inspiral-dominated (i.e. low mass) systems [104, 105]. When χ_{eff} is zero, the black holes may be non-spinning, or spinning entirely in the plane of the orbit. When χ_{eff} is +1 (−1), the black holes are spinning parallel to and in the direction of (opposite the direction of) the orbital angular momentum. In figure 5, we show marginal posteriors on χ_{eff} for a subset of our runs, sorted by a_1 and colored by the mass ratio. On the left, we show results for non-precessing signals, and on the right, results for precessing signals. For clarity, we only include the quieter signals at each pair of masses shown in table 1. The marginal posteriors on χ_{eff} look nearly identical at different SNRs with all else held equal (see appendix E for the marginal posteriors for the louder events). Instead, the mass ratio drives the measurement of spin parameters as observed in e.g. [99]. For non-precessing signals, we confidently recover χ_{eff} at all mass ratios and values of a_1 . We find our best result for the source with $q = 0.07$, $a_1 = 0.6$, with a 90% credible interval of 0.029, and the worst result for the source with $q = 1$, $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$, $a_1 = 0.9$ with a credible interval of 0.72. For precessing signals, we find that $\chi_{\text{eff}} = 0$ is confidently recovered for all sources, with a better measurement of χ_{eff} at lower q . Quantitatively, the posterior distributions for precessing sources are comparable to their non-precessing counterparts; our best (worst) recovery of χ_{eff} occurs for the precessing source with $q = 0.07$, $a_1 = 0.8$ ($q = 1$, $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$, $a_1 = 0.9$), yielding a credible interval of 0.025 (0.18).

We also observe that, for some precessing sources with lower spin, we would preferentially report the incorrect χ_{eff} . In particular, the marginal posterior on χ_{eff} for the precessing, $q = 0.56$ (green-blue color) signal at all spin magnitudes is consistent with zero, but peaks away from zero. Worse, at $a_1 = 0.6$, the $q = 0.1$ signal exhibits bimodality and the $q = 0.07$ result is inconsistent with $\chi_{\text{eff}} = 0$. However, these features are the result of posterior support for q slightly away from its true value, combined with large uncertainties in the spin magnitudes and tilts. We detail this behavior in appendix D.

Although it would be easier to directly measure the spin magnitudes as consistent or inconsistent with the maximal neutron star spin, we can indirectly constrain a_2 through χ_{eff}




Figure 5. Prior distribution $\pi(\chi_{\text{eff}})$ (top row) and marginal posterior distributions on the effective spin χ_{eff} for the quietest signals for a given set of intrinsic parameters $(\tilde{m}_1, \tilde{m}_2, a_1, \theta_1)$, injected into an O4-design sensitivity network of LIGO-Hanford, LIGO-Livingston, and Virgo (remaining rows). These are sorted, from top to bottom, by increasing dimensionless spin magnitude on the heavier black hole, a_1 . We color the posterior distributions by the true mass ratio of the source. In the left column, we show results for non-precessing ($\theta_1 = 0$) sources, where the truth is indicated with a dashed line colored according to the true q . In the right column, we show results for precessing ($\theta_1 = \pi/2$) sources, where the truth is indicated with a dashed black line. We also include two critical values of χ_{eff} above which a_2 is inconsistent with a neutron star spin, as derived in appendix C. These constraints are constructed assuming the spins are aligned (solid grey) and anti-aligned (dashed grey). We observe that χ_{eff} is best constrained for non-precessing systems, and at lower mass ratios.

if we assume some prior knowledge of the geometry of the system (a_2 being the spin of the sub-solar mass object in all of our simulated signals). In appendix C we find critical values of χ_{eff} for which $a_2 > a_{\text{NS}}$, assuming that we know the alignment of the black hole spins. We include these constraints as grey lines in both columns of figure 5. These constraints are constructed with conservative assumptions on the tilts of the black holes, the spin of the more massive black hole, and the mass ratio of the system. In particular, if $|\chi_{\text{eff}}|$ is larger than the spin-aligned constraint (solid grey), then the spin of the lighter object (which is a sub-solar mass black hole in all of our simulated sources) *must* be larger than the maximal neutron star spin. For non-precessing signals, we can exclude $|\chi_{\text{eff}}|$ below both constraints at the very lowest mass ratios. Otherwise, we would need to infer that the spins are anti-aligned to exclude spins consistent with a neutron star. For precessing signals, our inference of χ_{eff} provides no constraint on a_2 . This is due to the lossy nature of the χ_{eff} parameterization, as we are trying to recover zero which is degenerate with precessing signals (as we have simulated) as well as systems without any spin. Since the tilt angles are not well measured, it is difficult to break this degeneracy. These constraints may be relevant depending on the binary formation mechanism of sub-solar mass black holes and their associated spin distributions.




Figure 6. In the top row, we show the prior distribution $\pi(\chi_p)$ (black) and conditional priors $\pi(\chi_p|q)$ from [106] at the mass ratios listed in table 1. In the remaining rows, we show marginal posterior distributions on χ_p for the quietest signals for a given set of intrinsic parameters $(\tilde{m}_1, \tilde{m}_2, a_1, \theta_1)$, injected into an O4-design sensitivity network of LIGO-Hanford, LIGO-Livingston, and Virgo. These are sorted, from top to bottom, by increasing dimensionless spin magnitude on the heavier black hole, a_1 . We color the posterior distributions by the true mass ratio of the source. In the left column, we show results for non-precessing ($\theta_1 = 0$) sources, and in the right column, we show results for precessing ($\theta_1 = \pi/2$) sources. The truth is indicated with a dashed black line. We observe that χ_p is best constrained for precessing systems, and at lower mass ratios.

The effective spin precession is the average magnitude of the leading-order contribution to the precession of the orbital angular momentum of the binary, and can be expressed to leading order as [107],

$$\chi_p = \max \left(a_1 \sin \theta_1, q \frac{4q+3}{4+3q} a_2 \sin \theta_2 \right). \quad (3.2)$$

While this quantity does not capture all of the relativistic dynamics of precessing binaries (see e.g. [108]), it best captures precession in systems like those we have simulated, where a single component drives the orbital precession [107]. In figure 6, we plot marginal posteriors on χ_p for the same subset of runs shown in figure 5, organized in the same manner. We also include the prior $\pi(\chi_p)$ and conditional priors $\pi(\chi_p|q)$, derived in [106].

First, we note that we make a weak (though incorrect) measurement of χ_p for non-precessing systems at the most extreme mass ratio we studied, $q = 0.07$. For each non-precessing $q = 0.07$ source we recovered a posterior without the same “tail” shown in the priors on χ_p in figure 6. Further, for these sources, the 90% credible interval on the mass ratio is ~ 0.01 , with the true value at the $\sim 65\%$ percentile; i.e. most of the marginal posterior in q is in smaller mass ratios. Referring to the conditional priors $\pi(\chi_p|q)$, we see that lower q tends to increase support near χ_p of zero, but instead, we measure $\chi_p \sim 0.1$. Thus, the uncertainty on our estimate of q for these sources cannot explain the bias on our measurement




Figure 7. Marginal posterior distributions on the dimensionless spin magnitude of the heavier black hole a_1 , for the quietest signals for a given set of intrinsic parameters $(\tilde{m}_1, \tilde{m}_2, a_1, \theta_1)$, injected into a network of Cosmic Explorer and the Einstein Telescope. These are sorted, from top to bottom, by increasing values of a_1 . We color the posterior distributions by the true mass ratio of the source. In the left column, we show results for non-precessing ($\theta_1 = 0$) sources, and in the right column, we show results for precessing ($\theta_1 = \pi/2$) sources. The truth is indicated with a dashed black line. For reference, we include a solid grey line at the maximal neutron star spin, $a_{\text{NS}} = 0.4$. We observe that a_1 is measured away from these limits for non-precessing systems and at lower mass ratios.

of χ_p . For these systems, we may instead be dominated by the uncertainty in the tilt angle, and explaining some of the observed spin in the system with slightly misaligned compact object spins. For the remaining simulated signals, we can easily compare the posteriors on χ_p to $\pi(\chi_p)$ and $\pi(\chi_p|q)$, and we find that we only measure χ_p for precessing systems at the lowest mass ratios and the highest ($a_1 = 0.8, 0.9$) spins. This can be explained by remembering that slice of the binary parameter space is where precession will modify the inspiral of the binary the most.

3.2.2 Component spin magnitudes and tilts in next-generation detectors

In the previous section, we showed that effective spins may be measurable and provide some constraints on the nature of sub-solar mass compact objects in O4 gravitational-wave detectors. Unfortunately, individual spins cannot be well-constrained at low SNRs. Of the simulated signals studied in this work, the lowest network SNR achieved in the XG network is 236.6. These signals are orders of magnitude louder than any gravitational-wave event observed to date and enable us to directly measure the magnitude and angles of binary black hole spin vectors.

In figure 7, we show marginal posteriors on the spin magnitude of the more massive component, a_1 , for a subset of our signals injected into XG detectors, showing results only for the quieter signals. The posteriors are colored by the mass ratio of the source, and organized by the true value of a_1 , increasing from top to bottom. In all cases, we make a measurement consistent with the true value of a_1 . Comparing to the maximal neutron star spin $a_{\text{NS}} = 0.4$ (solid grey line), we observe that the best exclusion of such spins occurs for non-precessing




Figure 8. Marginal posterior distributions on the tilt angle between the spin vector of the heavier black hole and the orbital angular momentum, θ_1 , for the quietest signals for a given set of intrinsic parameters $(\tilde{m}_1, \tilde{m}_2, a_1, \theta_1)$, injected into a network of Cosmic Explorer and the Einstein Telescope. In the top row, we show the prior $\pi(\theta_1)$. For brevity, we only include results for sources with $q = 0.10$ (middle row) and $q = 1.0$ (bottom row). In the left column, we show results for non-precessing ($\theta_1 = 0$) sources, and in the right column, we show results for precessing ($\theta_1 = \pi/2$) sources. The truth is indicated with a dashed black line. We color the posterior distributions by the network SNR, and their linestyles reflect the true spin magnitude of the heavier black hole, a_1 . We observe that θ_1 is well-measured for non-precessing systems, and at low mass ratios. In addition, we note that the lack of posterior support at $\theta_1 = \pi/2$ for precessing, $q = 0.1$ signals reflects the prior effect demonstrated in appendix F.

systems. In addition, the constraints on a_1 improve with increasing mass ratio. Among non-precessing systems, we measure a_1 the best (worst) for the source with $q = 0.07$, $a_1 = 0.9$ ($q = 1$, $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$, $a_1 = 0.6$), which yields a 90% credible interval of 7.7×10^{-3} (0.28). Among precessing systems, we similarly measure a_1 the best (worst) for the $q = 0.07$, $a_1 = 0.8$ ($q = 1$, $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$, $a_1 = 0.8$) source, which yields a 90% credible interval of 8.2×10^{-4} (0.78). For systems approaching equal masses, we may not be able to exclude a_1 consistent with a neutron star if the system is highly precessing with a fully in-plane spin.

In figure 8, we show marginal posteriors on θ_1 for simulated signals with mass ratios $q = 0.1$ and $q = 1$ (only the $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$ runs, for clarity); posteriors at all q can be found in appendix E. These are colored by network SNR, with a linestyle denoting the value of a_1 for each system. Our analysis employed isotropic priors on θ_1 , shown in the top rows of both figure 8. Comparing the marginal posteriors shown to the priors on θ_1 , we find that we can make a measurement of the tilt angle for non-precessing systems, and for precessing systems up until $q = 1$ (for those systems, we make a measurement that excludes a totally aligned or anti-aligned spin with respect to the orbital angular momentum). For example, for the non-precessing, $q = 0.1$, $a_1 = 0.9$ source, we find a 90% credible interval of 0.01 radians, and for the non-precessing $q = 1$, $\tilde{m}_1 = \tilde{m}_2 = 0.9 M_\odot$, $a_1 = 0.9$ source with network SNR of 7.5, we find a credible interval of 0.26 radians. For our precessing sources, the measurement of θ_1 improves at $q = 0.10$ while dramatically broadening at $q = 1.00$; for example, the signal

with $q = 0.1$, $a_1 = 0.9$ yields a credible interval of 5.7×10^{-3} radians versus an interval of 0.86 radians for the $q = 1$, $a_1 = 0.9$ source with network SNR of 7.5. Finally, we note a peculiarity in the posteriors on θ_1 at $q = 0.01$; there, the true value is explicitly excluded, with the distribution peaking to either side of $\pi/2$. This is a non-physical effect reflecting a breakdown in the numerical approximation to the prior on the spin components aligned with the orbital angular momentum. In appendix F we demonstrate this prior effect by re-analyzing one simulated signal in the XG network with uniform priors on the tilt angles and magnitudes. Sampling directly in the spin magnitudes and tilts, we recover a posterior with a similar width as shown in figure 8 which peaks at the true θ_1 .

While there is no direct constraint on the value of θ_1 for a sub-solar mass black hole mimicker, better measurement of θ_1 improves our measurement of effective spin parameters like χ_{eff} and χ_p , which can indirectly constrain the nature of a compact object as discussed for signals in current-generation detectors. In appendix G, we observe a correlation between χ_p and the inclination angle of the system relative to the line of sight which also improves our measurement of binary spin physics.

3.3 Sky localization

Binary neutron star mergers and some neutron star-black hole mergers can produce an electromagnetic (EM) signal counterpart to gravitational-wave emission [114]. While the lack of an EM-counterpart does not strictly rule out a neutron star component, it can reduce the allowed binary parameter space. In particular, a neutron star in a low-mass ratio merger is expected to be tidally disrupted [115], and we expect to measure q particularly well for the low-mass ratio systems in this work (cf. section 3.1). So, if a source is well-localized and no counterpart is observed, a gravitational wave event may become easier to explain with a sub-solar mass black hole component.

In figure 9, we show the area on the sky to which each of our simulated signals can be localized as a function of their network SNRs. This localization area is the solid angle subtended by the 90% confidence contour on the marginal joint posterior for the right ascension and declination, computed using `ligo.skymap` [116]. For comparison, we include two reference scales: the area enclosed by the 90% confidence contour for the binary neutron star merger GW170817 [71] at a network SNR of ~ 32 (blue star), and the field of view of current- and next-generation optical, infrared, and microwave telescopes that could contribute to the electromagnetic follow-up of a sub-solar mass compact object merger (grey, various linestyles). We find that sub-solar mass compact object mergers achieve a higher sky localization “efficiency” than GW170817, in the sense that at a similar network SNR it is possible to localize these sub-solar mass mergers to a region on the sky that is an order of magnitude smaller (similar to the measurement efficiency α introduced in section 3.1). This follows from the fact that longer-duration signals are easier to localize and that all of our simulated systems (with either a lower total mass or mass ratio) will merge much more slowly than a binary neutron star merger like GW170817. In addition, we observe that the localization for every event, except those at the edge of detectability (at O4 network SNRs of 7.5), is within the field of view for at least one telescope, and most are within the field of view for multiple. In both scenarios (O4 and XG), we see that a compact object merger involving a




Figure 9. Localization areas i.e. 90% confidence regions as a function of network SNR for all simulated signals studied in this work. Signals injected into an O4-design sensitivity network of LIGO-Hanford, LIGO-Livingston, and Virgo are shown in purple. Those injected into a network of Cosmic Explorer and the Einstein Telescope are shown in gold. Grey lines indicate the field of view for electromagnetic follow-up observatories, including the Zwicky Transient Facility [109, ZTF; dash-dot-dotted], Vera C. Rubin Observatory [110, dotted], next-generation Very Large Array [111, ngVLA; dash-dotted], Dark Energy Camera [112, DECam; wide-spaced dots], and the Thirty Meter Telescope [113, TMT; dashed].

sub-solar mass component will be very well-localized compared to our electromagnetic follow-up capabilities, reducing the chance of missing an associated electromagnetic counterpart if something other than black holes is involved in the merger. We also note that our analysis did not include higher-order effects like the time-evolution of the detector antenna pattern due to the rotation of the Earth or the finite size of the detector relative to the wave; however, we expect these to further improve the localization of long-duration signals [117].

4 Discussion

Here, we have studied the measurability of the masses, spins, and sky location for a set of spinning, precessing and non-precessing, quasi-circular binary black hole mergers involving at least one sub-solar mass component. Using a spin-precessing waveform and heterodyned likelihood, we performed parameter estimation on these signals injected into a network of LIGO-Hanford, LIGO-Livingston, and Virgo at O4-design sensitivity, and a network of next-generation detectors, Cosmic Explorer and the Einstein Telescope.

We found that the long duration (~ 1000 's of seconds) of these signals engendered precise measurement of the black hole masses. We confidently identify a sub-solar mass compact object in all of the signals we studied, *at current design sensitivities*, down to network signal-to-noise ratios at the threshold of detectability. In particular, we confidently excluded the

least massive component as being $\geq 1 M_\odot$ in mass even for an equal-mass $0.9\text{-}0.9M_\odot$ merger with an SNR of 7.5, corresponding to a luminosity distance of ~ 125 Mpc, or roughly three times the distance of the binary neutron star merger GW170817. We observed that the SNR drives the measurement of the component masses over the spins, and noted an increasing, nonlinear relationship between the width of the 90% credible interval on m_1 or m_2 and the SNR. Driven by this improvement in measurement efficiency at high SNRs, we found that next-generation detectors will enable exquisite measurements of the source-frame component masses towards the $10^{-5} M_\odot$ level, enabling confident measurement of the compact objects we studied as super-, sub-, or solar in mass.

We then looked at the spins. In current generation detectors, we do not expect SNRs high enough to often confidently measure the component spin magnitudes nor their tilt angles with respect to the orbital angular momentum. However, we found that the leading-order effective spin parameters χ_{eff} and χ_p can be well-measured for non-precessing and precessing sources, respectively, in the O4 network. With strong assumptions on the binary geometry, we found that for non-precessing signals with the lowest mass ratios ($q = 0.07, 0.10$), we could use χ_{eff} to confidently exclude spins consistent with the fastest spinning neutron stars. In next-generation detectors, signals will be orders-of-magnitude louder at or beyond distances comparable to GW170817. This enabled us to directly measure the spin magnitude and tilt of the heavier component, and exclude spins consistent with theoretical maximal neutron star spins for non-precessing systems, and precessing systems at low mass ratios.

Binary neutron star and neutron star-black hole mergers may produce an electromagnetic counterpart; if none is found, especially at low mass ratios, that may rule out a neutron star component in a binary merger. Again enabled by the long duration of these signals, we found that nearly all of the sub-solar mass sources we studied in the three-detector, O4-design sensitivity network would be localized within the field of view of at least one current- or next-generation electromagnetic follow-up instrument. We noted that many of these sources are more “efficiently” localized than GW170817, achieving smaller 90% confidence regions on the sky at similar SNRs. All of the sources in the two-detector, next-generation gravitational wave network were localized within the field of view of four current- and next-generation electromagnetic follow-up instruments.

At least among the sources studied here, current-gravitational wave detectors would be able to confidently report the discovery of a sub-solar mass component of a compact object merger plus some measurement of the binary spin geometry and sky location, enabling unique constraints on the properties of dark matter and the physics of the early Universe. However, the “smoking gun” that characterizes a compact object is tidal deformability, which measures how the material of a compact object responds to the gravitational field of its companion. At present, there are no waveform models including tidal deformability that are calibrated for the low mass ratios, or small component masses studied in this work. Recently, [118] performed the first-ever numerical relativity simulation of a neutron star-sub-solar mass black hole merger. They found significant dephasing between the waveform predicted by numerical relativity and phenomenological or surrogate waveform predictions for the same merger. Once waveform models involving tides are developed for sub-solar mass compact object mergers, future work like ours could explore the measurability of those tides.

Acknowledgments

We thank Katerina Chatziioannou, Tim Detrich, and Carl Johan-Haster for enlightening discussions around the measurement and modeling of tidal deformability in neutron star/sub-solar mass black hole mergers, Geraint Pratten for assistance in understanding the **IMRPhenomX** family of waveforms, and Leo Singer for guidance on estimating sky localization regions. We additionally thank Divya Singh, Ester Ruiz-Morales, Aditya Vijaykumar, and our anonymous referee for their helpful comments on this manuscript. This work was inspired by the 2020 “Workshop on Gravitational Wave Constraints on Dark Compact Objects”, which was supported by a Gordon and Betty Moore Foundation Fundamental Physics Innovation Convening Award and the American Physical Society. 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 are grateful for computing resources provided by the California Institute of Technology and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. SV is partially supported by NSF through the award PHY-2045740. NEW is supported by the La Gattuta Physics Fund and the Henry W. Kendall (1955) Fellowship Fund. CT is supported by an MKI Kavli Fellowship.

Software: bilby v2 [75, 76], astropy v5.1 [119], pesummary v0.13.10 [120], healpy v1.16.1 [121, 122], ligo.skymap v1.0.3 [116], corner.py v2.2.1 [123], numpy v1.22.3 [124], scipy v1.8.1 [125], pandas v1.4.2 [126, 127], dynesty v2.1.0 [70], LALSimulation v4.0.2 [128, 129]

A Parameter estimation

A.1 Waveform model and heterodyned likelihood

We use the Whittle likelihood approximation in the frequency domain for the residual noise [130],

$$\mathcal{L}(d|\theta) = \prod_{i,j} \frac{1}{2\pi S_{ij}} \exp\left(-\frac{4}{T} \frac{|d_{ij} - h(\theta)_{ij}|^2}{S_{ij}}\right), \quad (\text{A.1})$$

where the products run over the gravitational wave detectors in a network and the frequency bins of the data, d_{ij} is the data, $h(\theta)_{ij}$ is the strain of the astrophysical signal, S_{ij} is the power spectral density, and T is the duration of the data. We inject simulated signals into zero noise realizations of the detectors, so the power spectral densities are (theoretical) design sensitivity curves. The strain is evaluated with the frequency-domain waveform model **IMRPhenomXP** which includes spin precession effects [91]. We choose a similar description of spin-precession as used in **IMRPhenomPv2** [131] by choosing the flag **PhenomXPrecVersion** = 104 (see table III, appendix F of [91]).

In practice, taking the log-likelihood is commonly expressed in terms of a complex-valued inner product between the data and the waveform strain, weighted by the power spectral density (see appendix B of [132]). Evaluating this inner product over $\mathcal{O}(10^6)$ (as we have for signals $\mathcal{O}(10^3)$ seconds in duration) is computationally expensive. A heterodyned likelihood [72] (an approximation also known as relative binning [73]) approximates equation (A.1)

by expanding the inner product between the data and the strain about its value at some fiducial values of θ , at a few frequencies whose spacing is chosen to minimize the change in waveform phase between frequencies. The `bilby` implementation of a heterodyned likelihood is described in ref. [77]. Following the description of [73], the heterodyned likelihood relies on hyperparameters ϵ and χ which control the tolerance for this inter-frequency dephasing (see equations 9 and 10 of [73]). In this work, we choose $\epsilon = 0.025$ and $\chi = 1$, which translates to an evaluation of the likelihood at only 1242 frequencies. For each simulated signal we study, we choose the fiducial values to be the true values chosen for each source.

A.2 Priors

Binary black hole parameter definitions follow those of table E1 in [76]. Conceptually, we adopt the following priors:

- *Component masses.* We adopt uniform priors on the detector-frame masses $\tilde{m}_1, \tilde{m}_2$; under these assumptions, we sample in the mass ratio q and detector-frame chirp mass $\tilde{\mathcal{M}}$ as these control the leading-order contributions to the gravitational wave strain. We additionally constrain these priors by requiring $\tilde{m}_1, \tilde{m}_2 \leq 2 M_\odot$ in all analyses. In principle, this allows us to recover super-solar component masses while reducing the volume of the parameter space the sampling procedure may explore.
- *Luminosity distance.* Uniform prior on the coalescence time in the source frame and the enclosed comoving volume, assuming the Planck15 cosmology [82].
- *Spins.* We assume isotropic priors on the spin vectors of the component black holes, parameterized in terms of the components of the spins aligned with the orbital angular momentum $\vec{L}$, $\chi_{1,2}$, the components pointing in the orbital plane $\chi_{1,2}^\perp$, the azimuthal angle (taking $\vec{L}$ to point along the z -direction) between the spins, ϕ_{12} , and the azimuthal angle between $\vec{L}$ and the total angular momentum.
- *Sky location and orientation.* We assume isotropic priors for the orientation and location of the binary on the sky, parameterized by the time of coalescence, azimuth, and zenith as defined at LIGO-Hanford. Since the time at a given detector is what is measured in practice, this parameterization is more efficient to sample in than directly sampling in the right ascension, declination, and time of coalescence as measured at the center of the Earth.

In all of our analyses, we choose priors that encompass the resultant posteriors.

A.3 Nested sampling in bilby

The strategy of nested sampling is to estimate the Bayesian evidence $\mathcal{Z}$ as the integral of the likelihood surface over the prior mass X ,

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X) dX \approx \sum_i w_i \mathcal{L}_i \quad (\text{A.2})$$

where i enumerates likelihood contours of value $\mathcal{L}_i$ and the weights w_i are the prior mass enclosed by those contours. The likelihood contours are chosen to enclose increasingly narrow

Sampler Argument	Value
<code>nlive</code>	500
<code>walks</code>	100
<code>naccept</code>	60
<code>sample</code>	‘acceptance-walk’
<code>proposals</code>	[‘diff’]
<code>bound</code>	‘live’

Table 2. Sampler arguments used in our analysis of simulated signals as defined for the `bilby` implementation of `dynesty` and Differential Evolution-MCMC. We note that a few runs use `nlive` = 2000 to achieve convergence.

regions of the total prior mass; this sorting allows us to estimate $\mathcal{Z}$ without explicitly referring to the complicated, high-dimensional geometry of the likelihood surface [69]. While evaluating the likelihood is a well-defined computation using equation (A.1) and given a model of the gravitational waveform, determining the prior mass enclosed by some likelihood iso-contour is non-trivial. We can estimate w_i with uncertainty e.g. with Markov Chain Monte Carlo (MCMC) chains sampling within some bounds of the prior mass. Here, we use the nested sampling algorithm implemented by `dynesty` [70] and employ Differential Evolution-MCMC (DE-MC), as implemented in `bilby`, to estimate w_i [133]. In table 2, we record the sampler settings provided to `bilby` when we conduct parameter estimation with nested sampling for our simulated signals.

B Marginal posteriors on chirp mass and mass ratio

In figures 10 and 11 we plot marginal posteriors on the source-frame chirp mass $\mathcal{M}$ and mass ratio q , respectively, for the signals simulated in the O4-design sensitivity network. In figures 12 and 13 we similarly plot marginal posteriors on $\mathcal{M}$ and q for signals simulated in the XG network. All of these figures are organized in the same manner as the plots of the marginal posteriors on the source-frame component masses in section 3.1. We note that in figure 12 the marginal posteriors on $\mathcal{M}$ for the louder set of non-precessing $q = 0.56$ signals plus all non-precessing, $q = 0.36$ signals do not recover the true values of the chirp mass, peaking away from these value by $\lesssim 10^{-2} M_\odot$. We emphasize that this is an artifact of marginalizing a 15-dimensional posterior distribution, and not a meaningful bias, as the value of the log-likelihood at the true parameters versus the parameters with highest posterior probability differs by roughly unity. The apparent offset in the source-frame masses is due to the fact that the luminosity distance posterior (needed to convert detector-frame mass to source-frame mass) only includes the true value at the edge of its tail, see figure 14. In turn, the exact shape of the distance posterior is driven by the interplay between the correlation of some of the binary parameters (specifically: distance, inclination angle and χ_p , see appendix G below) and their non-trivial priors. As shown in figure 6, for a system with small mass ratio, the χ_p prior is rather peaked away from zero, but due to correlations, that pushes θ_{JN} to smaller values, which in turn moves the distance to higher values.

Parameter	Units	Value
Intrinsic Parameters		
m_1, m_2	$M_\odot$	See table 1
a_1, a_2	—	a_1 is one of $\{0.6, 0.8, 0.9\}$ $a_2 = 0$
θ_1, θ_2	radians	θ_1 is one of $\{0, \pi/2\}$ $\theta_2 = 0$
ϕ_{12}	radians	1.7
ϕ_{jl}	radians	0.3
Extrinsic Parameters		
d_L	Mpc	Changed to achieve the SNRs in table 1.
θ_{JN}	radians	0.4
Right ascension	radians	1.375
Declination	radians	−1.2108
ψ	radians	2.659
t_c	GPS Time	1126259642.413
ϕ_c	seconds	1.3

Table 3. Complete set of binary black hole source parameters for the simulated signals studied in this work, including the intrinsic parameters of the black holes and the extrinsic parameters of the binary’s location and orientation. Parameters of the angular momenta are evaluated at a reference frequency of 20 Hz. Parameter definitions follow table E1 of [76].

C Effective spin constraints on spin magnitude

C.1 Spin-aligned systems

If we make assumptions about the spin geometry of the binary black hole system, we can constrain the spin of the lighter object a_2 under the most conservative assumptions about the spin of the heavier object, a_1 . First, we consider two black hole spins both aligned with the orbital angular momentum, or both in the opposite direction of the orbital angular momentum; here, we refer to both cases as “spin-aligned”. In these cases,

$$|\chi_{\text{eff}}| = \frac{a_1 + qa_2}{1 + q}. \quad (\text{C.1})$$

Rearranging in terms of a_2 ,

$$a_2 = \frac{|\chi_{\text{eff}}|(1 + q) - a_1}{q}, \quad (\text{C.2})$$

and notice that $a_1 \leq 1$, so,

$$a_2 \geq \frac{|\chi_{\text{eff}}|(1 + q) - 1}{q}. \quad (\text{C.3})$$




Figure 10. Marginal posterior distributions on the source-frame chirp mass $\mathcal{M}$ for the simulated signals injected into an O4-design sensitivity network of LIGO-Hanford, LIGO-Livingston, and Virgo. Results for non-precessing ($\theta_1 = 0$) sources are shown on the left and those for precessing sources ($\theta_1 = \pi/2$) on the right. The posteriors are colored by the network SNR of the signal. The linestyle of the posterior reflects the dimensionless spin magnitude of the more massive black hole, $a_1 = 0.6$ (dotted), 0.8 (dashed), or 0.9 (solid). Posteriors are organized by increasing mass ratio of the source, from top to bottom. Thin black lines rising from the $\mathcal{M}$ -axis indicate the true value of $\mathcal{M}$, and a grey line is included at the fiducial mass scale of $1 M_\odot$. We note that these posteriors are not normalized so that they may be visualized together.

Therefore, $a_2 > a_{\text{NS}}$ when the right-hand side of equation (C.3) is greater than a_{NS} . This is fulfilled for

$$|\chi_{\text{eff}}| > \frac{a_{\text{NS}}q + 1}{1 + q} \quad (\text{C.4})$$

which we recognize as $|\chi_{\text{eff}}|$ for spin-aligned black holes with $a_1 = 1, a_2 = a_{\text{NS}}$. For $a_{\text{NS}} = 0.4$ [78], this critical value of χ_{eff} is maximized at $q = 1$ and we find that $\chi_{\text{eff,crit}} = 0.7$. So, for spin-aligned systems, $|\chi_{\text{eff}}| > 0.7$ implies that a_2 *must* be larger than the maximal neutron star spin, assuming the spins are aligned.

C.2 Spin anti-aligned systems

We can construct a similar constraint in systems where one black hole spin is aligned with the orbital angular momentum and one points in the opposite direction. Here, we take




Figure 11. Marginal posterior distributions on the mass ratio q for the simulated signals injected into an O4-design sensitivity network of LIGO-Hanford, LIGO-Livingston, and Virgo. Results for non-precessing ($\theta_1 = 0$) sources are shown on the left and those for precessing sources ($\theta_1 = \pi/2$) on the right. The posteriors are colored by the network SNR of the signal. The linestyle of the posterior reflects the dimensionless spin magnitude of the more massive black hole, $a_1 = 0.6$ (dotted), 0.8 (dashed), or 0.9 (solid). Posteriors are organized by increasing mass ratio of the source, from top to bottom. Dashed black lines indicate the true value of q . We note that these posteriors are not normalized so that they may be visualized together.

the more massive object to spin in the direction of the angular momentum and the less massive object to spin opposite it (the constraint we construct turns out to be the same if these are switched). Then,

$$\chi_{\text{eff}} = \frac{a_1 - qa_2}{1 + q}, \quad (\text{C.5})$$

and so,

$$a_2 = \frac{a_1 - \chi_{\text{eff}}(1 + q)}{q} \geq -\frac{\chi_{\text{eff}}(1 + q)}{q} \quad (\text{C.6})$$

as $a_1 \geq 0$. Notice, for $a_1 = 0$ and with the spin of the lighter object pointing opposite the orbital angular momentum, $\chi_{\text{eff}} < 0$. Thus,

$$a_2 \geq \frac{|\chi_{\text{eff}}|(1 + q)}{q}, \quad (\text{C.7})$$




Figure 12. Marginal posterior distributions on the source-frame chirp mass $\mathcal{M}$ for the simulated signals injected into a network of Cosmic Explorer and the Einstein Telescope. Results for non-precessing ($\theta_1 = 0$) sources are shown on the left and those for precessing sources ($\theta_1 = \pi/2$) on the right. The posteriors are colored by the network SNR of the signal. The linestyle of the posterior reflects the dimensionless spin magnitude of the more massive black hole, $a_1 = 0.6$ (dotted), 0.8 (dashed), or 0.9 (solid). Posteriors are organized by increasing mass ratio of the source, from top to bottom. Thin black lines rising from the $\mathcal{M}$ -axis indicate the true value of $\mathcal{M}$, and a grey line is included at the fiducial mass scale of $1 M_\odot$. We note that these posteriors are not normalized so that they may be visualized together.

and similar to our analysis for spin-aligned systems, we that $a_2 > a_{\text{NS}}$ when

$$|\chi_{\text{eff}}| \geq \frac{a_{\text{NS}} q}{1 + q} \quad (\text{C.8})$$

which we recognize as the value of $|\chi_{\text{eff}}|$ for $a_1 = 0, a_2 = a_{\text{NS}}$. This critical value of $|\chi_{\text{eff}}|$ is maximized for $q = 1$, yielding $\chi_{\text{eff, crit}} = 0.2$. Again, above this threshold, a_2 *must* be larger than the maximal neutron star spin, assuming the spins are anti-aligned.

D Features in correlated mass ratio-effective spin posteriors

In figure 15, we show marginal posteriors for two precessing simulated signals with $a_1 = 0.6$ in the O4 detector network; the runs shown are at the lower SNRs listed in table 1. The source with $q = 0.10$, which has a network SNR of 14.3 (left panel, lime color) shows bimodality in




Figure 13. Marginal posterior distributions on the mass ratio q for the simulated signals injected into a network of Cosmic Explorer and the Einstein Telescope. Results for non-precessing ($\theta_1 = 0$) sources are shown on the left and those for precessing sources ($\theta_1 = \pi/2$) on the right. The posteriors are colored by the network SNR of the signal. The linestyle of the posterior reflects the dimensionless spin magnitude of the more massive black hole, $a_1 = 0.6$ (dotted), 0.8 (dashed), or 0.9 (solid). Posteriors are organized by increasing mass ratio of the source, from top to bottom. Dashed black lines indicate the true value of q . We note that these posteriors are not normalized so that they may be visualized together.

the posterior distribution for χ_{eff} driven by posterior support for q away from the true value, shown in black. As q and χ_{eff} are correlated, a small change in q drives samples to a new iso-likelihood ridge in the space of $q - \chi_{\text{eff}}$, introducing a new mode at larger χ_{eff} . In the 2D marginal $q - \chi_{\text{eff}}$ posterior for this source, the lower probability mode coincides with the true value. It does not dominate the posterior because the likelihood is especially “peaky” and, in a sense, “underresolved”. As q decreases, the length of the signal increases, and so even relatively minor deviations in the source parameters will result in a large mismatch between a proposal for the signal and the data; thus, the likelihood peaks around an especially narrow range of source parameters. Using **dynesty**, the number of nested sampling live points N_{live} can be heuristically related to our resolution of the total prior volume. If we have too few live points, it is less likely to initially place them within the narrow width of the maximum likelihood peak. The analysis for the precessing $q = 0.10$, $a_1 = 0.6$ source in the O4 network employed $N_{\text{live}} = 2000$. Using fewer live points, the true value lay at the edge of the 1σ level




Figure 14. Marginal posteriors on the source-frame chirp mass $\mathcal{M}$, luminosity distance d_L , inclination angle relative to the line of sight θ_{JN} , and effective spin precession χ_p for the non-precessing simulated signal with $q = 0.36$, $a_1 = 0.9$, and a network SNR of 665.5 in the XG network. True values are shown with black lines. We observe that the chirp mass is underestimated due to an overestimated distance, driven by correlations between d_L , θ_{JN} , and χ_p .




Figure 15. Marginal posterior distributions on mass ratio q and effective spin χ_{eff} for two precessing signals (with $a_1 = 0.6$) in the O4 detector network that do not peak at the true value of χ_{eff} , driven by the $q - \chi_{\text{eff}}$ correlation. True values are shown with black lines, and the distributions are colored by the mass ratio with $q = 0.10$ (left, lime) and $q = 0.56$ (right, green-blue), matching the colors of figure 5. The $q = 0.10$ run has a network SNR of 14.3, and the $q = 0.56$ run has an SNR of 12.5.

of the marginal posterior on q , and with increasing N_{live} we observe a (still subdominant) mode emerge centered on the true source parameters. It is likely that increasing N_{live} further would allow us to fully resolve a unimodal posterior centered on the true q and χ_{eff} , although this becomes increasingly computationally expensive.

The source with $q = 0.56$, with a network SNR of 12.5 (right panel of figure 15, green-blue color) underestimates χ_{eff} , below the true value of zero. However, inspecting the joint posterior on $q - \chi_{\text{eff}}$ we see that the true value lies along a ridge of similar probability, making it (roughly) equally as likely for χ_{eff} to be slightly negative, and with slightly larger q . Without better measurement of the spin magnitudes and tilts, a larger range of χ_{eff} is allowed. Since the marginal posterior on q shows additional support at larger values, we underestimate χ_{eff} .

E Additional marginal posteriors on spin geometry

In figures 16 and 17, we include additional marginal posteriors on some components of the binary spin geometry.




Figure 16. Marginal posterior distributions on χ_{eff} for the louder simulated signals in the O4 detector network. This figure is organized in the same manner as figure 5, with the marginal posteriors colored by the true mass ratio q and organized from top to bottom by the true value of a_1 , with the prior on χ_{eff} shown at the top. Non-precessing signals are on the left, with the true value of χ_{eff} shown with a dashed line colored by the true q . Precessing signals are shown on the left where the truth is shown with a dashed black line. We also include the two critical values of χ_{eff} above which a_2 is inconsistent with neutron star spin derived in appendix C, assuming that spins are aligned (solid grey) or anti-aligned (dashed grey). Note that there are no signals with $q \leq 0.1$ in this plot as those sources were only simulated with one SNR (see table 1).




Figure 17. Marginal posterior distributions on the tilt angle between the spin vector of the heavier black hole and the orbital angular momentum, θ_1 , for the quietest signals for a given set of intrinsic parameters $(\tilde{m}_1, \tilde{m}_2, a_1, \theta_1)$, injected into a network of Cosmic Explorer and the Einstein Telescope. In the top row, we show the prior on $\pi(\theta_1)$. The truth is indicated with a dashed black line. We color the posterior distributions by the network SNR, and their linestyles reflect the true spin magnitude of the heavier black hole, a_1 . We observe that θ_1 is well-measured for non-precessing systems, and at low mass ratios. As noted in section 3.2.2, the lack of posterior support at $\theta_1 = \pi/2$ for precessing signals reflects the aligned-spin prior effect investigated in appendix F.

F Biased tilt posteriors at high signal-to-noise ratios

Assuming isotropic priors on the spin magnitudes and tilt angles, the prior on the aligned components of the spins $\pi(\chi_i)$ formally diverges to infinity at $\chi_i = 0$, i.e. $\theta_i = \pi/2$ (see equation A7 of ref. [134]). Nested sampling in **dynesty** takes place in a unit hypercube, requiring the inverse of the cumulative distribution function (CDF) of the prior; however, the CDF of $\pi(\chi_i)$ is not analytically invertible. Instead, **bilby** constructs a numerical approximant to the CDF which insufficiently resolves the divergence at $\chi_i = 0$ in $\pi(\chi_i)$ for very high SNR signals, resulting in little to no posterior support at $\theta_i = \pi/2$ as observed in figures 8 and 17.

To verify that this bias in θ_1 is indeed a prior effect, we sample directly in the spin magnitudes and tilts to re-analyze the simulated signal with $q = 0.10$, $a_1 = 0.9$, and $\theta_1 = \pi/2$ in the




Figure 18. Marginal posterior on θ_1 from two analyses of the simulated source with $q = 0.10$, $a_1 = 0.9$, and $\theta_1 = \pi/2$ (dashed black line) injected into the XG network with an SNR of 490.3. Here, we repeat the posterior shown in figure 8 (solid line) and compare it to the posterior generated when adopting uniform priors on the spin magnitudes a_i and tilt angles θ_i (dash-dotted line). The uniform prior on θ_1 is shown in the black solid line. We observe that sampling in the spin magnitudes and tilts directly recovers a posterior of similar width as the posterior recovered when sampling in the aligned (χ_i) and in-plane spins ($\chi_i^\perp$), and it also recovers the true value of θ_1 .

XG network of Cosmic Explorer and the Einstein Telescope, with a network SNR of 490.3. In figure 18 we compare the marginal posteriors on θ_1 recovered by sampling in the aligned- and in-plane ($\chi_i^\perp$) spin components with isotropic priors (solid line) and in the spin magnitudes a_i and tilt angles θ_i with uniform priors (dash-dotted line). We note that, in the range of θ_1 considered very near $\pi/2$, this is nearly equivalent to an isotropic (sine) prior on θ_1 . The posterior generated by sampling in $\chi_i, \chi_i^\perp$ is repeated from the inset panel in the second row, right column of figure 8. Here, we see that the posterior generated by sampling in a_i, θ_i is of a similar width as the original result; sampling in $\chi_i, \chi_i^\perp$ we found a 90% credible interval on θ_1 of 3.3×10^{-3} radians versus a credible interval of 2.2×10^{-3} radians when sampling in the spin magnitude and tilt directly. Importantly, we also successfully recover the true value of θ_1 when using uniform priors on a_i, θ_i , indicating that the bias in θ_1 we observed in figure 8 is a prior effect.

G Correlation between spin precession and binary geometry

At the three lowest mass ratios of the simulated precessing signals considered in this work, $q = 0.07, 0.10, 0.36$, we observed a correlation between the effective spin precession χ_p and the zenith angles of the orbital angular momentum $\vec{L}$ and total angular momentum $\vec{J}$. To our knowledge, this correlation has not been previously reported. In figure 19 we show marginal posteriors on χ_p , θ_{JN} (zenith angle between $\vec{J}$ and the line of sight $\hat{N}$), θ_{JL} (zenith angle between $\vec{J}$ and $\vec{L}$), and ι (zenith angle between $\vec{L}$ and $\hat{N}$) for the precessing source with $q = 0.36$ and $a_1 = 0.9$ in the O4-design sensitivity network, which had a network SNR of 21.2. We compute θ_{JL} (also referred to as the opening angle, β) at a reference frequency of 20 Hz using `pesummary.gw.conversions.spins.opening_angle()`, which in turn implements methods from `LALSimulation`. Along the bottom row, we see linear correlations between χ_p and each of these angles.




Figure 19. Marginal posteriors on the effective spin precession χ_p and the zenith angles θ_{JN} (between the total angular momentum and line of sight), θ_{JL} (between the total and orbital angular momenta), and ι (between the orbital angular momentum and line of sight), from our analysis of the simulated $q = 0.36$, $a_1 = 0.9$ precessing signal in the O4-design sensitivity network. True values are shown in black lines.

Heuristically, this correlation can be understood as follows: consider a gravitational-wave signal consistent with precession. This necessarily implies a non-zero angle θ_{JL} between $\vec{L}$ and $\vec{J}$. If we infer $\theta_1 \neq 0$, we can explain larger (smaller) values of this angle with larger (smaller) values of a_1 which pushes $\vec{J}$ into (away from) the orbital plane and away from (towards) $\vec{L}$. Note that in all of the signals we simulated with $a_1 > 0, a_2 = 0$, we have $\chi_p = a_1 \sin \theta_1$ (cf. equation (3.2)).⁹ So, larger χ_p necessitates larger θ_{JL} , and thus these quantities are positively correlated. Alternatively, if we are uncertain of the degree to which we observe precession in a gravitational wave signal, we can fix all of the geometry of our system except for θ_{JL} and θ_{JN} , in particular fixing the component black hole spins and ι . As $\vec{L}$ precesses about $\vec{J}$, if we enlarge the cone of precession by increasing θ_{JL} then we must bring this cone closer to the line of sight by decreasing θ_{JN} to keep ι constant. In this way, θ_{JN} and θ_{JL} are negatively linearly correlated.

Combined, we have a negative linear correlation between χ_p and θ_{JN} . Similar heuristic arguments can explain the observed correlations with ι . In detail, this effect is likely because precession for these systems is in the “tropical region” noted by [79], where the plane of the orbit can wobble so violently that both the top and bottom are observed along the line of sight in the course of precession. This injects additional information on the binary geometry of the system into the gravitational wave signal, which has previously been observed to reduce uncertainty in the measurement of θ_{JN} (see e.g. figure 4 of [99] or figure 20 of [97]).

⁹Alternatively, we could increase θ_1 if we infer non-zero spin a_1 to achieve the same effect, although the change in χ_p with respect to θ_1 is much smaller near the true value of $\theta_1 = \pi/2$ when $\chi_p \propto \sin \theta_1$.

References

- [1] S. Bird et al., *Did LIGO detect dark matter?*, *Phys. Rev. Lett.* **116** (2016) 201301 [[arXiv:1603.00464](#)] [[INSPIRE](#)].
- [2] M. Sasaki, T. Suyama, T. Tanaka and S. Yokoyama, *Primordial Black Hole Scenario for the Gravitational-Wave Event GW150914*, *Phys. Rev. Lett.* **117** (2016) 061101 [Erratum *ibid.* **121** (2018) 059901] [[arXiv:1603.08338](#)] [[INSPIRE](#)].
- [3] S. Clesse and J. García-Bellido, *The clustering of massive Primordial Black Holes as Dark Matter: measuring their mass distribution with Advanced LIGO*, *Phys. Dark Univ.* **15** (2017) 142 [[arXiv:1603.05234](#)] [[INSPIRE](#)].
- [4] S. Clesse and J. García-Bellido, *GW190425, GW190521 and GW190814: Three candidate mergers of primordial black holes from the QCD epoch*, *Phys. Dark Univ.* **38** (2022) 101111 [[arXiv:2007.06481](#)] [[INSPIRE](#)].
- [5] Y. Ali-Haïmoud, E.D. Kovetz and M. Kamionkowski, *Merger rate of primordial black-hole binaries*, *Phys. Rev. D* **96** (2017) 123523 [[arXiv:1709.06576](#)] [[INSPIRE](#)].
- [6] S. Clesse and J. García-Bellido, *Seven Hints for Primordial Black Hole Dark Matter*, *Phys. Dark Univ.* **22** (2018) 137 [[arXiv:1711.10458](#)] [[INSPIRE](#)].
- [7] K.W.K. Wong et al., *Constraining the primordial black hole scenario with Bayesian inference and machine learning: the GWTC-2 gravitational wave catalog*, *Phys. Rev. D* **103** (2021) 023026 [[arXiv:2011.01865](#)] [[INSPIRE](#)].
- [8] V. De Luca, G. Franciolini, P. Pani and A. Riotto, *Primordial Black Holes Confront LIGO/Virgo data: Current situation*, *JCAP* **06** (2020) 044 [[arXiv:2005.05641](#)] [[INSPIRE](#)].
- [9] V. De Luca, G. Franciolini, P. Pani and A. Riotto, *Bayesian Evidence for Both Astrophysical and Primordial Black Holes: Mapping the GWTC-2 Catalog to Third-Generation Detectors*, *JCAP* **05** (2021) 003 [[arXiv:2102.03809](#)] [[INSPIRE](#)].
- [10] Z.-C. Chen, C. Yuan and Q.-G. Huang, *Confronting the primordial black hole scenario with the gravitational-wave events detected by LIGO-Virgo*, *Phys. Lett. B* **829** (2022) 137040 [[arXiv:2108.11740](#)] [[INSPIRE](#)].
- [11] G. Franciolini, I. Musco, P. Pani and A. Urbano, *From inflation to black hole mergers and back again: Gravitational-wave data-driven constraints on inflationary scenarios with a first-principle model of primordial black holes across the QCD epoch*, *Phys. Rev. D* **106** (2022) 123526 [[arXiv:2209.05959](#)] [[INSPIRE](#)].
- [12] LIGO SCIENTIFIC collaboration, *Advanced LIGO*, *Class. Quant. Grav.* **32** (2015) 074001 [[arXiv:1411.4547](#)] [[INSPIRE](#)].
- [13] VIRGO collaboration, *Advanced Virgo: a second-generation interferometric gravitational wave detector*, *Class. Quant. Grav.* **32** (2015) 024001 [[arXiv:1408.3978](#)] [[INSPIRE](#)].
- [14] KAGRA collaboration, *Overview of KAGRA: Calibration, detector characterization, physical environmental monitors, and the geophysics interferometer*, *PTEP* **2021** (2021) 05A102 [[arXiv:2009.09305](#)] [[INSPIRE](#)].
- [15] M. Evans et al., *A Horizon Study for Cosmic Explorer: Science, Observatories, and Community*, [arXiv:2109.09882](#) [[INSPIRE](#)].
- [16] M. Punturo et al., *The Einstein Telescope: A third-generation gravitational wave observatory*, *Class. Quant. Grav.* **27** (2010) 194002 [[INSPIRE](#)].

- [17] S. Chandrasekhar, *The maximum mass of ideal white dwarfs*, *Astrophys. J.* **74** (1931) 81 [[INSPIRE](#)].
- [18] P. Villanueva-Domingo, O. Mena and S. Palomares-Ruiz, *A brief review on primordial black holes as dark matter*, *Front. Astron. Space Sci.* **8** (2021) 87 [[arXiv:2103.12087](#)] [[INSPIRE](#)].
- [19] Y.B. Zel'dovich and I.D. Novikov, *The Hypothesis of Cores Retarded during Expansion and the Hot Cosmological Model*, *Soviet Astron. AJ (Engl. Transl.)* **10** (1967) 602 [[INSPIRE](#)].
- [20] S. Hawking, *Gravitationally collapsed objects of very low mass*, *Mon. Not. Roy. Astron. Soc.* **152** (1971) 75 [[INSPIRE](#)].
- [21] B. Carr and F. Kühnel, *Primordial black holes as dark matter candidates*, *SciPost Phys. Lect. Notes* **48** (2022) 1 [[arXiv:2110.02821](#)] [[INSPIRE](#)].
- [22] J.C. Niemeyer and K. Jedamzik, *Near-critical gravitational collapse and the initial mass function of primordial black holes*, *Phys. Rev. Lett.* **80** (1998) 5481 [[astro-ph/9709072](#)] [[INSPIRE](#)].
- [23] J.C. Niemeyer and K. Jedamzik, *Dynamics of primordial black hole formation*, *Phys. Rev. D* **59** (1999) 124013 [[astro-ph/9901292](#)] [[INSPIRE](#)].
- [24] I. Musco, J.C. Miller and L. Rezzolla, *Computations of primordial black hole formation*, *Class. Quant. Grav.* **22** (2005) 1405 [[gr-qc/0412063](#)] [[INSPIRE](#)].
- [25] B. Carr, S. Clesse, J. García-Bellido and F. Kühnel, *Cosmic conundra explained by thermal history and primordial black holes*, *Phys. Dark Univ.* **31** (2021) 100755 [[arXiv:1906.08217](#)] [[INSPIRE](#)].
- [26] S. Young and C.T. Byrnes, *Primordial black holes in non-Gaussian regimes*, *JCAP* **08** (2013) 052 [[arXiv:1307.4995](#)] [[INSPIRE](#)].
- [27] E.V. Bugaev and P.A. Klimai, *Primordial black hole constraints for curvaton models with predicted large non-Gaussianity*, *Int. J. Mod. Phys. D* **22** (2013) 1350034 [[arXiv:1303.3146](#)] [[INSPIRE](#)].
- [28] G. Franciolini, A. Kehagias, S. Matarrese and A. Riotto, *Primordial Black Holes from Inflation and non-Gaussianity*, *JCAP* **03** (2018) 016 [[arXiv:1801.09415](#)] [[INSPIRE](#)].
- [29] A. Dolgov and J. Silk, *Baryon isocurvature fluctuations at small scales and baryonic dark matter*, *Phys. Rev. D* **47** (1993) 4244 [[INSPIRE](#)].
- [30] B.J. Carr and J.E. Lidsey, *Primordial black holes and generalized constraints on chaotic inflation*, *Phys. Rev. D* **48** (1993) 543 [[INSPIRE](#)].
- [31] P. Ivanov, P. Naselsky and I. Novikov, *Inflation and primordial black holes as dark matter*, *Phys. Rev. D* **50** (1994) 7173 [[INSPIRE](#)].
- [32] L. Randall, M. Soljacic and A.H. Guth, *Supernatural inflation: Inflation from supersymmetry with no (very) small parameters*, *Nucl. Phys. B* **472** (1996) 377 [[hep-ph/9512439](#)] [[INSPIRE](#)].
- [33] J. García-Bellido, A.D. Linde and D. Wands, *Density perturbations and black hole formation in hybrid inflation*, *Phys. Rev. D* **54** (1996) 6040 [[astro-ph/9605094](#)] [[INSPIRE](#)].
- [34] S.W. Hawking, *Black Holes From Cosmic Strings*, *Phys. Lett. B* **231** (1989) 237 [[INSPIRE](#)].
- [35] M.Y. Khlopov, R.V. Konoplich, S.G. Rubin and A.S. Sakharov, *First order phase transitions as a source of black holes in the early universe*, *Grav. Cosmol.* **2** (1999) S1 [[hep-ph/9912422](#)] [[INSPIRE](#)].
- [36] J. Garriga, A. Vilenkin and J. Zhang, *Black holes and the multiverse*, *JCAP* **02** (2016) 064 [[arXiv:1512.01819](#)] [[INSPIRE](#)].

- [37] M. Raidal, C. Spethmann, V. Vaskonen and H. Veermäe, *Formation and Evolution of Primordial Black Hole Binaries in the Early Universe*, *JCAP* **02** (2019) 018 [[arXiv:1812.01930](#)] [[INSPIRE](#)].
- [38] T. Matsubara, T. Terada, K. Kohri and S. Yokoyama, *Clustering of primordial black holes formed in a matter-dominated epoch*, *Phys. Rev. D* **100** (2019) 123544 [[arXiv:1909.04053](#)] [[INSPIRE](#)].
- [39] T. Suyama and S. Yokoyama, *Clustering of primordial black holes with non-Gaussian initial fluctuations*, *PTEP* **2019** (2019) 103E02 [[arXiv:1906.04958](#)] [[INSPIRE](#)].
- [40] M. Trashorras, J. García-Bellido and S. Nesseris, *The clustering dynamics of primordial black holes in N -body simulations*, *Universe* **7** (2021) 18 [[arXiv:2006.15018](#)] [[INSPIRE](#)].
- [41] G. Hütsi, M. Raidal, V. Vaskonen and H. Veermäe, *Two populations of LIGO-Virgo black holes*, *JCAP* **03** (2021) 068 [[arXiv:2012.02786](#)] [[INSPIRE](#)].
- [42] K.S. Phukon et al., *The hunt for sub-solar primordial black holes in low mass ratio binaries is open*, [arXiv:2105.11449](#) [[INSPIRE](#)].
- [43] K. Jedamzik, *Consistency of Primordial Black Hole Dark Matter with LIGO/Virgo Merger Rates*, *Phys. Rev. Lett.* **126** (2021) 051302 [[arXiv:2007.03565](#)] [[INSPIRE](#)].
- [44] G. D’Amico et al., *Massive Black Holes from Dissipative Dark Matter*, *Mon. Not. Roy. Astron. Soc.* **473** (2018) 328 [[arXiv:1707.03419](#)] [[INSPIRE](#)].
- [45] S. Shandera, D. Jeong and H.S.G. Gebhardt, *Gravitational Waves from Binary Mergers of Subsolar Mass Dark Black Holes*, *Phys. Rev. Lett.* **120** (2018) 241102 [[arXiv:1802.08206](#)] [[INSPIRE](#)].
- [46] J.H. Chang, D. Egana-Ugrinovic, R. Essig and C. Kouvaris, *Structure Formation and Exotic Compact Objects in a Dissipative Dark Sector*, *JCAP* **03** (2019) 036 [[arXiv:1812.07000](#)] [[INSPIRE](#)].
- [47] J. Choquette, J.M. Cline and J.M. Cornell, *Early formation of supermassive black holes via dark matter self-interactions*, *JCAP* **07** (2019) 036 [[arXiv:1812.05088](#)] [[INSPIRE](#)].
- [48] M.A. Latif et al., *Black hole formation in the context of dissipative dark matter*, *Mon. Not. Roy. Astron. Soc.* **485** (2019) 3352 [[arXiv:1812.03104](#)] [[INSPIRE](#)].
- [49] R. Essig, S.D. McDermott, H.-B. Yu and Y.-M. Zhong, *Constraining Dissipative Dark Matter Self-Interactions*, *Phys. Rev. Lett.* **123** (2019) 121102 [[arXiv:1809.01144](#)] [[INSPIRE](#)].
- [50] M. Ryan, J. Gurian, S. Shandera and D. Jeong, *Molecular Chemistry for Dark Matter*, *Astrophys. J.* **934** (2022) 120 [[arXiv:2106.13245](#)] [[INSPIRE](#)].
- [51] M. Hippert et al., *Mirror neutron stars*, *Phys. Rev. D* **106** (2022) 035025 [[arXiv:2103.01965](#)] [[INSPIRE](#)].
- [52] J. Gurian et al., *A Lower Bound on the Mass of Compact Objects from Dissipative Dark Matter*, *Astrophys. J. Lett.* **939** (2022) L12 [Erratum *ibid.* **949** (2023) L44] [[arXiv:2209.00064](#)] [[INSPIRE](#)].
- [53] I. de Martino et al., *Dark matters on the scale of galaxies*, *Universe* **6** (2020) 107 [[arXiv:2007.15539](#)] [[INSPIRE](#)].
- [54] LIGO SCIENTIFIC et al. collaborations, *Search for subsolar-mass black hole binaries in the second part of Advanced LIGO’s and Advanced Virgo’s third observing run*, [arXiv:2212.01477](#) [[INSPIRE](#)].

- [55] A.H. Nitz and Y.-F. Wang, *Broad search for gravitational waves from subsolar-mass binaries through LIGO and Virgo’s third observing run*, *Phys. Rev. D* **106** (2022) 023024 [[arXiv:2202.11024](#)] [[INSPIRE](#)].
- [56] LIGO SCIENTIFIC et al. collaborations, *Search for Subsolar-Mass Binaries in the First Half of Advanced LIGO’s and Advanced Virgo’s Third Observing Run*, *Phys. Rev. Lett.* **129** (2022) 061104 [[arXiv:2109.12197](#)] [[INSPIRE](#)].
- [57] A.H. Nitz and Y.-F. Wang, *Search for Gravitational Waves from the Coalescence of Subsolar-Mass Binaries in the First Half of Advanced LIGO and Virgo’s Third Observing Run*, *Phys. Rev. Lett.* **127** (2021) 151101 [[arXiv:2106.08979](#)] [[INSPIRE](#)].
- [58] M. Pieroni, A. Ricciardone and E. Barausse, *Detectability and parameter estimation of stellar origin black hole binaries with next generation gravitational wave detectors*, *Sci. Rep.* **12** (2022) 17940 [[arXiv:2203.12586](#)] [[INSPIRE](#)].
- [59] A. Bandopadhyay et al., *Detectability of subsolar mass neutron stars through a template bank search*, *Phys. Rev. D* **107** (2023) 103012 [[arXiv:2212.03855](#)] [[INSPIRE](#)].
- [60] S.L. Liebling and C. Palenzuela, *Dynamical boson stars*, *Living Rev. Rel.* **26** (2023) 1 [[arXiv:1202.5809](#)] [[INSPIRE](#)].
- [61] V. Doroshenko, V. Suleimanov, G. Pühlhofer and A. Santangelo, *A strangely light neutron star within a supernova remnant*, *Nature Astron.* **6** (2022) 1444.
- [62] P. Landry and J.S. Read, *The Mass Distribution of Neutron Stars in Gravitational-wave Binaries*, *Astrophys. J. Lett.* **921** (2021) L25 [[arXiv:2107.04559](#)] [[INSPIRE](#)].
- [63] M.M. Meskhi et al., *A New Constraint on the Nuclear Equation of State from Statistical Distributions of Compact Remnants of Supernovae*, *Astrophys. J. Lett.* **932** (2022) L3 [[arXiv:2111.01815](#)] [[INSPIRE](#)].
- [64] G.F. Giudice, M. McCullough and A. Urbano, *Hunting for Dark Particles with Gravitational Waves*, *JCAP* **10** (2016) 001 [[arXiv:1605.01209](#)] [[INSPIRE](#)].
- [65] G. Morras et al., *Analysis of a subsolar-mass compact binary candidate from the second observing run of Advanced LIGO*, [arXiv:2301.11619](#) [[DOI:10.1016/j.dark.2023.101285](#)] [[INSPIRE](#)].
- [66] M. Hannam et al., *When can gravitational-wave observations distinguish between black holes and neutron stars?*, *Astrophys. J. Lett.* **766** (2013) L14 [[arXiv:1301.5616](#)] [[INSPIRE](#)].
- [67] S. Wu and A.H. Nitz, *Mock data study for next-generation ground-based detectors: The performance loss of matched filtering due to correlated confusion noise*, *Phys. Rev. D* **107** (2023) 063022 [[arXiv:2209.03135](#)] [[INSPIRE](#)].
- [68] J. Skilling, *Nested Sampling*, *AIP Conf. Proc.* **735** (2004) 395.
- [69] J. Skilling, *Nested sampling for general Bayesian computation*, *Bayesian Anal.* **1** (2006) 833 [[INSPIRE](#)].
- [70] J.S. Speagle, *dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences*, *Mon. Not. Roy. Astron. Soc.* **493** (2020) 3132 [[arXiv:1904.02180](#)] [[INSPIRE](#)].
- [71] LIGO SCIENTIFIC and VIRGO collaborations, *Properties of the binary neutron star merger GW170817*, *Phys. Rev. X* **9** (2019) 011001 [[arXiv:1805.11579](#)] [[INSPIRE](#)].
- [72] N.J. Cornish, *Heterodyned likelihood for rapid gravitational wave parameter inference*, *Phys. Rev. D* **104** (2021) 104054 [[arXiv:2109.02728](#)] [[INSPIRE](#)].
- [73] B. Zackay, L. Dai and T. Venumadhav, *Relative Binning and Fast Likelihood Evaluation for Gravitational Wave Parameter Estimation*, [arXiv:1806.08792](#) [[INSPIRE](#)].

- [74] N. Leslie, L. Dai and G. Pratten, *Mode-by-mode relative binning: Fast likelihood estimation for gravitational waveforms with spin-orbit precession and multiple harmonics*, *Phys. Rev. D* **104** (2021) 123030 [[arXiv:2109.09872](#)] [[INSPIRE](#)].
- [75] G. Ashton et al., *BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy*, *Astrophys. J. Suppl.* **241** (2019) 27 [[arXiv:1811.02042](#)] [[INSPIRE](#)].
- [76] I.M. Romero-Shaw et al., *Bayesian inference for compact binary coalescences with bilby: validation and application to the first LIGO-Virgo gravitational-wave transient catalogue*, *Mon. Not. Roy. Astron. Soc.* **499** (2020) 3295 [[arXiv:2006.00714](#)] [[INSPIRE](#)].
- [77] K. Krishna, A. Vijaykumar, A. Ganguly and P. Ajith, *Relative Binning in BILBY*, in preparation.
- [78] J.W.T. Hessels et al., *A radio pulsar spinning at 716-hz*, *Science* **311** (2006) 1901 [[astro-ph/0601337](#)] [[INSPIRE](#)].
- [79] T.A. Apostolatos, C. Cutler, G.J. Sussman and K.S. Thorne, *Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries*, *Phys. Rev. D* **49** (1994) 6274 [[INSPIRE](#)].
- [80] V. De Luca, G. Franciolini, P. Pani and A. Riotto, *The evolution of primordial black holes and their final observable spins*, *JCAP* **04** (2020) 052 [[arXiv:2003.02778](#)] [[INSPIRE](#)].
- [81] T. Harada, C.-M. Yoo, K. Kohri and K.-I. Nakao, *Spins of primordial black holes formed in the matter-dominated phase of the Universe*, *Phys. Rev. D* **96** (2017) 083517 [Erratum *ibid.* **99** (2019) 069904] [[arXiv:1707.03595](#)] [[INSPIRE](#)].
- [82] PLANCK collaboration, *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [[arXiv:1502.01589](#)] [[INSPIRE](#)].
- [83] LIGO SCIENTIFIC, VIRGO and KAGRA collaborations, *Noise curves used for Simulations in the update of the Observing Scenarios Paper*, [LIGO-T2000012](#).
- [84] KAGRA, LIGO SCIENTIFIC and VIRGO collaborations, *Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA*, *Living Rev. Rel.* **21** (2018) 3 [[arXiv:1304.0670](#)] [[INSPIRE](#)].
- [85] M. Evans, J. Harms and S. Vitale, *Exploring the Sensitivity of Next Generation Gravitational Wave Detectors*, [LIGO-P1600143](#).
- [86] LIGO SCIENTIFIC collaboration, *Exploring the Sensitivity of Next Generation Gravitational Wave Detectors*, *Class. Quant. Grav.* **34** (2017) 044001 [[arXiv:1607.08697](#)] [[INSPIRE](#)].
- [87] S. Nissanke et al., *Exploring short gamma-ray bursts as gravitational-wave standard sirens*, *Astrophys. J.* **725** (2010) 496 [[arXiv:0904.1017](#)] [[INSPIRE](#)].
- [88] M. Vallisneri, *Beyond Fisher: exact sampling distributions of the maximum-likelihood estimator in gravitational-wave parameter estimation*, *Phys. Rev. Lett.* **107** (2011) 191104 [[arXiv:1108.1158](#)] [[INSPIRE](#)].
- [89] C.L. Rodriguez et al., *Basic Parameter Estimation of Binary Neutron Star Systems by the Advanced LIGO/Virgo Network*, *Astrophys. J.* **784** (2014) 119 [[arXiv:1309.3273](#)] [[INSPIRE](#)].
- [90] C. Pankow et al., *Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817*, *Phys. Rev. D* **98** (2018) 084016 [[arXiv:1808.03619](#)] [[INSPIRE](#)].
- [91] G. Pratten et al., *Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes*, *Phys. Rev. D* **103** (2021) 104056 [[arXiv:2004.06503](#)] [[INSPIRE](#)].

- [92] J. Calderón Bustillo, S.H.W. Leong, T. Dietrich and P.D. Lasky, *Mapping the Universe Expansion: Enabling Percent-level Measurements of the Hubble Constant with a Single Binary Neutron-star Merger Detection*, *Astrophys. J. Lett.* **912** (2021) L10 [[arXiv:2006.11525](#)] [[INSPIRE](#)].
- [93] G. Pratten et al., *Setting the cornerstone for a family of models for gravitational waves from compact binaries: The dominant harmonic for nonprecessing quasicircular black holes*, *Phys. Rev. D* **102** (2020) 064001 [[arXiv:2001.11412](#)] [[INSPIRE](#)].
- [94] C. García-Quirós et al., *Multimode frequency-domain model for the gravitational wave signal from nonprecessing black-hole binaries*, *Phys. Rev. D* **102** (2020) 064002 [[arXiv:2001.10914](#)] [[INSPIRE](#)].
- [95] K.G. Arun, B.R. Iyer, B.S. Sathyaprakash and P.A. Sundararajan, *Parameter estimation of inspiralling compact binaries using 3.5 post-Newtonian gravitational wave phasing: The Non-spinning case*, *Phys. Rev. D* **71** (2005) 084008 [Erratum *ibid.* **72** (2005) 069903] [[gr-qc/0411146](#)] [[INSPIRE](#)].
- [96] B. Farr et al., *Parameter estimation on gravitational waves from neutron-star binaries with spinning components*, *Astrophys. J.* **825** (2016) 116 [[arXiv:1508.05336](#)] [[INSPIRE](#)].
- [97] S. Vitale et al., *Parameter estimation for heavy binary-black holes with networks of second-generation gravitational-wave detectors*, *Phys. Rev. D* **95** (2017) 064053 [[arXiv:1611.01122](#)] [[INSPIRE](#)].
- [98] LIGO SCIENTIFIC et al. collaborations, *GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run*, [arXiv:2111.03606](#) [[INSPIRE](#)].
- [99] S. Vitale et al., *Measuring the spin of black holes in binary systems using gravitational waves*, *Phys. Rev. Lett.* **112** (2014) 251101 [[arXiv:1403.0129](#)] [[INSPIRE](#)].
- [100] M. Pürrer, M. Hannam and F. Ohme, *Can we measure individual black-hole spins from gravitational-wave observations?*, *Phys. Rev. D* **93** (2016) 084042 [[arXiv:1512.04955](#)] [[INSPIRE](#)].
- [101] T. Damour, *Coalescence of two spinning black holes: an effective one-body approach*, *Phys. Rev. D* **64** (2001) 124013 [[gr-qc/0103018](#)] [[INSPIRE](#)].
- [102] P. Ajith et al., *Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins*, *Phys. Rev. Lett.* **106** (2011) 241101 [[arXiv:0909.2867](#)] [[INSPIRE](#)].
- [103] L. Santamaria et al., *Matching post-Newtonian and numerical relativity waveforms: systematic errors and a new phenomenological model for non-precessing black hole binaries*, *Phys. Rev. D* **82** (2010) 064016 [[arXiv:1005.3306](#)] [[INSPIRE](#)].
- [104] E. Baird, S. Fairhurst, M. Hannam and P. Murphy, *Degeneracy between mass and spin in black-hole-binary waveforms*, *Phys. Rev. D* **87** (2013) 024035 [[arXiv:1211.0546](#)] [[INSPIRE](#)].
- [105] K.K.Y. Ng et al., *Gravitational-wave astrophysics with effective-spin measurements: asymmetries and selection biases*, *Phys. Rev. D* **98** (2018) 083007 [[arXiv:1805.03046](#)] [[INSPIRE](#)].
- [106] T. Callister, *A Thesaurus for Common Priors in Gravitational-Wave Astronomy*, [arXiv:2104.09508](#) [[INSPIRE](#)].
- [107] P. Schmidt, F. Ohme and M. Hannam, *Towards models of gravitational waveforms from generic binaries II: Modelling precession effects with a single effective precession parameter*, *Phys. Rev. D* **91** (2015) 024043 [[arXiv:1408.1810](#)] [[INSPIRE](#)].
- [108] D. Gerosa et al., *A generalized precession parameter χ_p to interpret gravitational-wave data*, *Phys. Rev. D* **103** (2021) 064067 [[arXiv:2011.11948](#)] [[INSPIRE](#)].

- [109] R. Dekany et al., *The Zwicky Transient Facility: Observing System*, *Publ. Astron. Soc. Pac.* **132** (2020) 038001 [[arXiv:2008.04923](#)] [[INSPIRE](#)].
- [110] LSST collaboration, *LSST: from Science Drivers to Reference Design and Anticipated Data Products*, *Astrophys. J.* **873** (2019) 111 [[arXiv:0805.2366](#)] [[INSPIRE](#)].
- [111] R. Selina et al., *Science with an ngVLA: The ngVLA Reference Design*, [arXiv:1810.08197](#) [[DOI:10.48550/ARXIV.1810.08197](#)].
- [112] DES collaboration, *The Dark Energy Camera*, *Astron. J.* **150** (2015) 150 [[arXiv:1504.02900](#)] [[INSPIRE](#)].
- [113] W. Skidmore, G.C. Anupama and R. Srianand, *The Thirty Meter Telescope International Observatory facilitating transformative astrophysical science*, [arXiv:1806.02481](#) [[DOI:10.48550/ARXIV.1806.02481](#)].
- [114] S. Biscoveanu, P. Landry and S. Vitale, *Population properties and multimessenger prospects of neutron star-black hole mergers following GWTC-3*, *Mon. Not. Roy. Astron. Soc.* **518** (2022) 5298 [[arXiv:2207.01568](#)] [[INSPIRE](#)].
- [115] 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** (2018) 081501 [[arXiv:1807.00011](#)] [[INSPIRE](#)].
- [116] L.P. Singer and L.R. Price, *Rapid Bayesian position reconstruction for gravitational-wave transients*, *Phys. Rev. D* **93** (2016) 024013 [[arXiv:1508.03634](#)] [[INSPIRE](#)].
- [117] P. Baral, S. Morisaki, I. Magaña Hernandez and J. Creighton, *Localization of binary neutron star mergers with a single cosmic explorer*, *Phys. Rev. D* **108** (2023) 043010 [[arXiv:2304.09889](#)] [[INSPIRE](#)].
- [118] I. Markin et al., *General-Relativistic Hydrodynamics Simulation of a Neutron Star — Sub-Solar-Mass Black Hole Merger*, [arXiv:2304.11642](#) [[INSPIRE](#)].
- [119] ASTROPY collaboration, *The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package*, *Astrophys. J.* **935** (2022) 167 [[arXiv:2206.14220](#)] [[INSPIRE](#)].
- [120] C. Hoy and V. Raymond, *PESummary: the code agnostic Parameter Estimation Summary page builder*, *SoftwareX* **15** (2021) 100765 [[arXiv:2006.06639](#)] [[INSPIRE](#)].
- [121] K.M. Górski et al., *HEALPix — A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[INSPIRE](#)].
- [122] A. Zonca et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python*, *J. Open Source Softw.* **4** (2019) 1298 [[INSPIRE](#)].
- [123] D. Foreman-Mackey, *corner.py: Scatterplot matrices in Python*, *J. Open Source Softw.* **1** (2016) 24.
- [124] C.R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357 [[arXiv:2006.10256](#)] [[INSPIRE](#)].
- [125] P. Virtanen et al., *SciPy 1.0 — Fundamental Algorithms for Scientific Computing in Python*, *Nature Meth.* **17** (2020) 261 [[arXiv:1907.10121](#)] [[INSPIRE](#)].
- [126] J. Reback et al., *pandas-dev/pandas: Pandas 1.4.2*, [DOI:10.5281/ZENODO.6408044](#) (2022).

- [127] W. McKinney, *Data Structures for Statistical Computing in Python*, in the proceedings of the *9th Python in Science Conference*, S. van der Walt and J. Millman eds., (2010), p. 56–61 [[DOI:10.25080/majora-92bf1922-00a](https://doi.org/10.25080/majora-92bf1922-00a)] [[INSPIRE](#)].
- [128] LIGO SCIENTIFIC collaboration, *LIGO Algorithm Library*, <https://git.ligo.org/lscsoft/lalsuite> [[DOI:10.7935/GT1W-FZ16](https://doi.org/10.7935/GT1W-FZ16)].
- [129] K. Wette, *SWIGLAL: Python and Octave interfaces to the LALSuite gravitational-wave data analysis libraries*, *SoftwareX* **12** (2020) 100634 [[arXiv:2012.09552](https://arxiv.org/abs/2012.09552)] [[INSPIRE](#)].
- [130] J.D. Romano and N.J. Cornish, *Detection methods for stochastic gravitational-wave backgrounds: a unified treatment*, *Living Rev. Rel.* **20** (2017) 2 [[arXiv:1608.06889](https://arxiv.org/abs/1608.06889)] [[INSPIRE](#)].
- [131] S. Khan, K. Chatziioannou, M. Hannam and F. Ohme, *Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects*, *Phys. Rev. D* **100** (2019) 024059 [[arXiv:1809.10113](https://arxiv.org/abs/1809.10113)] [[INSPIRE](#)].
- [132] E. Thrane and C. Talbot, *An introduction to Bayesian inference in gravitational-wave astronomy: parameter estimation, model selection, and hierarchical models*, *Publ. Astron. Soc. Austral.* **36** (2019) e010 [[arXiv:1809.02293](https://arxiv.org/abs/1809.02293)] [[INSPIRE](#)].
- [133] C.J.F.T. Braak, *A Markov Chain Monte Carlo version of the genetic algorithm Differential Evolution: easy Bayesian computing for real parameter spaces*, *Stat. Comput.* **16** (2006) 239.
- [134] J. Lange, R. O’Shaughnessy and M. Rizzo, *Rapid and accurate parameter inference for coalescing, precessing compact binaries*, [arXiv:1805.10457](https://arxiv.org/abs/1805.10457) [[INSPIRE](#)].