

Results from an injection and recovery study involving 35 numerical relativity simulations and three waveform models

Sarp Akçay¹, Charlie Hoy², and Jake Mac Uilliam^{1,3}

¹*University College Dublin, Belfield, D4, Dublin, Ireland*

²*Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth, PO1 3FX, United Kingdom*

³*Irish Centre for High-End Computing, 2, 7/F, The Tower, Trinity Technology and Enterprise Campus, Grand Canal Dock, D2, Dublin, Ireland*



(Received 27 June 2025; accepted 15 August 2025; published 14 October 2025)

We present Bayesian inference results from an extensive injection-recovery campaign to test the validity of three state of the art quasicircular gravitational waveform models: SEOBNRv5PHM, IMRPhenomTPHM, IMRPhenomXPHM, the latter with the SpinTaylorT4 implementation for its precession dynamics. We analyze 35 strongly precessing binary black hole numerical relativity simulations with all available harmonic content. Ten simulations have a mass ratio of 4:1 and five, mass ratio of 8:1. Overall, we find that SEOBNRv5PHM is the most consistent model to numerical relativity, with the majority of true source properties lying within the inferred 90% credible interval. However, we find that none of the models can reliably infer the true source properties for binaries with mass ratio 8:1 systems. We additionally conduct inspiral-merger-ringdown (IMR) consistency tests to determine if our chosen state of the art waveform models infer consistent properties when analysing only the inspiral (low frequency) and ringdown (high frequency) portions of the signal. For the simulations considered in this work, we find that the IMR consistency test depends on the frequency that separates the inspiral and ringdown regimes. For two sensible choices of the cutoff frequency, we report that IMRPhenomXPHM can produce false GR deviations. Meanwhile, we find that IMRPhenomTPHM is the most reliable model under the IMR consistency test. Finally, we reanalyze the same 35 simulations, but this time we incorporate model accuracy into our Bayesian inference. Consistent with the work in Hoy *et al.* Incorporating model accuracy into gravitational-wave Bayesian inference, *Nat. Astron.* **9**, 1256 (2025)., we find that this approach generally yields more accurate inferred properties for binary black holes with less biases compared to methods that combine model-dependent posterior distributions based on their evidence, or with equal weight.

DOI: [10.1103/rkgr-psrt](https://doi.org/10.1103/rkgr-psrt)

I. INTRODUCTION

With a decade of observations, $\mathcal{O}(100)$ confirmed gravitational-wave (GW) detections [1–5] and $\mathcal{O}(200)$ potential GW candidates [6] under “its belt,” GW astronomy has matured into an established branch of astronomy. At the forefront of this endeavor are the GW interferometers [7–9] (IFOs) operated by the LIGO-Virgo-KAGRA Collaboration (LVK), with plans for an additional detector in India [10]. An overwhelming majority of the detected events have been sourced by the inspiral, merger, and ringdown (IMR) of binary black hole systems (BBHs) [1,11–13]. Assuming vacuum general relativity and negligible orbital eccentricity

at the time the GW signal enters the sensitive region of the IFOs, ca. 20 Hz, each system should be fully characterized by 15 parameters: eight intrinsic (two masses and two Euclidean spin vectors) and seven extrinsic (sky position, orbit orientation, time, and phase shifts).

Gravitational-wave Bayesian analyses are crucial for extracting source properties from noisy data. Typical Bayesian analyses rely on stochastic sampling of the 15-dimensional parameter space [14] with $\sim \mathcal{O}(10^7)$ waveform evaluations per analysis. This implies that fast and faithful GW models are needed. The faithfulness of a GW model is gauged with respect to state of the art numerical relativity (NR) simulations, such as those produced by the SXS [15–18], BAM [19–21], RIT [22–25], NINJA [26], NRAR [27], MAYA Collaborations [28,29] as well as the GR-Athena++ catalog [30] and the works of Refs. [31]. Each NR simulation is expensive to produce, taking days to months to complete on a high performance computing clusters. This means that direct use of NR waveforms is not

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

feasible for GW data analysis (although see, e.g., Ref. [32]). Instead, waveform models based on analytical, semianalytical and phenomenological methods are employed in parameter estimation analyses. Such methods allow waveform generation to be fast at the expense of faithfulness to NR simulations, which may cause biases in parameter estimation.

It is customary to perform injection and recovery Bayesian analyses to assess when a waveform model may produce biased estimates for the true underlying source properties (see, e.g., Refs. [33–36]). These are procedures whereby a template waveform, usually from an NR simulation, is injected into a single or multiple IFOs and the resulting signal analyzed with different waveform models through Bayesian methods [see, e.g., [37]]. In order to solely focus on systematic biases induced by waveform models, the template waveforms are generally injected into zero detector noise, though the dependence of the likelihood on detector noise is retained (see Sec. III). Such studies usually involve a few injections as they are computationally expensive.

Here, we undertake a systematic survey of waveform model recovery for numerical relativity injections, prioritizing strongly precessing BBHs. A similar injection campaign was conducted in Ref. [38] to assess the systematic biases in the waveform models used in GW150914. That particular study focused mostly on the injection-recovery of nonprecessing systems. Our singling out strongly precessing cases for our campaign is motivated by:

- (i) Waveform systematics for nonprecessing, quasicircular binary black holes are well “under control” with the state of the art models achieving 10^{-3} level of faithfulness (median) to numerical relativity simulations [39–46]. Models older than these were found to cause no systematic biases in the recovery of the systems with injected parameters similar to GW150914 [38], i.e., either no precession or weak precession [37].
- (ii) Precessing waveform models now attain $\lesssim 10^{-2.5}$ faithfulness (median) for systems with weak to mild precession (see, e.g., Ref. [47]).
- (iii) Last, although they are thought to be rare [48,49], (vacuum) binary black hole systems exhibiting strong spin precession are astrophysically interesting as they may point to a different formation channel [50–55]. Therefore, the understanding of the expected level of bias in the inferred source properties for these systems will have implications on our ability to infer the binary black hole history.

So far, GW200129_065458 has been the only observation which exhibits spin-precession under the assumption of circular orbits [1,56–58]. The rarity of detecting GW200129_065458-like systems has been quantified by Ref. [49] to be once out of every $\lesssim 50$ events. At the time of GW200129_065458, it was shown in Ref. [56] that the numerical relativity surrogate model, NRSUR7DQ4 [59], was

the least likely to produce biased parameter estimates for a strongly precessing BBH with dimensionless primary spin magnitude of 0.9 and spin tilt angle of nearly 90° (possibly due to incorporating multipole-asymmetries [60]) also confirmed recently by Ref. [58]. N.B.: the consensus on the detectability of spin precession for this event is not universal and can change depending on the assumptions made [61–64].

Therefore, our goal is to assess the performance of the latest precessing waveform models developed since 2020, specifically, SEOBNRv5PHM [36] (hereafter SEOB), IMRPhenomTPHM (thenceforth TPHM) [34] and IMRPhenomXPHM [33,65] (henceforward XPHM), for Bayesian analyses. Given that the primary investigation of the only published O4 event, GW230529_181500, uses SEOBNRv5PHM and IMRPhenomXPHM [66], our choosing of these two particular waveform models is appropriate (also motivated by Ref. [67]). We also include IMRPhenomTPHM in our work as it represents a relatively fast time-domain model, which bridge the gap between the time-domain effective one-body [68,69], and the frequency-domain phenomenological approaches [70,71] both in terms of accuracy [47] and speed. Finally, we also employ IMRPhenomXO4a [72] for some selected cases for which XPHM exhibits poor recovery performance. Note that we work with the version of XPHM with the Spin-Taylor prescription for the spin dynamics [65], which is more accurate than its MSA version [33,73].

Our assessment involves the injection of 35 numerical relativity waveforms from the SXS¹-BAM databases [16,19] into zero-noise two-detector advanced LIGO network with design O4 amplitude spectral densities [74] (ASD). The simulations are divided into three subsets of size 10 and a fourth subset of size five. The 10-element subsets are separated by their mass ratios: 1:1, 2:1, 4:1 with the five-element subset containing simulations with mass ratio 8:1. We have chosen simulations that exhibit strong precession, with extrinsic parameters chosen such that a quarter of the accumulated network signal to noise ratio (SNR) is purely due to the precession part of the waveform [75]. This ensures that the precession signal is not diminished due to the orbital configuration of the BBHs with respect to the detector network.

Another important assessment of waveform systematics is the inspiral-merger-ringdown (IMR) consistency tests [76–78], whereby the final mass and spin of the merger-product black hole is computed from both the inspiral (premerger) and the ringdown (postmerger) phases independently. Assuming general relativity to be the correct theory of gravity, a robust waveform model should produce consistent results for the inferred values of the final mass and the spin. This test is routinely performed on GW events [78–80]. Here, we perform a systematic investigation of the

¹The updated and enlarged 2025 SXS catalog [18] became available as we were completing this work. Therefore, we employ waveforms from the 2019 catalog.

waveform model robustness under the IMR consistency tests over our sample of 35 numerical relativity signals. As the IMR consistency tests are computationally expensive, we employ the fastest of the three waveform models considered, XPHM, for the entire sample set. We reperform the IMR tests with SEOB, TPHM and IMRPhenomXO4a for cases in which XPHM shows noteworthy biases. We then repeat the IMR consistency test using a different frequency to distinguish between the inspiral and ringdown phases to assess how this change affects our results. We again employ XPHM for all 35 cases, and perform supplementary runs with the other waveform models as needed. As far as we are aware, our undertaking here is the most systematic and comprehensive application of the IMR consistency test.

This work provides us with model-dependent posterior distributions that can be combined into multimodel posteriors to marginalize over waveform uncertainty. The standard way of doing is to assign equal weights to each model's posterior as has been done routinely by the LVK [1]. Alternatively, Ref. [81] suggested combining the posteriors based on the Bayesian evidence obtained by each waveform model. We will refer to this as the evidence-informed method. In Ref. [82], we proposed yet another method for combining multimodel posteriors based on each waveform model's faithfulness to NR in the parameter space explored by the stochastic sampler(s). The large number of posteriors generated by our study allows us to make a thorough comparison between these three methods for combining posteriors. As the method in Ref. [82], dubbed NR-informed, is currently restricted to BBHs between mass ratios of 1:1 to 4:1, we do not apply it to the five BAM injections of mass ratio 8:1. To our knowledge, our work here is the *first* large-scale systematic study of the latter two methods for combining posteriors from multiple precessing waveform models.

This article is organized as follows: in Sec. III, we present a general introduction to injection-recovery studies and an overview of the specific choices and settings for the runs that we performed. Section IV showcases the results of the recoveries by SEOBNRv5PHM, IMRPhenomTPHM, and IMRPhenomXPHM for the 35 numerical relativity injections. Section V reports on the results of our IMR consistency tests, performed on IMRPhenomXPHM for all 35 simulations with additional results presented for all three models for eight cases. Section VI compares the results of our NR-informed parameter estimation runs with those obtained via the evidence-informed and standard methods. We summarize our most important findings in Sec. VII and discuss the implications. All the relevant parameter, data and configuration files have been collected in a GIT repository which can be accessed at [83].

II. PHYSICAL SETUP AND NOTATION

A BBH system comprised of two compact objects with masses m_1 and m_2 has total mass $M = m_1 + m_2$, chirp

mass $M_c = (m_1 m_2)^{3/5} / M^{1/5}$ and symmetric mass ratio $\eta = m_1 m_2 / M^2$. The ratio of binary masses can either be ≥ 1 , defined by the large mass ratio $Q := m_1 / m_2$, or ≤ 1 , defined by the small mass ratio $q = 1 / Q$. We quote all masses and frequencies in the detector frame. Source-frame masses are related to their detector frame quantities via $\text{mass}_s = (1 + z)^{-1}(\text{mass})$, where z is the cosmological redshift of the source. We use geometrized units $G = c = 1$ except when quoting values for the total binary mass which we do in units of solar masses and the GW frequency f which will be given in Hertz.

The orbital angular momentum of the binary is given by $\mathbf{L} = L \hat{\mathbf{L}} = \mathbf{r} \times \mathbf{p}$ where $\mathbf{r}$, $\mathbf{p}$ are the relative separation and momentum of the binary respectively, and $L = |\mathbf{L}|$. The black hole spins are represented by the Euclidean three-vectors $\mathbf{S}_i$ which have magnitudes $|\mathbf{S}_i| = m_i^2 a_i$ with $0 \leq a_i < 1$ for $i = 1, 2$, where we deliberately avoid the extremal Kerr spin limit $a_i = 1$. θ_1, θ_2 denote the spin tilt angles with respect to a reference axis which is taken to be the direction of the Newtonian orbital angular momentum vector at some reference frequency, e.g., 20 Hz [84]. Similarly, ϕ_1, ϕ_2 represent the azimuthal spin angles in the orbital plane with $\Delta\phi := \phi_2 - \phi_1$.

The phenomenology of spin precession in compact binary inspirals was thoroughly exposed in the seminal work of Ref. [85]. Using the large disparity between the orbital, precession and radiation reaction timescales, and the assumption of conserved total angular momentum, $\mathbf{J} = \mathbf{L} + \mathbf{S}_1 + \mathbf{S}_2$, a picture emerges whereby $\mathbf{L}, \mathbf{S}_1, \mathbf{S}_2$ all precess around $\mathbf{J}$, tracing out cones over a precession period. We show the two-dimensional projection of two such cones for the case of the $Q = 2$ binary black hole simulation (SXS:BBH:1200) [16] in the left panel of Fig. 1 (see also, e.g., Fig. 1 of Ref. [86]). The relevant parameters of this system can be found in Tables I and VI. In the right panel of the figure, we show the GW signal produced by this merging system as seen in the $\mathbf{L}$ frame and in the $\mathbf{J}$ frame. The $\mathbf{L}$ frame is standard for NR, and in the $\mathbf{J}$ frame the amplitude modulations due to precession are in general less pronounced. However, this is hard to discern for this particular case because the precession cone opening angles for $\mathbf{L}$ ($\mathbf{J}$) are small with respect to their frames $\mathbf{L}_0$ ($\mathbf{J}_0$). The differences between the two waveforms become more pronounced when the spins point mostly in the opposite direction to $\mathbf{L}_0$ (see, e.g., Fig. 3 of Ref. [87].)

Often it is convenient to describe the individual spin vectors by effective inspiral and precessing spins [88–92]. The effective inspiral spin is defined as

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

and the effective precession spin




FIG. 1. Left panel: the precession of $\mathbf{S}_1$ and $\mathbf{S}_2$ around the total angular momentum vector $\mathbf{J}$ which points out of the page through the origin. As the spin vectors precess, they trace out precession cones around $\mathbf{J}$ whose projections are plotted in the figure as the blue and orange curves. The small fluctuations in the trajectories are due to nutation. As the mass ratio of this system (SXS:BBH:1200) is 2:1 and the dimensionless spins are equal, the magnitude of $\mathbf{S}_1$ is approximately $2^2 = 4$ times larger than the magnitude of $\mathbf{S}_2$. The dots mark the initial projections of the spin vectors. Right panel: the numerical relativity waveform strain $\text{Re}(h_{\text{NR}}(t))$ for this simulation seen both in the $\mathbf{L}$ frame (top row) and the $\mathbf{J}$ frame (bottom row). In this figure, the time units have been geometrized (adimensionalized) via $G = c = M = 1$. For $M = 53.6M_\odot$ given for this system in Table I, a time interval of $4000M$ converts to 1.06 second. Though the waveforms in the two frames look very similar, they are not identical as can be seen in the lower right subfigure where we overlaid the orange $\mathbf{L}$ -frame merger-ringdown waveform on top of its $\mathbf{J}$ -frame counterpart (blue).

TABLE I. The most relevant intrinsic parameters for the 30 sxs binary black hole simulations used throughout this work. Column one lists the sxs simulation code with columns two and three displaying the total system mass (in $M_\odot$) and the mass ratio, respectively. Columns four and five give the total dimensionless spin magnitudes for the black holes while columns six, seven, and eight list the spin tilt angles and the relative azimuthal separation of the spin vectors.

SXS:BBH	$M(M_\odot)$	Q	a_1	a_2	θ_1	θ_2	$\Delta\phi$
0764	51.5	1.0	0.80	0.80	92.9°	83.5°	1.53
0767	51.3	1.0	0.80	0.80	87.3°	92.0°	3.66
0838	51.5	1.0	0.80	0.80	93.0°	83.5°	1.56
0841	51.3	1.0	0.80	0.80	87.1°	92.1°	3.68
0925	51.7	1.0	0.80	0.80	83.2°	92.7°	4.76
00935	51.3	1.0	0.80	0.80	89.7°	90.2°	2.60
0965	51.7	1.0	0.80	0.80	82.8°	91.3°	5.24
0982	51.3	1.0	0.80	0.80	93.3°	85.0°	2.06
1205	51.4	1.0	0.85	0.85	84.4°	93.4°	4.19
1217	51.8	1.0	0.85	0.80	86.2°	86.2°	6.28
0716	53.6	2.0	0.80	0.80	86.2°	84.8°	0.06

(Table continued)

TABLE I. (Continued)

SXS:BBH	$M(M_\odot)$	Q	a_1	a_2	θ_1	θ_2	$\Delta\phi$
0717	53.7	2.0	0.80	0.80	86.0°	96.0°	4.30
0812	53.5	2.0	0.80	0.80	89.1°	91.3°	3.24
0814	53.9	2.0	0.80	0.80	83.9°	92.8°	5.36
0926	53.5	2.0	0.80	0.80	84.8°	95.7°	4.84
0936	53.5	2.0	0.80	0.80	90.9°	86.9°	2.69
0966	53.1	2.0	0.80	0.80	85.5°	94.5°	5.23
0976	53.4	2.0	0.80	0.80	87.8°	93.9°	3.73
1197	52.7	2.0	0.85	0.85	92.5°	83.1°	2.04
1200	53.6	2.0	0.85	0.85	85.9°	84.3°	0.06
1916	58.3	4.0	0.80	0.80	87.5°	98.7°	4.18
1921	59.0	4.0	0.80	0.80	87.6°	92.8°	3.28
1922	58.3	4.0	0.80	0.80	88.9°	84.3°	2.36
1923	59.1	4.0	0.80	0.80	83.4°	88.4°	5.97
2000	59.8	4.0	0.80	0.80	84.0°	97.7°	5.13
2004	58.2	4.0	0.80	0.80	88.4°	96.3°	3.77
2070	58.4	4.0	0.80	0.80	89.5°	86.3°	2.61
2074	59.1	4.0	0.80	0.80	84.7°	83.1°	0.23
2075	58.6	4.0	0.80	0.80	86.8°	95.2°	5.53
2079	58.9	4.0	0.80	0.80	87.3°	98.7°	4.07

$$\chi_p = \max \left(a_1 \sin \theta_1, q \frac{4q+3}{4+3q} a_2 \sin \theta_2 \right). \quad (2)$$

These obey the bounds $-1 < \chi_{\text{eff}} < 1$ and $0 \leq \chi_p < 1$. Additionally, we can define the effective parallel spin [82]

$$\chi_{\parallel} := \frac{1}{(1+q)^2} (a_1 \cos \theta_1 + q^2 a_2 \cos \theta_2) \quad (3)$$

and the effective perpendicular spin [47,82,87]

$$\chi_{\perp} = \frac{|\mathbf{S}_{1,\perp} + \mathbf{S}_{2,\perp}|}{M^2}, \quad (4)$$

where $\mathbf{S}_{i,\perp} = a_i m_i^2 (\sin \theta_i \cos \phi_i, \sin \theta_i \sin \phi_i, 0)^T$ for $i = 1, 2$.

The masses and spins combine to make up the eight intrinsic parameters, λ_{int} , of a given BBH. There are additionally seven extrinsic parameters in the case of quasicircular (quasispherical) inspirals: the luminosity distance d_L , the right ascension and declination angles $\{\alpha, \delta\}$, an overall constant time shift t_{ref} , and polar (inclination), azimuthal, and polarization angles $\{\iota, \varphi_{\text{ref}}, \psi\}$ parametrizing the orientation of the orbit and the polarization of the GWs with respect to the detector frame.

With the relevant parameters of the system specified, the gravitational-wave strain in the time domain can be written in its multipolar form as

$$h(t) = \frac{1}{d_L} \sum_{\ell, m} h_{\ell m}(t, \lambda_{\text{int}}) {}_{-2}Y^{\ell m}(\iota, \varphi_{\text{ref}}), \quad (5)$$

where ${}_{-2}Y^{\ell m}$ are spin-weighted spherical harmonics and $h_{\ell m}$ are the GW multipoles. The three models of interest here use similar prescriptions to frame-rotate the coprecessing multipoles $h_{\ell m}^{\text{coprec}}$ to $h_{\ell m}$ using the orientation of $\mathbf{L}(t)$ with respect to $\mathbf{J}_0$ during the inspiral to determine these Euler angles [93–95]. The models use more differing approximations to extend the Euler angles into the merger-ringdown regime. The coprecessing multipoles can be modeled using the aligned-spin (nonprecessing) multipoles as is done, e.g., for TPHM and XPHM. SEOB has a more sophisticated prescription whereby partial precessing-spin information is incorporated into the EOB Hamiltonian used in the coprecessing frame [96]. An important commonality in all three models is that they all employ the so-called multipole symmetry for their coprecessing multipoles, i.e.,

$$h_{\ell, -m}^{\text{coprec}} = (-1)^{\ell} [h_{\ell m}^{\text{coprec}}]^*, \quad (6)$$

where $*$ denotes complex conjugation.

III. A PRIMER ON INJECTION/RECOVERY STUDIES

The accuracy of a GW model is commonly assessed by performing Bayesian inference on a simulated GW signal, h , produced from a coalescing binary of known parameters λ_{inj} [see, e.g., [36,59,65,72]]; Bayesian inference is the process of estimating a model-dependent *posterior probability distribution*, which describes the probability of the binary having a specific set of properties $\lambda = \{\lambda_1, \lambda_2, \dots, \lambda_N\}$ given the observed data d and model $\mathfrak{M}$ [14,97]. The model-dependent posterior distribution is calculated through Bayes' theorem as,

$$p(\lambda|d, \mathfrak{M}) = \frac{\Pi(\lambda|\mathfrak{M})\mathcal{L}(d|\lambda, \mathfrak{M})}{\mathcal{Z}}, \quad (7)$$

where $\Pi(\lambda|\mathfrak{M})$ is the probability of the parameters λ given the model $\mathfrak{M}$, otherwise known as the prior; $\mathcal{L}(d|\lambda, \mathfrak{M})$ is the probability of observing the data given the parameters λ and model $\mathfrak{M}$, otherwise known as the likelihood; and $\mathcal{Z}$ is the probability of observing the data given the model $\mathcal{Z} = \int \Pi(\lambda|\mathfrak{M})\mathcal{L}(d|\lambda, \mathfrak{M})d\lambda$, otherwise known as the evidence.

For the case of GW astronomy, the likelihood is known and depends on the specified GW detector network and individual detector sensitivities [14,97]—typically characterized by the power spectral density (PSD) (the square of the ASD). Under the assumption of Gaussian and stationary noise, the likelihood is simply²

$$\mathcal{L}(d|\lambda, \mathfrak{M}) = \prod_k \exp \left(-\frac{1}{2} (d_k - \mathfrak{M}(\lambda) | d_k - \mathfrak{M}(\lambda)) \right), \quad (8)$$

where k denotes the individual detectors in the network, $(a|b)$ refers to the inner product between two frequency series $a(f)$ and $b(f)$,

$$(a|b) = 4\Re \int df \frac{a(f)b^*(f)}{S_k(f)}, \quad (9)$$

and $S_k(f)$ is the PSD of the detector k . Despite the likelihood being well known, it is often not possible to analytically calculate the model dependent posterior distribution. The reason is because the evidence involves computing the likelihood and prior for all points in the parameter space. As a result, stochastic sampling techniques, such as Markov-chain Monte-Carlo (MCMC) [99] and nested sampling [100,101], were developed to return a set of independent

²The likelihood also includes an additional term describing the noise covariance. Under the assumption that the noise is Gaussian and stationary, the noise covariance matrix is the identity matrix and often excluded for simplicity. When this assumption is no longer valid, the noise covariance should be included, see, e.g., Ref. [98] for details.

draws from the unknown posterior distribution; although see, e.g., Refs. [102–111] for other approaches.

Often we wish to consider a specific configuration of GW detectors. As such, a simulated GW signal is often injected into real or synthetic GW strain data $n(t)$, or injected into “zero-noise” where $n(t) = 0 \ \forall \ t$. For either case, the likelihood for a single detector network³ reduces to,

$$\mathcal{L}(d|\lambda, \mathfrak{M}) \propto A(n) \exp(-|h - \mathfrak{M}(\lambda)|^2), \quad (10)$$

where $|\cdot| = \sqrt{\langle \cdot | \cdot \rangle}$. The likelihood can be further simplified by introducing the mismatch: a quantity which characterizes how similar two GWs are to one another. The mismatch between waveforms h and $\mathfrak{M}(\lambda)$ ranges between 0 and 1, where 0 implies the waveforms are identical (up to an overall amplitude rescaling), and 1 implies the waveforms are orthogonal. The mismatch is defined as

$$\mathcal{M} = 1 - \max_{dt, d\phi} \frac{(h|\mathfrak{M}(\lambda))}{|h||\mathfrak{M}(\lambda)|}, \quad (11)$$

where we maximize over time, dt , and phase shifts, $d\phi$. When restricting attention to a reduced subspace that neglects the phase and time, the likelihood takes the simplified form [112],⁴

$$\mathcal{L}(d|\lambda, \mathfrak{M}) \propto A(n) \exp(-|h|^2(1 - \mathcal{M}^2)). \quad (12)$$

For a numerical relativity simulation injected in zero-noise, Eq. (12) implies that the likelihood will peak in the region of the parameter space where the mismatch between the simulation and $\mathfrak{M}(\lambda)$ is minimized. For the idealized case, where $\mathfrak{M}$ perfectly describes numerical relativity, the likelihood will peak at the true parameters of the simulated GW signal, λ_{inj} . Of course, in reality we expect to observe a bias due to mismodeling; a phenomenon known as waveform systematics. For this case, the maximum likelihood (and hence the minimum mismatch between h and $\mathfrak{M}$) will be at the parameters $\lambda_{\text{max}} \neq \lambda_{\text{inj}}$. Although in general the maximum likelihood may not peak at the true parameters, the 90% credible interval of the posterior distribution may encase λ_{inj} . Likewise, given the dependence of the prior distribution in Eq. (7), the posterior distribution may not peak at λ_{inj} , unless the prior distribution is uniform in all parameters. This may not always be true, especially for derived quantities.

³The likelihood may also be simplified for a multidetector network. However, we only consider a single detector network for simplicity.

⁴It is not necessary to restrict attention to a reduced subspace that neglects the phase and time. If the full space is considered, the match, $1 - \mathcal{M}$, is replaced by the overlap. The overlap follows the same functional form as the match but does not maximize over the phase and time shifts.

If we assume the high-SNR limit, where the likelihood is Gaussian, a conservative limit (see Ref. [113] for details) for when two waveforms will be distinguishable at 90% confidence is when

$$\mathcal{M} \lesssim \frac{3.12}{\rho^2}, \quad (13)$$

where $\mathcal{M}$ is the mismatch between the waveforms h_1 and h_2 and ρ is the SNR of the signal [112, 114]. For the case of a real gravitational-wave signal h and a model evaluated at the parameters λ , we can obtain an approximate contour containing 90% of the posterior distribution by combining Eqs. (12) and (13), see, e.g., Ref. [111]. Although the 90% likelihood surface may encase the true parameters, the posterior distribution may not owing to shifts caused by the prior. For a numerical relativity simulation injected into synthetic GW strain data, we expect to observe a systematic bias even in the idealized case where the model perfectly describes numerical relativity. However, averaging the results from many analyses of the same simulation injected into different instances of synthetic GW strain data is expected to obtain the same result as a zero-noise injection.

When an ensemble of models exist, $\mathfrak{M} = \{\mathfrak{M}_1, \mathfrak{M}_2, \dots, \mathfrak{M}_i\}$, individual model-dependent posterior distributions can be compared and contrasted against the true source properties in order to determine which model more accurately describes the simulated signal (see Sec. VI for details about how to combine individual model-dependent posterior distributions to marginalize over model uncertainty). To quantify this a recovery score for each dimension λ_j can be calculated, which accounts for the width of the marginalized one-dimensional posterior as well as the injected value. The recovery score for a given model $\mathfrak{M}_i$ and simulation K is defined as

$$\mathbf{r}_K(\mathfrak{M}_i) := \frac{\sigma_{\lambda_j}(\mathfrak{M}_i)}{C(\lambda_j, \mathfrak{M}_i)}. \quad (14)$$

where σ_{λ} is the standard deviation of the one-dimensional marginalized posterior distribution defined as,

$$p(\lambda_j|d, \mathfrak{M}_i) = \int p(\lambda|d, \mathfrak{M}_i) d\lambda_1 \dots d\lambda_{j-1} d\lambda_{j+1} \dots d\lambda_N \quad (15)$$

and $C(\lambda_j, \mathfrak{M}_i)$ is a cost function given by Ref. [115] as

$$C(\lambda_j, \mathfrak{M}_i) = \int_{-\infty}^{\infty} p(\lambda_j|d, \mathfrak{M}_i) (\lambda_j - \lambda_{j,\text{inj}})^2 d\lambda_j \quad (16)$$

with $\lambda_{j,\text{inj}}$ representing the injected value. If the mean of the one-dimensional posterior equals $\lambda_{j,\text{inj}}$ then the cost function equals σ_{λ_j} , otherwise it is greater than it. This assures that $\mathbf{r} \in (0, 1]$ with values close to unity highlighting better recovery performance. For example, assuming a normal

distribution with zero mean and standard deviation of σ , we obtain $\mathbf{r} \approx \{0.71, 0.45, 0.32\}$ for $\lambda_{\text{inj}} = \{1, 2, 3\}\sigma$. Thus, $\mathbf{r} > 0.7$ indicates robust recovery within ± 1 standard deviation. Note that the particular integral (16) for the cost function can penalize narrow asymmetric (skewed) posteriors more heavily than wide symmetric posteriors. We will encounter examples of this in Sec. VI.

In this work, we perform Bayesian inference on 35 numerical relativity simulations injected into zero-noise to assess the accuracy of three state of the art precessing GW models. Throughout this work, we use a theoretical PSD for Advanced LIGO's O4 design sensitivity [74], and assume a two-detector network consisting of LIGO Hanford (H1) and LIGO Livingston (L1) [7]. Although numerous packages are now available to stochastically sample the parameter space [14,97,116,117], we use the DYNesty nested sampler [118] via BILBY [97] and BILBY_PIPE [116]. For all analyses, we use 1000 live points along with the BILBY-implemented ACCEPTANCE-WALK sampling algorithm with an average of 60 accepted steps per MCMC. We use sufficiently wide and uninformative priors for all parameters. Specifically we use the BILBY function `UniformInComponentsChirpMass` with the range $M_c \in [10, 40]M_\odot$ and the `UniformInComponentsMassRatio` function with range $q \in [0.083, 1]$ (as well as a constraint on $m_i \in [1, 1000]M_\odot$). The luminosity distance prior range is given by $d_L \in [100, 5000]$ Mpc via the function `UniformSourceFrame` using Λ CDM cosmology parametrized by Planck 2015 data [119]. The spin magnitudes are uniform in the range $\chi_i \in [0, 0.99]$ while the spin tilt angles are uniform in their sines: $\sin \theta_i \in [0, 1]$, and the **J**-frame inclination angle is uniform in cosine: $\cos \vartheta_{\text{JN},0} \in [-1, 1]$, where $\vartheta_{\text{JN},0}$ is the angle between **J** and $\hat{\mathbf{N}}$ at $f = f_0$. Other angles are all uniform in their respective ranges. We consistently integrate the likelihood between 23 – 1024 Hz, with the injected numerical relativity waveform always generated from $f_0 = 20$ Hz to avoid artefacts in the fast Fourier transforms. Although this means that we will be missing power from, e.g., the $\ell < 5$ higher order multipoles for frequencies < 50 Hz, this was chosen to maximize the power in the dominant quadrupole.

IV. RESULTS I: RECOVERY PERFORMANCE OF THE WAVEFORM MODELS

As our work presents injection-recovery results of 15 parameters by three waveform models for almost three dozen numerical relativity simulations, we calculate and compare recovery scores for each dimension. We note that the recovery score, as defined in Eq. (14), has its limitations. Namely, when the injected value lies at the edge of the parameter space, as is the case with the recovery of the mass ratio for the $Q = q = 1$ subset of simulations, the recovery score will be ≈ 0.55 as opposed $\gtrsim 0.7$ for $Q > 1$

cases. Similarly, if the underlying posterior is heavily skewed, as can happen even without posterior railing, the recovery score will be reduced.

Another limitation of Eq. (14) is that it yields unity whether the posteriors are narrow or wide as long as the sample mean coincides with the injected value. To distinguish between such possibilities, we introduce a measure of how sharply peaked the posterior is via a quantity that we dub the recovery width

$$\mathbf{w} := \frac{\Sigma_{90}(p(\lambda_i|d, \mathfrak{M}_i))}{\max(\Pi(\lambda_i)) - \min(\Pi(\lambda_i))} \quad (17)$$

with $\Sigma_{90}(p(\lambda_i|d, \mathfrak{M}_i))$ representing the width of the 90% credible interval of the one-dimensional marginalized posterior distribution for the selected parameter λ_i . We divide this quantity by the length of the interval for the prior $\Pi(\lambda_i)$ of each parameter. Note that given the chirp mass and mass ratio prior ranges of $[10M_\odot, 40M_\odot]$ and $[0.083, 1]$, we obtain a rather wide range for the total mass: $\Pi(M) \in [23M_\odot, 196M_\odot]$. As a result, the recovery width values for the total mass may appear small. However, we wish to keep the denominator of Eq. (17) the same for a given parameter as we increase the mass asymmetry of the systems. This also raises the question of what constitutes a good value for the recovery width which is not so clear as in the case of the recovery score. In general, it makes sense to compare values of $\mathbf{w}$ for the same parameter between different models or as the mass ratio is changed, rather than comparing the values of $\mathbf{w}$ between two different parameters.

To further aid in comparing model performance, we additionally introduce

$$\mathbf{r}_{\text{av}}(\mathfrak{M}_i) = \frac{1}{N} \sum_K \mathbf{r}_K(\mathfrak{M}_i), \quad (18)$$

where K represents the set of N numerical simulations to average over. The important thing to keep in mind when looking at the values of $\mathbf{r}$ and $\mathbf{w}$ is their magnitude for a given model relative to another model or for a given subset of simulations, e.g., $Q = 1$, vs another subset, e.g., $Q = 2$.

A. Cases with mass ratio between 1:1 and 4:1

We analyzed a subset of 30 numerical relativity simulations produced by the SXS collaboration with mass ratios between 1:1 and 4:1. We selected these simulations according to the following criteria. We start with the requirement that the number of orbital cycles is > 18 , residual NR eccentricity less than 10^{-3} , $|a_i| > 0.79$ and $1 < \cos^{-1} \theta_i < 2$. This yields 18 simulations with $Q < 1.2$, 57 with $1.2 < Q < 2.2$ and 24 with $Q > 2.2$. We reduce the size of the $1.2 < Q < 2.2$ portion by limiting the corresponding tilt angles to be within $\pm 1\%$ of the range of the corresponding $Q < 1.2$ angles while also restricting the range of $\Delta\phi$ to be 5%–95% of the range of the $Q < 1.2$




FIG. 2. One-dimensional marginalized posterior distributions obtained for the inferred total detector-frame mass (first column), mass ratio (second column), effective parallel (third column) and effective perpendicular spins (fourth column) for the 30 sxs binary black hole simulations used throughout this work. The red vertical lines indicate the true values.

values. This yields 10 cases with $1.99958 \leq Q \leq 2.00011$. We obtain 10 cases with $3.99911 \leq Q \leq 4.00046$ in a similar fashion, but without needing to restrict $\Delta\phi$. Finally, we relax the $\Delta\phi$ bound enough to obtain 10 cases with $1.00001 \leq Q \leq 1.00006$. These steps produce a dataset that has very similar parameters in the dimensionless spin space, but this is deliberate as we want to focus on keeping the strength of precession more or less fixed while we increase the mass asymmetry. Finally, we checked that none of our chosen simulations have been listed as deprecated in the updated SXS catalog [18].

The most relevant intrinsic parameters for the resulting 30 BBHs are given in Table I. The values for the detector-frame binary total mass are chosen such that the detector-frame reference GW frequency of the SXS simulation is 20 Hz. Note that all dimensionless spin magnitudes are set to the same high value, i.e., $a_1 = a_2 = 0.8$, but the spin tilt angles are slightly different in each case though they are all between 80° and 100° , resulting in systems with rather strong precession, i.e., $\chi_p > 0.75$. The values of $\{\iota, \alpha, \delta, \psi, d_L\}$ for each of the 30 BBHs are chosen such that the total network SNR, ρ_{tot} , is 40, with the precession SNR [75], ρ_p , equalling 10. The specific values of the extrinsic parameters can be found in Table VI in the Appendix where it can be seen that many different combinations of extrinsic parameters can yield $\rho_p = 0.25\rho_{\text{tot}}$.

In Fig. 2, we present the one-dimensional marginalized posterior distributions for $\{M, q, \chi_{\text{eff}}, \chi_p\}$ as recovered by SEOB, TPHM, and XPHM. The figure contains 30 rows divided into subsets of 10 by mass ratio and four columns, one for each parameter. The top (bottom) 10 rows show the $q = 1(1/4)$ results with the middle 10 displaying the $q = 1/2$ posteriors. The number for each SXS simulation is at the left end of each row. We summarize our key findings below.

1. Mass ratio 1:1 systems

For the $Q = q = 1$ simulations, all three models recover the detector-frame total mass within their respective 90% credible intervals (CIs), with SEOB's posteriors generally peaking closer to the injected values. Similarly, the models recover the injected value of the mass ratio, i.e., $q = 1$ well modulo the usual issues with having the injected value at the boundary of the parameter space. Interestingly, for SXS:1217 TPHM's posteriors completely miss the true value. However, because TPHM severely misidentifies the mass ratio of the binary, while inferring M well, TPHM obtains biased estimates for the components masses m_1 and m_2 . It is not clear why the model behaves this way for this particular simulation. We checked the signal power in the subdominant multipoles [120], but found that $\rho_{44} \approx 1.5$ with other multipole SNRs less than one. As a further check, we re injected this simulation with the SXS:1205's extrinsic parameters (ensuring a fixed total SNR, but varying precession SNR) and recovered posteriors that

are consistent, albeit broader, with the injection. We also conducted the converse experiment, i.e., re injecting SXS:1205 with SXS:1217's extrinsic parameters. This time, we observed the expected posterior railing against $q = 1$ as opposed to what Fig. 2 shows for SXS:1217. As a final check, we re performed the PE analysis with more aggressive settings to ensure that the posterior distribution is converged (e.g., twice as many live points, see the config file in our public data release). We found that the resulting posteriors are statistically indistinguishable. Moreover, when plotting the 90% credible interval for the reconstructed [whitened] waveforms we find that the recovered TPHM posterior overlaps better with the injected SXS simulation than when generating the injection with TPHM (i.e., when passing the injected parameters to the TPHM waveform model). Additionally, the maximum-likelihood TPHM waveform yields a marginally larger match of 0.976 to the injected SXS simulation, compared to 0.975 when matching the injection generated with TPHM against the injected SXS simulation. Therefore, we conclude that this unexpected recovery result by TPHM is due to specific combination of intrinsic and extrinsic parameters exposing a well-known fact of GW data analysis: the intrinsic parameters for the maximum-likelihood waveform need not agree well with the injected intrinsic parameters. This is a generic feature of the stochastic nature of the data analysis methods, and not limited to just one waveform model.

The models show more variability in their recovery of χ_{eff} with SEOB always recovering the true source properties within its 90% CI, and peaking close to the true value half the time. Interestingly, TPHM and XPHM seem to recover χ_{eff} worse when SEOB recovers it near its peak as can be seen, e.g., in the cases of SXS:0764, 0838, 0965. These simulations all share very similar values for θ_1, θ_2 that yield $|\chi_{\text{eff}}| \lesssim 0.05$. It is also noteworthy that SEOB's performance degrades for a similar configuration, SXS:0982, compared to, e.g., SXS:0838. This turns out to be caused by the railing of SEOB's posteriors for a_1 against zero, thus missing the true value of 0.8 by more than three standard deviations (SDs or σ 's). Therefore, the resulting SEOB posterior favors $\chi_{\text{eff}} > 0$ at 90% confidence. Nonetheless, this is marginally better than the output of TPHM and XPHM.

The models' recovery performance for χ_p is rather complicated as this depends on the recovery of $\{m_1, m_2, a_1, a_2, \theta_1, \theta_2\}$. Let us first focus on the first seven $Q = 1$ cases, i.e., SXS:0764 to 0965 for which all three models behave the same way. We see that they seem to alternate between either recovering the injected values close to their peaks or underestimating it by one to two SDs. A deeper look reveals that all three models recover a_1 and a_2 poorly for 767 and 841, but produce narrow posteriors for θ_1, θ_2 . This means that $\chi_{\text{eff}} \approx 0$ can still be well recovered while χ_p recovery is poor. The only exception seems to be SXS:0935 where the posteriors for $\{a_1, a_2, \theta_1, \theta_2\}$ contain the injected values within their 90% CIs. However, the posteriors for θ_1, θ_2 are much

TABLE II. The recovery score (14) of each waveform model for the parameters listed in column one. Columns two to four list the averages over the entire set of 30 simulations for SEOB, TPHM, XPHM, respectively labeled as s , t , x here. The remaining columns show the average of the same quantity over the size- $10Q=1, 2, 4$ subsets. A value of $r > 0.7$ indicates a recovery of the injected parameter within one standard deviation of the mean of the posterior distribution. Note that the lower recovery score for the mass ratio q for the $q=1$ subset is due to the fact that the injected value lies at the boundary of the domain hence causing the posteriors to rail as can be seen in Fig. 2. The closer the recovery score to unity, the better.

Θ	r_{av}^S	r_{av}^T	r_{av}^X	$r_{av}^S[Q=1]$	$r_{av}^T[Q=1]$	$r_{av}^X[Q=1]$	$r_{av}^S[Q=2]$	$r_{av}^T[Q=2]$	$r_{av}^X[Q=2]$	$r_{av}^S[Q=4]$	$r_{av}^T[Q=4]$	$r_{av}^X[Q=4]$
M	0.834	0.723	0.625	0.917	0.699	0.733	0.793	0.740	0.686	0.793	0.730	0.456
q	0.756	0.708	0.593	0.555	0.531	0.556	0.881	0.783	0.712	0.831	0.810	0.512
χ_{eff}	0.807	0.532	0.718	0.834	0.518	0.559	0.786	0.720	0.872	0.803	0.358	0.724
χ_p	0.692	0.512	0.521	0.698	0.634	0.645	0.676	0.538	0.607	0.702	0.365	0.310

TABLE III. Same as Table II, but for the recovery width w . Smaller recovery widths indicate narrower posterior distributions. As the denominator of Eq. (17) yields much larger values for M than $\{q, \chi_{eff}, \chi_p\}$ relative to the numerator, the resulting recovery widths for M may be interpreted to be artificially small. An alternative is replacing the denominator with mean $\langle p(M) \rangle$ which yields values of $w \approx \mathcal{O}(10^{-2})$ for the total mass.

Θ	w_{av}^S	w_{av}^T	w_{av}^X	$w_{av}^S[Q=1]$	$w_{av}^T[Q=1]$	$w_{av}^X[Q=1]$	$w_{av}^S[Q=2]$	$w_{av}^T[Q=2]$	$w_{av}^X[Q=2]$	$w_{av}^S[Q=4]$	$w_{av}^T[Q=4]$	$w_{av}^X[Q=4]$
M	0.00326	0.00322	0.00354	0.00330	0.00306	0.00349	0.00343	0.00351	0.00386	0.00304	0.00309	0.00327
q	0.614	0.607	0.667	0.623	0.577	0.658	0.647	0.662	0.727	0.573	0.583	0.616
χ_{eff}	0.282	0.278	0.306	0.286	0.265	0.302	0.297	0.304	0.334	0.263	0.267	0.283
χ_p	0.563	0.557	0.612	0.571	0.529	0.603	0.593	0.607	0.667	0.525	0.534	0.565

broadener in this case than the other six cases, hence the poor χ_p recovery. We have also noticed that for this particular case TPHM yields posteriors that strongly peak at $\theta_1 \approx 2.5$, $\theta_2 \approx 0.5$, thus further skewing its χ_p posterior⁵ TPHM additionally produces divergent results for SXS:0982, a simulation where all models struggle with recovery. Interestingly, XPHM marginally outperforms SEOB for this case. A detailed investigation reveals that TPHM poorly recovers each of the four parameters in the set $\{a_1, a_2, \theta_1, \theta_2\}$ with SEOB and XPHM recovering $\{a_2, \theta_1, \theta_2\}$ reasonably well. SXS:0982 is most similar to SXS:0838 when it comes to its intrinsic parameters so it is puzzling that SEOB should perform so differently between these two simulations whereas TPHM and XPHM behave somewhat consistently.

Recalling that the models considered here obey the multipolar symmetry (6), an explanation for the poor χ_p recovery for at least SXS:0767,0841 and possibly SXS:0935 is the fact that these simulations have their values of $\Delta\phi$ closest to π as listed in Table I. Mass ratio 1:1 systems with spins purely in the orbital plane pointing in the opposite directions, i.e., $\Delta\phi = \pi$, are well known to be in the so-called superkick configuration [121–123], having the largest amount of multipole asymmetry.

We can quantify some of the above statements further via the values for the recovery score r and recovery width w presented in Tables II and III, specifically columns five

through seven, showing the averaged values for the $Q = 1$ subset of simulations. We see that SEOB's average r exceeds 0.9 for M corroborating what Fig. 2 shows. For the mass ratio q , we obtain very similar values of r_{av} from all three models with TPHM's score slightly worse because of the aforementioned issues related to SXS:1217. As TPHM and XPHM χ_{eff} posteriors are very similar, the resulting r_{av} 's are in the range 0.5–0.6 with SEOB's exceeding 0.8, indicating very good recovery performance. Finally, since all three models produce very similar χ_p posteriors for seven of the $10Q = 1$ simulations, the resulting r_{av} are also similar with some differences due to model performance for SXS:0982, 1205, and 1217.

The recovery widths of the model posteriors are in general quite comparable for the $Q = 1$ subset with TPHM (XPHM) producing the narrowest (widest) posteriors for all four parameters considered in Table III. However, recall that a small recovery width is only meaningful when coupled with a good recovery score. For example, though TPHM has the narrowest posteriors for χ_{eff} , it yields $r_{av} = 0.518$, much lower than SEOB's value of 0.834.

2. Mass ratio 2:1 cases

Turning our attention to the $Q = 2 (q = 1/2)$ subset, we return to Fig. 2 but now focusing on the middle 10 panels. In general, the models recover M and q within their 90% CIs with the exception of SXS:0814 where all three models fail to recover these parameters. The intrinsic spin parameters of this system are most similar to the $Q = 1$ SXS:0925 case,

⁵The 1D posteriors for all the system parameters can be seen on our aforementioned GIT repository.

where we had recorded robust recovery of $\{M, q\}$ by all three models. Therefore, it may be another case of the extrinsic parameters conspiring to make this simulation more challenging. Indeed, we find that when we reinject SXS:814 with SXS:925's extrinsic parameters (again ensuring the same network SNR but varying precession SNR) for another recovery run with XPHM, the resulting posteriors for q shift close enough to the injected value $q = 1/2$ that the recovery is now within 1σ , albeit with broader posteriors than before. More quantitatively, the q recovery score jumps from $r = 0.50$ to 0.78 , and from 0.38 to 0.66 for M .

Beside SXS:0814, XPHM also fails to recover $\{M, q\}$ within its 90% CI for SXS:0976 and SXS:1200. It significantly underestimates M while overestimating q . This is not unusual as the best recovered parameter is often the chirp mass whose [marginalized] 2D posteriors form a narrow banana-shaped region in the $m_1 - m_2$ space with very similar likelihoods along the “banana” (for high mass, short duration simulations, the total mass is often the best recovered parameter due analyses only being sensitive to the merger and ringdown portions of the signal). Therefore, a model recovering $\mathcal{M}_c$ within $\pm 1\sigma$ can yield $\{M, q\}$ posteriors due to an under(over)estimation of m_1 (m_2) such as the ones mentioned here.

The values in Table II corroborate these findings with SEOB yielding the best recovery scores for both M and q , followed by TPHM then XPHM. SEOB also produces the narrowest M, q posteriors with TPHM the broadest as given by the recovery widths in Table III.

The recovery of χ_{eff} is in general quite robust among the models. XPHM has the highest average recovery score which can also be seen in Fig. 2. SEOB's recovery is similar to XPHM's, but ever so slightly worse in some cases, thus its recovery score being about 10% lower. TPHM's performance is the least consistent with the model sometimes recovering very well with $r \approx 1$, but also not recovering the injections within the 90% CI for four of the 10 cases.

It is difficult to draw any general conclusions about the models' recovery performance for spin precession. What is clear from Fig. 2, even at a mass ratio of 2:1 and a precession SNR of 10, is that all three models struggle to produce a consistently faithful recovery of the injection. It also seems that time-domain SEOB and TPHM models behave more similarly to each other than either to the frequency-domain XPHM model. A detailed investigation of the recovery of $\{a_1, a_2, \theta_1, \theta_2\}$ reveals, e.g., that XPHM's performance for SXS:0716 and 0926 is due to its underestimating a_1 by 2σ to 3σ . TPHM behaves the same way for SXS:0812 and 0926. SEOB's a_1 recovery similarly makes it underestimate χ_p by 2σ for SXS:0812. For these three simulations, SEOB and TPHM recover θ_1 robustly, as does XPHM except for SXS:0716. TPHM even seems to recover θ_2 very well with the injected values recovered within $\pm\sigma$. So

the underlying issue, as was the case with the relevant $Q = 1$ subset, is the unreliable inference of a_1 for certain cases.

The recovery scores of the models for χ_p , given in Table II, quantitatively affirm the above statements. First, of the four parameters presented in the Table for $Q = 2$, the χ_p recovery scores are significantly lower than the rest for all models though SEOB's average recovery score of 0.676 is high enough to be considered reliable. Similarly, the recovery widths w for the χ_p posteriors are roughly twice those of χ_{eff} implying much broader posteriors. This is expected as spin components perpendicular to $\mathbf{L}$ are subdominant to parallel components in the waveform phase by half a post-Newtonian order [85,124].

3. Mass ratio 4:1 systems

A quick glance to the bottom 10 rows of Fig. 2 reveals that the posteriors from different models now overlap much less than they did for the $Q = 1, 2$ cases, indicating increased model disagreement for higher-mass ratio configurations. For M and q , SEOB yields the best recovery performance with $r_{\text{av}} = 0.793, 0.831$, respectively. The mass ratio recovery is especially impressive with SEOB capturing the injection within 1σ for eight out of the 10 cases. TPHM's performance is nearly comparable, reflected by its corresponding recovery scores being less than 10% lower than SEOB's. Additionally, TPHM recovers better scores than SEOB for two cases, SXS:2000 and SXS:2075, where SEOB's recovery scores are the lowest. The recovery widths of SEOB and TPHM are also quite close for this mass ratio as can be seen from Table III. XPHM's recovery performance of $\{M, q\}$ degrades significantly at this mass ratio with $M(q)$ under(over)estimated by more than 2σ in six cases, which is also reflected by the average recovery scores of 0.456 (0.512) given in the last column of Table II.

Next, focusing on the recovery of the effective spins, we see that SEOB's χ_{eff} recovery is very robust with $r_{\text{av}} = 0.803$. XPHM also recovers χ_{eff} well with $r_{\text{av}} = 0.724$ whereas TPHM's recovery is rather poor. We think this may be due to our specific choice of $Q \leq 4$ simulations, all of which have $|\chi_{\text{eff}}| < 0.02$. This poor χ_{eff} recovery performance of TPHM was already observed in Ref. [47] for the $\chi_{\text{eff}} = 0.001$ injection SXS:0050. For non-negligible values of χ_{eff} , TPHM should recover χ_{eff} well as can be seen, e.g., in Fig. 4 of Ref. [34]. This is partly supported by the results of our $Q = 8$ injections (Sec. IV B): TPHM infers χ_{eff} without biases for two out of three $|\chi_{\text{eff}}| > 0.17$ cases.

As for the recovery of the effective precession spin, we observe that SEOB significantly outperforms TPHM and XPHM with a $Q = 4$ average recovery score of 0.702 vs (< 0.4) for the other two models. However, SEOB does not recover χ_p well for SXS:1916, 1923, 2074. In fact, none of the models can recover χ_p within 2σ for the SXS:1923, 2074 injections, though SEOB and XPHM do recover χ_{eff} . These two cases have very similar intrinsic and extrinsic parameters with the exception of the sky position and polarization angles. We therefore

recommend that future waveform models analyse the SXS:1923 and 2074 simulations to assess their improved accuracy. What is more puzzling is the case of SXS:1916 which has intrinsic parameters very similar to SXS:2079 for which SEOB recovers χ_p well (but not TPHM or XPHM). This turns out to be due to the $2\text{-}\sigma$ biased recovery of a_1 for this case, which does not occur for SXS:2079. Given the similarity of the intrinsic parameters, we suspect once again that particular combinations of the extrinsic parameters cause SEOB to underestimate χ_p for SXS:1916 though the χ_{eff} recovery remains robust. After all, given that the parameter space is 15-dimensional, many combinations of parameters can result in waveforms that match the data very well.

A potential explanation for SEOB's deteriorated χ_p recovery for SXS:1923, 2070, 2074 may once again be the multipole symmetry (6). Although we do not systematically explore the consequences of this symmetry on the recovery performance of SEOB (or the other models), studies have been conducted using NRSUR7DQ4 and SEOB [58,60]. Especially relevant here is the injection-recovery of SXS:2070 by Ref. [58] using the same version of SEOB as in here and an improved version that incorporates mode asymmetry for the $\ell = m \leq 4$ multipoles. In particular, they find that the χ_p posteriors shift by more than 3σ from overestimating to underestimating the injected value when the multipole symmetry is removed. Though the "asymmetric" SEOB model shows no improvement for this particular case, it does so for two other injections and a reanalysis of GW200129 [58].

To supplement their findings, we recovered the same injection with the model IMRPhenomXO4a which has the $(2, 2) \rightarrow (2, -2)$ multipole asymmetry. Interestingly, we find that, as shown in Fig. 12, the resulting χ_p posteriors form a narrow distribution strongly peaking slightly farther away from the injected value than the mode-symmetric SEOB. Perhaps, this is due to the fact that the LALSimulation version of IMRPhenomXO4a that we use is missing the phase offset fix implemented in Ref. [125].

For the same set of $Q = 4$ simulations, XPHM severely underestimates χ_p , by more than three standard deviations in some cases, with the exception of SXS:1916 which it recovers very well. This performance is reflected in its average recovery score of 0.310. Though TPHM's performance seems marginally better based on its recovery score of 0.365, Fig. 2 reveals that TPHM's posteriors underestimate χ_p much less severely than XPHM's do. The average recovery scores are comparable because of XPHM's excellent recovery of χ_p ($\tau \approx 1$) for SXS:1916, which skews its average from roughly 0.2 to above 0.3.

4. Summary of findings for mass ratios $< 4:1$ simulations

In summary, for all 30 simulations considered here, we can confidently state that SEOB recovers the set of parameters

$\{M, q, \chi_{\text{eff}}, \chi_p\}$ well as indicated by the overall average recovery scores given in columns 2 to 5 of Table II. Its $\tau_{\text{av}} = 0.692$ for χ_p translates to a ± 1 standard deviation recovery of the effective precession spin from dozens of synthetic GW events with a wide range of extrinsic parameters. As such, in our opinion, it is reliable to conduct parameter estimation for highly precessing binary black holes up to mass ratios of 4:1 and SNRs of 40 with SEOB. Up to mass ratios of 2:1, TPHM and XPHM also provide robust recovery though not necessarily as reliable as SEOB on average. Nonetheless, as can be seen for several marginalized posterior distributions in Fig. 2, these two models outperform SEOB on occasion. This reminds us of the necessity of multimodel inference. However, the general trends of our results show that SEOB is in general more reliable and thus should be given more weight in multimodel PE when appropriate. We explore this in detail in Sec. VI.

Let us conclude this section with the interesting observation that the $\{M, q, \chi_{\text{eff}}, \chi_p\}$ posteriors from the $Q = 2$ subset are wider than their corresponding $Q = 1, 4$ -subset posteriors. This holds for each model as can be gathered from Table III. This may be because the likelihood gets more sharply peaked with decreasing q at a fixed SNR. Indeed, explicitly comparing the width of the $q = 1/2$ posteriors with those of the $q = 1/4, 1$ posteriors, we find the former to be 30% to twice larger on average than the latter two. This can also be gleaned from Refs. [33,34] though it is less clear-cut as the total SNRs are different, as would be the precession SNRs. Since χ_{eff} and χ_p are both functions of q , this behavior is also present in their posteriors as seen in Table III and Fig. 2. This is even manifest in the M posteriors due to the fact that what is directly inferred is the chirp mass M_c , which gives us the total mass via another mass-ratio-dependent relation $M = [q/(1+q)]^{-3/5} M_c$. Therefore, the resulting M posteriors for the $Q = 2$ subset are wider than their $Q = 1, 4$ counterparts as corroborated once again by Table III and Fig. 2.

B. Cases with mass ratio 8:1

Due to the abundance of the publicly available SXS simulations with $Q \leq 4$, performance studies of waveform models usually stop at this mass ratio. However, we wish to

TABLE IV. The most relevant intrinsic parameters for the five BAM binary black hole simulations used throughout this work. The columns represent the same quantities as in Table I. Note that these are all single-spin simulations, i.e., $a_2 = 0$.

CF _—	$M(M_\odot)$	Q	a_1	θ_1
67	93.2	8.0	0.40	59.9°
69	85.1	8.0	0.40	119.8°
73	89.3	8.0	0.60	89.8°
77	97.5	8.0	0.80	59.7°
78	89.7	8.0	0.80	89.6°




FIG. 3. Same as Fig. 2 but for the five single-spin $Q = 8$ BAM binary black hole simulations of Sec. IV B. The red vertical lines indicate the true values.

assess the performance of the models for BBHs with more asymmetric masses. To this end, we select five simulations from the BAM catalog. These are CF_67, CF_69, CF_73, CF_77, and CF_78. Their relevant intrinsic parameters are provided in Table IV which makes it clear that these are all single-spin simulations. Note that we chose this particular subset of simulations so as to slowly increase χ_p from 0.346 for CF_67 to $\chi_p = 0.801$ for CF_78 while still keeping the total SNR at 40 with the precession SNR ρ_p at 10.

We begin with Fig. 3 which displays the three-model recovery posteriors for the set of parameters $\{M, q, \chi_{\text{eff}}, \chi_p\}$ resulting from the injection of the $Q = 8$ BAM simulations listed in Table IV. The color coding is the same as in Fig. 2. For the recovery of the binary total mass, SEOB and TPMH exhibit alternating performances in the sense that if one model recovers M within $\pm 1\sigma$, the other does so within $\pm(2-3)\sigma$. Regardless, the combined posteriors between the two models do recover M reliably. These statements also hold true for the SEOB-TPMH recovery of the mass ratio q as can be seen in the second column of Fig. 3.

Despite similarities in their recovery of $\{M, q\}$, SEOB outperforms TPMH in inferring χ_{eff} . This is mostly because of the aforementioned drop in TPMH's in performance for when spin magnitude(s) are somewhat large and tilt angle(s) very close to $\pi/2$. Indeed, the two simulations, CF_73 and CF_78, for which TPMH overestimates χ_{eff} by more than 3σ have $a_1 \geq 0.6$ and $\theta_1 \approx \pi/2$ resulting in $|\chi_{\text{eff}}| \lesssim 10^{-2}$ while $\chi_p \gtrsim 0.6$. We had already highlighted this issue with TPMH for the $Q = 4$ cases in Sec. IV A 3 and it seems to have become even more severe for the $Q = 8$ cases. SEOB and TPMH again behave similarly in their recovery of χ_p with the two models unable to recover χ_p reliably either individually or combined together.⁶ This necessitates

further model calibration in this regime of the intrinsic parameter space.

XPHM's recovery of $\{M, q\}$ is robust for the two cases with $a_1 = 0.4$. It seems to also recover χ_{eff} within $\pm 2\sigma$ for these two cases, but once the spin magnitude exceeds 0.4, the model's reliability drops significantly. However, it recovers χ_p well for the two cases that are most challenging for SEOB and TPMH: CF_73 and CF_78 with $a_1 \geq 0.6$ and $\theta_1 \approx \pi/2$. We should caution that this is not a conclusive finding since it is based on two cases.

Interestingly, the combined SEOB-TPMH posteriors reliably recover $\{M, q, \chi_{\text{eff}}\}$ for all five cases and recover χ_p only for three of the five simulations. Therefore, at least for strongly single-spin precessing BBHs with high mass asymmetry, we advise use of at least SEOB and TPMH in the parameter estimation. More injection-recovery studies should be conducted at this mass ratio for establishing efficient PE strategies for future events.

In summary, waveform model recovery performance deteriorates further as we go from mass ratio 4:1 to mass ratio 8:1 systems. SEOB still provides marginally better recoveries than TPMH while XPHM mostly produces posteriors that are biased by more than two standard deviations. We hypothesize that the reason for this trend is due to higher order multipoles in the GW emission which become more prominent as the mass ratio of the binary increases [120]. Since XPHM does not include calibration to numerical relativity simulations for the higher multipoles in precessing systems, it is possible that errors in these high-multipole contributions could cause biases in our parameter estimation analyses. As the amplitude of each higher multipole becomes more significant, this will become more of an issue. However, we note that SEOB and TPMH are also not calibrated to numerical relativity simulations for the higher multipoles in precessing systems yet both models perform better than XPHM. Nevertheless, what is clearly evident is that no single model can be said to

⁶Here, by “reliably” we mean that the posteriors should recover χ_p within $\pm 2\sigma$ for all five injections.

be reliable in this regime. Now that many more precessing $Q = 8$ simulations have become available [18], a study similar to our $Q = 4$ subset should be conducted on these models.

V. RESULTS II: INSPIRAL-MERGER/RINGDOWN CONSISTENCY TESTS OF THE MODELS

We begin with a qualitative description of the inspiral-merger/ringdown (IMR) consistency test. We highlight how estimates for the properties of the remnant black hole can be made from independent measurements from the inspiral and ringdown portions of the signal and therefore checked for consistency.

The early inspiral portion of the GW signal is well described as a post-Newtonian series with the chirp mass $\mathcal{M}_c$ contributing to the GW phasing at the leading (Newtonian) order and the symmetric mass ratio η entering at $1/2$ PN order higher. The spins first enter at the relative 1.5PN order via their components parallel to the angular momentum: $\chi_i = a_i \cos \theta_i$, for $i = 1, 2$ [126,127]. Based solely on the amplitude, phase and frequency evolution of the inspiral portion of the GW signal, estimates for the binary parameters can be obtained [see, e.g., [111]]. These parameters can be mapped to estimates for the properties of the remnant black hole by evaluating numerical-relativity fits [128–130], as has been done previously [76,77,131].

On the other hand, the ringdown portion of the GW signal is well described by black hole perturbation theory [132–140]. Here, a perturbed Kerr black hole radiates GWs via a superposition of exponentially damped quasinormal modes (QNMs), where the QNM frequencies and damping times are functions of only the mass and spin of the unperturbed final Kerr black hole (see Refs. [139,141,142] for reviews). Based solely on a measurement of the frequency and the decay time of the fundamental QNM, estimates for the binary parameters can be obtained [143,144].

We therefore have two independent methods for obtaining the final mass and spin of the remnant black hole. In the case of zero noise and assuming GR to be the correct theory of gravity, these measurements should agree modulo modeling systematics. To quantify this agreement, we introduce the dimensionless quantities [77]

$$\Delta M_f / \bar{M}_f, \quad \Delta a_f / \bar{a}_f, \quad (19)$$

where an overbar denotes the average between the inspiral (I) and the merger-ringdown (MR) values, viz.,

$$\bar{M}_f := (M_f^I + M_f^{\text{MR}})/2, \quad (20)$$

$$\bar{a}_f := (a_f^I + a_f^{\text{MR}})/2 \quad (21)$$

and

$$\Delta M_f := M_f^I - M_f^{\text{MR}}, \quad (22)$$

$$\Delta a_f := a_f^I - a_f^{\text{MR}}. \quad (23)$$

If there are no systematic biases due to mismodeling then we expect the posteriors for $\Delta M_f / \bar{M}_f$ and $\Delta a_f / \bar{a}_f$ to be centered around zero [145]. In this case, the inspiral values from post-Newtonian theory are said to be consistent with those from the ringdown based on black hole perturbation theory.

Given that waveform models employ fitting formulas calibrated to NR simulations to obtain M_f, a_f as functions of the initial masses and spins, the IMR consistency check is a good way to expose systematic biases. Such fits were already used to estimate M_f, a_f for GW150914 [146]. Spin precession only complicates the matter as, e.g., the waveform models extend the postinspiral Euler angles (used in the frame rotation to go to the final a_f -frame) in different ways [33,35,36,41,65,72,147]. An especially alarming scenario is systematic biases mimicking false beyond-GR signatures in parameter estimation [67,148].⁷ Therefore, it is of crucial importance to systematically survey the performance of different waveform models under the IMR consistency test.

This test has become a standard check for the LVK Collaboration [78–80,131]. The details of the implementation within a Bayesian framework can be found in Ref. [76] (also see Refs. [77,145,150,151]), so we provide a brief introduction here. First, we must select a cutoff frequency, $\bar{f}$, to separate the inspiral regime from the MR part of the signal. One option is to set

$$\bar{f} = f_{\text{ISCO}}^{\text{Sch}} = (\pi)^{-1} 6^{-3/2} (GM/c^3)^{-1} \approx (60M_\odot/M) 73 \text{ Hz}, \quad (24)$$

i.e., the GW frequency corresponding to the innermost stable circular orbit (ISCO) in Schwarzschild spacetime. We have checked that this is less than the frequency of the minimum energy circular orbit (MECO) [42] for all our injections using the `lal.SimIMRPhenomXfMECO` function. Another choice is to set

$$\bar{f} = f_{\text{ISCO}}^{\text{Kerr}} = \frac{c^3}{\pi G} \frac{\sqrt{M}}{r_{\text{ISCO}}^{3/2} + a\sqrt{M}}, \quad (25)$$

where r_{ISCO} is obtained from the solution to [152]

$$r(r - 6M) + 8a\sqrt{Mr} - 3a^2 = 0. \quad (26)$$

This second choice corresponds to the ISCO frequency of an equatorial, prograde timelike orbit around a Kerr black hole of mass M and spin $a = \chi M$ with $\chi \in [0, 1)$ [76].

Based on GW signals with only the dominant quadrupolar contribution Ref. [76] argues that the specified cutoff

⁷Alternatively, an eccentric BBH signal analyzed with quasi-circular templates can also yield a false beyond-GR signature [149].

frequency should not have a significant effect on the IMR consistency test, as long as a reasonable selection is made. We test their assertion here for waveform models that include precession effects and higher multipoles by reconducting the same IMR consistency tests with both of these frequency cutoffs. For a nonspinning BBH with $q = 1$, $a_1 = a_2 = 0$, the resulting black hole's final spin is approximately $2M/3$ [153]. Thus, for $M = 60M_\odot$, $a = 2M/3$, Eq. (25) gives us approximately 150 Hz, roughly twice the value from Eq. (24). We list both cutoff frequencies in Table VII in the Appendix, where one sees that the Kerr frequency is, in general, 2 to 2.5 times larger than the Schwarzschild one.

We note that a $60M_\odot$, $a = 2M/3$ Kerr black hole has its $\ell = 2$ fundamental ($n = 0$) QNM frequencies ranging from ~ 150 Hz to ~ 250 Hz [141], the former being the QNM frequency of the $(2, -2)$ multipole, and the latter $(2, 2)$. So although the Kerr ISCO may be an appropriate cutoff for IMR consistency tests involving only the $(2, 2)$ mode, and indeed is commonly used by the LVK [80], this cutoff may be too high when other higher order multipoles are considered. Be that as it may, we explore how the results of the IMR consistency test change depending on the choice of cutoff.

For the inspiral PE analyses, the likelihood is integrated from 23 Hz to $\tilde{f}$ (the lower bound is consistent with the full IMR analyses, see Sec. III for details). For the merger-ringdown analyses, we integrate the likelihood from $\tilde{f}$ to 2048 Hz. The final mass and spin from the inspiral part are

obtained from the posteriors for the binary masses and spins via waveform-specific fits to NR simulations [128–130]. M_f , a_f are also derived from the inferred fundamental QNM of the MR injection analyses. As in Sec. IV, our Bayesian inference analyses are performed using the BILBY and BILBY_PIPE packages with the DYNESTY_{quantum} sampler (using consistent settings as before). Given that we are using consistent priors as in Sec. IV, this leads to nonuniform effective priors for $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$. Ref. [79] reweighted the inspiral and MR analyses to uniform priors for the deviation parameters, arguing that this more clearly conveys the information gained from the data. Consistent with Ref. [78], we do not perform any reweighting in this work. We finally use the PESUMMARY package [154] to calculate posteriors for $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$, via the SUMMARYTGR executable.

A. Cases with $Q \leq 4$

For expediency, we perform the IMR consistency test of all 30 NR injections with xPHM. We present results using other waveform models for selected cases owing to computational restrictions. The 1D marginalized posteriors for $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$ from the xPHM recovery of the $30Q \leq 4$ injections using both cutoff frequencies are displayed in the top and bottom rows of Fig. 4. Each row is further divided into three sections containing 10 plots separated by the mass ratio with the $Q = 1$ ($Q = 4$) subset on the left (right).




FIG. 4. The posteriors for $\Delta M_f/\bar{M}_f$ (top row) and $\Delta a_f/\bar{a}_f$ (bottom row) from the IMR consistency test performed with xPHM for the $30Q \leq 4$ numerical relativity simulations whose parameters are given in Tables I and VI. See Eqs. (20)–(23) for the relevant definitions. Each row is further grouped into three clusters of 10 subfigures each separated by the mass ratio of the binary black hole simulations with $Q = 1, 2, 4$ clusters from left to right. Note that the vertical scale is different in each cluster. In each plot, we show the IMR results obtained with two different choices for the cutoff frequency, see Eqs. (24) and (25). In orange (cyan) we show the results obtained with a cutoff frequency equal to the Schwarzschild (Kerr) ISCO frequency.

To quantify the performance of XPHM in the IMR consistency test, we additionally introduce a new metric for the posterior recovery in the context of the IMR consistency test: $\mathcal{D}^{2D}$. This is the probability that gives the fraction of the 2D posterior probability distribution enclosed by the isoprobability contour that passes through 0 (see, e.g., Ref. [155] for details). Under this metric, smaller probabilities indicate better consistency with GR. For example, a 2D normal distribution that captures 0 within two standard deviations along a single dimension (with the cross section along second dimension remaining centered around 0) yields $\mathcal{D}^{2D} \leq 88\%$; this increases to $\mathcal{D}^{2D} \leq 99\%$ if the 2D distribution captures 0 within two standard deviations in both dimensions. We list the values we obtain for $\mathcal{D}^{2D}$ for all 30 cases using both cutoff frequencies in Table VII.

Note that GR deviations are seen more frequently when studying the marginalized 1D posterior distributions than the 2D posteriors. Consider, e.g., SXS:0965 whose 1D posteriors in Fig. 4 for both $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$ point to GR violations using the Schwarzschild ISCO cutoff. However, the corresponding value for $\mathcal{D}^{2D}$ in Table VII is 85% indicating, somewhat marginally, that the final mass and spin are consistent with GR. Since a true-positive GR violation is a rather unexpected result, we opt for the most conservative metric: we declare the IMR consistency test (IMRCT) a failure only when $\mathcal{D}^{2D} > 88\%$ and when both 1D posteriors “miss” the GR value, 0, by more than two standard deviations. This is still a frequency dependent claim as we explain below.

The first observation from Fig. 4 is that, in general, the posteriors obtained using two different cutoffs are not consistent. In other words, the results of the IMRCT *depend* on the choice of the cutoff frequency, at least for signals composed of multiple harmonics emitted by strongly precessing BBHs. The 1D posteriors obtained using the cutoff based on the Kerr black hole are, in general, broader than the posteriors obtained using the Schwarzschild cutoff. This is reasonable since there is less information/SNR in the postinspiral phase of the signal when employing the larger Kerr cutoff, and therefore wider posterior distributions are expected. As a result, the Kerr ISCO results tend to be more consistent with GR as confirmed by the values of $\mathcal{D}^{2D}$ in Table VII. However, for the $Q = 2$ subset of BBHs, we report that in general the Schwarzschild ISCO results are more consistent with GR.

The discussion above raises the question of whether a single clear case of synthetic GR violation can be found. For example, using the Kerr ISCO cutoff and the 1D cyan posteriors of Fig. 4, we could declare a GR violation for SXS:1921. However, we see that the results using the Schwarzschild cutoff contradict this. Returning to Table VII, we observe that there is not a single occurrence of both $\mathcal{D}_{\text{Sch}}^{2D}$ and $\mathcal{D}_{\text{Kerr}}^{2D}$ exceeding 88% for the same case. If we relax this bound somewhat, we can place SXS:0838, 2000 in

a list of cases of interest. Looking at the 1D posteriors in Fig. 4, we can expand this list to also include SXS:0926, 0965, 1197, 1217, 1922, 2070 giving us eight special cases to explore in greater detail. Given that in some cases we observed large differences in the obtained posterior distributions for different waveform models in Sec. IV, two natural questions are (a) how much the differences between the Schwarzschild and Kerr cutoff frequencies depend on the choice of a particular waveform model, and (b) whether or not model systematics can impact GR violation statements. To this end, we reran the IMRCT with TPHM and SEOB for the above-list of eight cases.

With regards to the first question, we obtained consistent conclusions for TPHM and SEOB as with XPHM, i.e., the Schwarzschild ISCO posteriors are narrower than the Kerr results, and the Kerr results yield posteriors that are more consistent with GR as shown in Fig. 9 in the Appendix. We next focus on the second question. Rather than presenting results from two different cutoff frequencies, which show broadly the same features among the waveform models considered, we solely focus on the Schwarzschild ISCO results since they show the worst case scenario. In Fig. 5, we show a comparison of the performance of these three models in the IMRCT. We see significant differences between the models, especially in the $\Delta a_f/\bar{a}_f$ posteriors. It is worth noting that when this analysis is repeated with the Kerr cutoff, we observe more consistent results owing to the wider posterior distributions. It is clear from Fig. 5 that SEOB and TPHM perform better in the IMRCT than




FIG. 5. Model performance under the IMR consistency test for the eight special cases mentioned in Sec. VA using the Schwarzschild ISCO as the frequency cutoff [Eq. (24)].

XPHM. Moreover, TPHM recovers both $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$ within $\pm 2\sigma$ every time. $\Delta a_f/\bar{a}_f$ consistency is especially robust as evident from the last column of the figure. SEOB can also be said to mostly pass this test, but it seems to produce less consistent final spins for more mass symmetric, i.e., $Q \leq 2$ binaries than TPHM, while its posteriors of $\Delta a_f/\bar{a}_f$ for the three $Q = 4$, shown in Fig. 5, are very consistent. Perhaps this is indicative of the model's performance toward the single-spin limit which we explore more in Sec. VB.

We additionally checked the performance of the recently developed frequency-domain model IMRPhenomXO4a (xO4A) [72] for these eight IMRCTs. We present the results in Fig. 10 in the Appendix where we compare this model's performance with that of XPHM. xO4A seems to bring some improvement with the most important differences being in the posteriors for $\Delta a_f/\bar{a}_f$ for the three $Q = 4$ cases: SXS:1922, 2000, 2070. In particular, xO4A's posteriors are much narrower and peak much closer to $\Delta a_f = 0$ than for XPHM. However, the former are still biased by $\gtrsim 2\sigma$.

As for the drop in XPHM's performance for these eight cases, and especially the three $Q = 4$ cases, this may be due to the mismodeling of QNM frequencies in the coprecessing frame which is exacerbated for BBHs with $Q \gtrsim 4$, $a_1 \gtrsim 0.7$ and $\theta_1 \gtrsim \pi/2$, and is worst for $\theta_1 \approx 150^\circ$ in the single-spin case [156]. This is consistent with the findings of Ref. [157] where it is shown that for single-spin systems, $\theta_1 \approx 150^\circ$ corresponds roughly to $\cos\langle\beta\rangle \approx 0$ where $\beta = \cos^{-1}(\hat{\mathbf{J}} \cdot \hat{\mathbf{L}})$ and the angular brackets represent its average in the frequency domain between $f(t = t_c)$ and $f(t = t_c + \Delta t)$ with t_c denoting the coalescence (merger) time and $\Delta t \in [40, 90]M$. Essentially, when $\cos\beta = 0$, the shift away from the true QNM ringdown frequency is maximized [156]. Further support for this hypothesis is the fact that both SEOB and xO4A have improved merger-ringdown attachments based on the work of Ref. [157]. This might explain why they produce $\Delta a_f/\bar{a}_f$ posteriors more consistent with GR than TPHM for the aforementioned three $Q = 4$ cases. However, TPHM has a similar ringdown prescription to XPHM, but performs better for the $Q = 4$ BBHs under the IMRCT. So the missing QNM corrections offer a partial explanation for the XPHM results. We will return to this discussion in the next section, where we consider systems even closer to the single-spin limit ($Q = 8$).

Focusing on XPHM with the Schwarzschild ISCO cutoff, we performed a few more checks motivated by the results in Ref. [158], where they observe GR violations even when injecting and recovering with the same waveform model. They argue that the presence of higher-order multipoles may be the cause as there have not been extensive tests on the IMRCTs for waveform models including higher-order multipoles. To investigate whether this may be the reason why we observe inconsistent recoveries for M_f and a_f , we

reperform the IMRCT, but now inject XPHM waveforms that have the same parameters as the previously injected SXS waveforms. By doing so, we remove the possibility of model systematics biasing results. In this fashion, we reanalyze these eight aforementioned SXS simulations. As can be seen in Fig. 11 in the Appendix, we find that in general XPHM performs much better in the IMRCT for simulations that are produced using XPHM injections compared with NR. This implies that the M_f and a_f inconsistencies that we observe in Fig. 4 are indeed due to model systematics, and hence a consequence of XPHM not faithfully representing NR for strongly precessing GW signals.

We recommend that these eight simulations be used as the IMR consistency test set for future precessing waveform models, keeping in mind that the results are sensitive to the choice of the cutoff frequency.

B. Cases with mass ratio 8:1

As in the last section, we begin with results obtained using XPHM. These are shown in Fig. 6 for the five $Q = 8$ injections of Table IV. We see from the figure that the posteriors are quite distinguishable from each other, i.e., the IMRCT results continue to be sensitive to the cutoff choice as q decreases.

In Fig. 7, we show the results of the IMR consistency test on the three waveform models using the Schwarzschild ISCO cutoff. As was the case with the $Q \leq 4$ injections, TPHM produces the most consistent posteriors for both $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$ in four out of the five cases. SEOB




FIG. 6. Results of the IMR consistency test of Sec. V with XPHM for our $Q = 8$ binary black hole injections whose parameters are given in Table IV. In each plot, we show the IMR results obtained with two different choices for the cutoff frequency, as in Fig. 4.

performs similarly for spin consistency, but yields $\Delta M_f/\bar{M}_f$ posteriors for CF_67 that are inconsistent at 3σ . XPHM produces very broad, mostly uninformative, posteriors for $\Delta a_f/\bar{a}_f$, but is, in principle, consistent with $\Delta a_f = 0$ to 2σ for all five simulations. On the other hand, it consistently points to $\Delta M_f < 0$ at 1σ to 3σ .

Previous studies have shown that $Q \gtrsim 4$, $a_1 \gtrsim 0.7$, $\theta_1 \approx 150^\circ$ is the most challenging for XPHM in the IMRCT [156,157]. In our sample set, the simulation CF_69 with $Q = 8$, $a_1 = 0.4$ and $\theta_1 \approx 120^\circ$ is has parameters closest to the region. Consistent with previous works, we observe the worst [2D] recovery using both cutoffs for this simulation as shown in Fig. 6. On the other hand, the XPHM posteriors recovered for CF_77 are mostly more consistent with GR despite this simulation having $a_1 = 0.8$, possibly because $\theta_1 \approx 60^\circ$.

Another interesting finding is the inability of the models to pass the ΔM_f consistency check for CF_73 when using the Schwarzschild ISCO cutoff. SEOB and TPHM also fail the test for Δa_f and XPHM produces a posterior for $\Delta a_f/\bar{a}_f$ that ranges from -1 to 1 . This simulation has $a_1 = 0.6$ and $\theta_1 \approx 90^\circ$, so it has rather strong precession, but not as much as CF_78 with $a_1 = 0.8$, $\theta_1 \approx 90^\circ$, which produces much more consistent results. We performed several additional checks/runs to understand this behavior: (i) We visually compared the 90%-credible time-domain waveforms generated by the three models for CF_73 to those for CF_78 and found no smoking-gun signatures explaining the issue. (ii) We reinjected the BBH with intrinsic parameters of CF_73 (CF_78) and extrinsic parameters of CF_78

(CF_73) and recovered posteriors consistent with $\Delta M_f = \Delta a_f = 0$ for all three models similar to what is shown for CF_78 in Fig. 7 (again using the Schwarzschild ISCO cutoff). (iii) We computed the precession SNR, ρ_p , for these two new injections, as well as the original injections of CF_73 and CF_78. The troublesome case yields the lowest value with $\rho_p = 9.6$ with the original CF_78 injection giving $\rho_p = 9.9$ and the two new injections mentioned in (ii) yielding 11.0 and 10.7, respectively. (iv) We finally looked at how the injected SNR is distributed across the LIGO L1-H1 network. For CF_73, we have $\text{SNR}_{\text{H1}} = 11$, $\text{SNR}_{\text{L1}} = 38$; for CF_78, $\text{SNR}_{\text{H1}} = 31$, $\text{SNR}_{\text{L1}} = 25$. However, for the reinjections of (ii), we have 31, 25 and 10, 38, respectively, consistent with the original cases. In short, we have not been able to determine why the models pass the IMRCT for CF_78, but not for CF_73. In summary, for the strongly double-spin precessing BBHs considered here, the results of the IMR consistency test depend on the choice of the cutoff frequency. The posteriors obtained using the cutoff based on the Kerr ISCO are generally more consistent with GR, albeit broader. The results obtained using the cutoff based on the Schwarzschild ISCO can be taken to be more informative though can sometimes point to false deviations of GR.

Focusing on the Schwarzschild cutoff, we see that TPHM passes the IMR consistency test convincingly for mass ratios up to $Q = 4$. SEOB mostly produces consistent posteriors as well. XPHM shows consistency for 22 out of the 30 BBHs. For the remaining eight cases, this model yields biased posteriors that could be mistaken as violations of GR. The more recent phenomenological model IMRPhenomXO4a offers a discernible improvement, but still points to false GR violations based on the $\Delta a_f/\bar{a}_f$ posteriors for four cases (see Fig. 10). For the single-spin precessing $Q = 8$ BBHs, no single waveform model can be relied on for IMR consistency though the combined SEOB-TPHM posteriors mostly yield consistent results. Additional injections using the new $Q = 8$ double-spin precessing simulations of Ref. [18] should be conducted to further test the combined SEOB-TPHM performance using at least two different cutoffs. What we can state with confidence is that strong claims of GR violations based on the IMR consistency test should be based on more than one reasonable frequency cutoff and should employ at least two waveform models, preferably more. Finally, waveform systematics should always be checked using state of the art NR simulations.




FIG. 7. Model performance under the IMR consistency test for our $Q = 8$ binary black hole injections whose parameters are given in Table IV and discussed in Sec. V. For these cases we consider a cutoff frequency based on the Schwarzschild ISCO.

VI. RESULTS III: INCORPORATING MODEL UNCERTAINTY INTO THE INJECTION STUDY

As introduced in Sec. III, the model-dependent posterior distribution is calculated through Bayes' theorem using Eq. (7). If we have an ensemble of N models and we wish to produce a single distribution that marginalizes over model uncertainty, Bayesian model averaging can be used,




FIG. 8. One-dimensional marginalized posterior distributions obtained for the inferred total mass (first column), mass ratio (second column), effective parallel (third column) and effective perpendicular spins (fourth column) for the 30 sxs binary black hole simulations used throughout this work. The red vertical lines indicate the true values. The {blue, orange, green} posteriors are obtained via the NR-informed, evidence-informed and standard methods for combining PE results from multiple waveform models (see Sec. VI).

$$p(\lambda|d) = \sum_{i=1}^N p(\lambda|d, \mathfrak{M}_i) p(\mathfrak{M}_i|d), \quad (27)$$

where $p(\mathfrak{M}_i|d)$, is the probability of the model $\mathfrak{M}_i$ given the data, and $p(\lambda|d, \mathfrak{M}_i)$ is the model-dependent posterior distribution introduced in Eq. (7). By reapplying Bayes' theorem to Eq. (27), it can be shown that Bayesian model averaging simply averages the model-dependent posterior distributions, weighted by their respective evidence [81], hereafter referred to as the *evidence informed approach*

$$\begin{aligned} p(\lambda|d) &= \sum_{i=1}^N \left[\frac{\mathcal{Z}_i \Pi(\mathfrak{M}_i)}{\sum_{j=1}^N \mathcal{Z}_j \Pi(\mathfrak{M}_j)} \right] p(\lambda|d, \mathfrak{M}_i), \\ &= \sum_{i=1}^N w_i p(\lambda|d, \mathfrak{M}_i), \end{aligned} \quad (28)$$

where $\Pi(\mathfrak{M}_i)$ is the discrete prior probability for the choice of model, which is often assumed to constant and equal to $1/N$, and w_i is the weight given the each model dependent posterior distribution. Alternatively, we can remain completely agnostic and simply mix the model dependent posterior distributions with equal weight (see e.g. Ref. [1]), i.e., $p(\mathfrak{M}_i|d) = 1/N$, hereafter referred to as the *standard approach*.

An alternative technique for marginalizing over model uncertainty is to simultaneously infer the model and model properties in a single *joint* Bayesian analysis [159]. This technique expands the binary properties $\lambda = \{\lambda_1, \lambda_2, \dots, \lambda_N\}$ to include an additional parameter, m , which is mapped to a model $\mathfrak{M}_i$ during the sampling. As with other parameters, a prior must be defined for choice of model.

Defining a prior probability for the choice of model is challenging as the accuracy of each model varies across the parameter space. For this reason Ref. [82] proposed using a parameter-space dependent prior informed by the model's accuracy to numerical relativity simulations. This implies that the most accurate GW model will more likely be used to evaluate the likelihood in each region of parameter space. Although other functional forms are possible, Ref. [82]

suggested the following prior,

$$\Pi(\mathfrak{M}_i|\lambda) \propto \mathcal{M}_i(\lambda)^{-4} \quad (29)$$

where $\mathcal{M}_i$ is the mismatch for model $\mathfrak{M}_i$ as defined in Eq. (11). Unfortunately, calculating the mismatch exactly is computationally expensive. Consequently, an interpolant can be built prior to the analysis which is fast to evaluate [82]. In this work, we employ this technique and use the same interpolant as introduced in Ref. [82]. We hereafter refer to this algorithm as the *NR-informed method*.

The mismatch interpolant constructed in Ref. [82] has only been trained on binaries with mass ratios $1/4 \leq q \leq 1$. Therefore, in this section we focus only on simulations whose intrinsic and extrinsic parameters are listed in Tables I and VI.

We present the main results in Fig. 8 where we recover the 30 injections of Table I using the three methods summarized above. Note that the color coding has now changed: blue corresponds to the NR-informed method [82], green to the evidence-informed method [81] and orange to the standard method [1] for combining posteriors from multiple waveform models. For most of the panels in the figure, we see that the three approaches mostly yield overlapping posteriors. In what follows, we will focus on cases (or parameters) where a discernible disagreement is observed. The figure is further supplemented by the recovery scores given in Table V presented in the same format as Table II.

Starting with the $Q = 1$ subset, we see that the NR-informed method produces slightly better posteriors for the total binary mass M . The q posteriors strongly overlap with each other with the most prominent disagreement occurring for sxs:0982 which proved to be a very challenging injection for all three waveform models as exhibited in Fig. 2. The NR-informed method shows more discernible improvement compared to the other two methods for inferring χ_{eff} . On the other hand, the χ_p posteriors are mostly indistinguishable except for two cases where the NR-informed method produces slightly better posteriors. These findings are quantified in Table V in terms of the recovery score, $\mathbf{r}$, once again averaged over the 10 ($Q = 1$) simulation

TABLE V. The recovery score (14) of each method for combining posteriors from different waveform models for the parameters listed in column one. Columns two to four list the averages over the entire set of 30 simulations for NR-informed, evidence-informed and the standard methods respectively labeled as N, E, S here. The remaining columns show the average of the same quantity over the size-10 $Q = 1, 2, 4$ subsets. A value of $\mathbf{r} > 0.7$ indicates a recovery of the injected parameter within one standard deviation of the mean of the posterior distribution. Note that the lower recovery score for the mass ratio q for the $Q = 1$ subset is due to the fact that the injected value lies at the boundary of the domain hence causing the posteriors to rail as can be seen in Fig. 8. The closer the recovery score to unity, the better.

Θ	$\mathbf{r}_{\text{av}}^{\text{N}}$	$\mathbf{r}_{\text{av}}^{\text{E}}$	$\mathbf{r}_{\text{av}}^{\text{S}}$	$\mathbf{r}_{\text{av}}^{\text{N}}[Q=1]$	$\mathbf{r}_{\text{av}}^{\text{E}}[Q=1]$	$\mathbf{r}_{\text{av}}^{\text{S}}[Q=1]$	$\mathbf{r}_{\text{av}}^{\text{N}}[Q=2]$	$\mathbf{r}_{\text{av}}^{\text{E}}[Q=2]$	$\mathbf{r}_{\text{av}}^{\text{S}}[Q=2]$	$\mathbf{r}_{\text{av}}^{\text{N}}[Q=4]$	$\mathbf{r}_{\text{av}}^{\text{E}}[Q=4]$	$\mathbf{r}_{\text{av}}^{\text{S}}[Q=4]$
M	0.827	0.750	0.845	0.922	0.813	0.808	0.810	0.828	0.834	0.749	0.610	0.892
q	0.748	0.691	0.769	0.553	0.559	0.556	0.905	0.851	0.867	0.787	0.662	0.885
χ_{eff}	0.804	0.763	0.793	0.846	0.672	0.692	0.763	0.832	0.889	0.802	0.785	0.797
χ_p	0.698	0.675	0.633	0.713	0.674	0.667	0.675	0.669	0.650	0.707	0.681	0.583

posteriors for each parameter of interest. We also note that the ($Q = 1$)-average recovery scores produced by the NR-informed method are also slightly better (closer to unity) than the best recovery score produced by an individual waveform model as given in Table II with the exception of q posteriors, which are virtually the same between the three different methods.

For the $Q = 2$ subset, the standard method seems to yield a slightly better average recovery score for the total binary mass M than the other two methods. This is partly due to the aforementioned property of the cost function (16) penalizing asymmetric posteriors more severely, which is most evident for SXS:1200. Visually, the NR-informed (blue) posterior seems like a better recovery than the standard (green) one, but as the former is much less symmetric than the latter, it returns a higher cost function with respect the corresponding standard deviation. The evidence-informed method also produces a marginally higher average recovery score for M than the NR-informed one. However, the latter method does yield the highest values of $\mathbf{r}$ for the mass ratio q as can be corroborated by Table V and Fig. 8. In fact, $\mathbf{r}_{\text{av}}$ is once again higher than the highest average recovery produced by an individual waveform model (see Table II). Overall, all three methods yield recovery scores above 0.7 for all $Q = 2$ simulations except for one: SXS:0814. This is not surprising as none of the waveform models can recover the injected value $q = 1/2$ within $\pm\sigma$ for this case (see Fig. 2).

As for the recovery of the effective spins, the standard method yields the highest recovery scores for χ_{eff} owing to the fact that it produces the widest and the most symmetric posteriors for this parameter. Given that for $Q = 2$, it was the model XPHM that yielded the highest recovery scores, we expect that neither the NR-informed, nor the evidence-informed methods produce the best recoveries since these approaches tend to mostly prefer SEOB and/or TPHM over XPHM when generating their respective posteriors. For the recovery of χ_p , the three methods produce comparable recovery scores with the NR-informed method yielding marginally the highest average recovery score of $\mathbf{r}_{\text{av}} = 0.675$ close to the best average produced by SEOB ($\mathbf{r}_{\text{av}} = 0.702$). This outcome is consistent assuming that both the NR-informed and evidence-informed methods heavily favor SEOB as both TPHM and XPHM yield $\mathbf{r}_{\text{av}} \approx 0.3$. The standard method, on the other hand, produces wider posteriors, thus also capturing the injected values within 1σ for most cases.

For the $Q = 4$ subset, the standard method yields the highest recovery scores owing to the fact that it yields the widest posteriors for M and q , some of which are bimodal due to the fact that the underlying waveform models produced strongly disagreeing posteriors as can be seen from the bottom left section of Fig. 2. This is most evident for SXS:1921 and 2070. The NR-informed method, on the other hand, produces narrower posteriors, which recover

the injections well for eight of the 10 cases except for SXS:1916 and 2000. For the latter, the method's posteriors are most similar to SEOB's as it employs SEOB 98.4% of the time in the parameter estimation analysis due to the fact that SEOB yields the lowest mismatches to NR in the relevant parameter space. And though $\{M, q\}$ recovery is less than ideal, M_c is recovered within 1σ and η within 2σ . In this sense, the recovery is still robust since M_c and η are the parameters most relevant to the waveform phase.

For the case of SXS:1916, we see that the NR-informed analysis infers $\{M, q\}$ posteriors that are similar to SEOB's, but not identical. This is despite the fact that the method uses SEOB 94.7% of the time (TPHM 4.7% and XPHM 1%). As a first check, we reran this PE analysis with more aggressive settings and obtained statistically indistinguishable posteriors indicating that the run results have converged. As can be gathered from Figs. 2 and 8, the NR-informed analysis performs better than the SEOB-only analysis when estimating the spins, but worse when looking at the $M - q$ ($m_1 - m_2$) projection of the 15D parameter space. However, when we translate the $M - q$ posteriors to $M_c - \eta$ posteriors, we observe that the chirp mass is recovered with a score of $\mathbf{r} > 0.9$ and the difference between the injected value of η and the median of the posterior is $\approx 5\%$. In short, though one may naively expect the NR-informed method to agree with SEOB for this case, given (i) the high dimensionality of the parameter space, (ii) the stochastic nature of the parameter estimation process, and (iii) the fact that the NR-informed method does not use SEOB 100% of the time, the most likely parameters need not agree with those from an SEOB PE run.

Another interesting result of note is that the evidence-informed method yields posteriors for SXS:1916 that are very similar to those from the XPHM run, consistent with the fact that XPHM has overwhelmingly the largest Bayesian evidence for this simulation: when computing the weights in Eq. (28), we find that evidence-informed gives a mixing fraction of 99.99% for XPHM. In a similar manner, this method gives a mixing fraction of 99.99% for SEOB for SXS:2000, hence the similarity in the posteriors between SEOB and the evidence-informed method for this case. We find it intriguing that this method so heavily favors two different models for two seemingly very similar sets of intrinsic parameters (however, the extrinsic parameters differ significantly).

When it comes to inferring spin information, the standard method outputs the most uninformative posteriors exhibiting bimodalities and even trimodalities in the recovery of χ_{eff} . Therefore, its slightly better recovery score compared to the evidence-informed method should be taken in this context. The method's recovery of χ_p further degrades due to the fact that the individual model posteriors overlap weakly resulting in a 90% CI width of ≈ 0.5 for the standard method for injected values of

$\chi_p \simeq 0.8$. On the other hand, both the NR and evidence-informed methods recover the injections well with the former yielding $r_{av} \simeq 0.8, 0.7$, respectively for χ_{eff}, χ_p , and the latter method giving slightly lower values. The average recovery scores coming from the NR-informed method further equal or marginally exceed the best $Q = 4$ recovery scores from Table II.

In short, our findings indicate that for mass ratios of 1:1 and 4:1, the NR-informed method provides the most consistent parameter recovery with respect to the injections. For mass ratios of 2:1, this method still gives the most reliable recovery for q and χ_p with its recovery of the total mass M only slightly inferior to the other two methods. This is because the method mostly employs the waveform model SEOBNRV5PHM as it is the most accurate (non-surrogate) model for most of the parameter space of precessing BBHs [47]. However, it does not recover χ_{eff} as well as IMRPhenomXPHM for the $Q = 2$ set of simulations injected here. As a result, the NR-informed method does not provide the best recovery for χ_{eff} . This is one of the main takeaways of this section: the most accurate waveform model does not necessarily yield the most accurate posteriors for a single parameter, but it is likely to provide a more accurate 15-dimensional posterior distribution.

It is also possible that the accuracy of the NR-informed method is being compromised by the limitations of the interpolant used to predict model faithfulness to NR in the parameter space. As detailed in Ref. [82], the interpolant is rather coarse: it is constructed from 250 points in the effective spin space and on average it is 5% accurate to the logarithm of the true mismatch. Therefore, it remains a possibility that the method may sometimes favor the less accurate waveform model in certain parts of the parameter space, although we did not find this to be the case for the BBHs analyzed in Ref. [82].

VII. SUMMARY AND OUTLOOK

In this article, we tested the robustness of three state of the art waveform models, SEOBNRV5PHM, IMRPhenomTPHM, and IMRPhenomXPHM, by performing Bayesian inference on 35 strongly precessing numerical relativity simulations. 30 of these simulations featured BBHs with significant spin magnitudes: $\{a_1, a_2\} \in [0.8, 0.85]$, nearly planar spin tilts: $|\theta_i - \pi/2| \leq 10^\circ, i = 1, 2$, and mass ratios of 1:1, 2:1, and 4:1. The remaining five simulations were systems with single spin precession and mass ratio of 8:1. We arranged the extrinsic parameters in each case such that the total SNR in an advanced LIGO O4 network consisting of the Hanford and Livingston observatories equaled 40 with a precession SNR of 10.

We first focused on the recovery performance of the individual waveform models as detailed in Sec. IV. For the sake of brevity, we presented one-dimensional (marginalized) model posterior distributions for only four intrinsic

parameters: the total mass M , the small mass ratio q , the effective inspiral spin χ_{eff} and the effective precession spin χ_p . The full model dependent posterior distribution can be found in our open-access data repository.⁸ We then looked at model performance under the inspiral-merger-ringdown consistency test in Sec. V, specifically documenting this in terms of the posteriors for $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$ given by Eqs. (20)–(23). We additionally investigated how the results of this test depend on the choice of the cutoff frequency. Finally, we compared three different methods for combining parameter posteriors from multiple waveform models in Sec. VI. We presented the results from this comparison once again in terms of the marginalized 1D posteriors of $\{M, q, \chi_{eff}, \chi_p\}$.

Though we noted many interesting findings from our study throughout this article, we summarize our most robust conclusions and recommendations below. However, we note that as with any injection-recovery analysis our conclusions are dependent on the simulations analyzed, i.e., our specific choice of strongly precessing systems, on the values of extrinsic parameters selected and on the chosen waveform models.

- (i) SEOB provides the least biased parameter recovery for systems with mass ratios less than or equal to 4:1. Though we reported a few occurrences of biases, its overall mean recovery scores for $\{M, q, \chi_{eff}, \chi_p\}$ testify that, on average, it captures each injection within one standard deviation of the posterior median.
- (ii) TPHM and XPHM perform reasonably well for the systems with mass ratios of 1:1 and 2:1. As indicated by their recovery scores given in Table II, we see that the true source properties are mostly captured within one standard deviation of the median. However, these waveform models produce a biased inference for χ_p for mass ratios of 4:1. We understand that this is because these models infer biased estimates for the spin magnitudes of each black hole, although they recover the spin tilt angles well. This is also the case for SEOB for the few systems where it exhibits biases in χ_p .
- (iii) One likely cause of biases in the χ_p inference is the multipole symmetry (6) that the models employ. This is only an approximation in the case of coprecessing multipoles and needs to be abandoned to improve the models.
- (iv) The waveform models yield much more biased posteriors for BBHs with mass ratios of 8:1. However, the combined SEOB-TPHM posteriors capture $\{M, q, \chi_{eff}\}$ well. Nonetheless, in light of the new catalog of numerical relativity simulations, we recommend that this should be further tested with other double-spin precessing simulations.

⁸https://github.com/akcays2/Injection_Campaign.

TABLE VI. Extrinsic parameters for the 30 SXS BBH systems for which we provide the relevant intrinsic parameters in Table I. See Sec. II for the definition of each symbol. All angles are given in radians.

SXS:BBH	ρ_p	d_L (Mpc)	α	δ	ψ	$\vartheta_{\text{LN},0}$	$\vartheta_{\text{JN},0}$	ϕ_{JL}
0764	10.00	301.25	0.35	0.52	0.63	1.57	1.31	4.78
0767	9.98	507.56	2.44	0.52	0.94	1.05	1.07	2.82
0838	10.02	440.72	0.70	-0.52	0.31	1.57	1.32	4.77
0841	10.00	314.61	0.35	1.22	0.31	1.05	1.08	2.80
0925	10.01	347.30	2.09	0.52	0.63	1.57	1.80	2.10
0935	9.99	177.02	1.40	0.17	0.63	1.05	0.95	4.94
0965	10.00	73.10	2.44	-0.87	0.94	1.57	1.79	2.35
0982	10.04	379.71	1.05	-0.87	0.94	1.57	1.39	4.50
1205	9.99	212.96	1.40	-0.17	0.94	1.57	1.65	2.74
1217	10.00	378.72	1.05	-0.52	0.31	1.57	1.33	5.56
0716	10.01	374.75	0.00	0.17	0.00	1.05	1.23	2.62
0717	10.00	403.21	2.09	0.87	0.00	1.05	1.33	1.07
0812	10.00	617.22	1.05	-1.22	0.94	0.52	0.27	5.21
0814	10.00	855.66	2.79	0.52	0.00	0.20	0.58	1.34
0926	9.99	440.74	0.00	-1.22	0.00	1.57	1.94	1.54
0936	9.99	374.01	1.40	-0.17	0.63	0.52	0.32	5.74
0966	10.00	133.31	2.79	-0.87	0.63	1.05	0.74	3.90
0976	10.00	583.42	0.70	1.57	0.31	0.52	0.78	1.18
1197	10.00	311.98	1.05	0.17	0.63	0.52	0.85	1.77
1200	10.00	759.11	0.35	-0.17	0.31	0.20	0.55	1.94
1916	10.00	524.54	2.44	-1.57	0.00	0.20	0.80	1.76
1921	11.46	693.67	0.00	-0.87	0.63	0.52	1.06	1.17
1922	9.95	168.36	2.09	-0.87	0.63	1.57	1.08	5.46
1923	10.12	744.55	0.35	-0.52	0.94	0.20	0.53	1.85
2000	9.98	283.26	1.05	1.57	1.26	1.57	0.97	5.21
2004	10.00	263.35	1.05	0.87	0.31	1.05	0.56	5.67
2070	10.02	290.78	1.75	0.87	0.63	1.57	1.69	2.97
2074	10.09	754.28	2.79	0.87	0.00	0.20	0.46	1.69
2075	10.01	468.36	3.14	0.87	0.31	1.05	1.56	0.89
2079	10.27	555.25	2.09	0.87	0.31	0.20	0.75	1.83

- (v) The results of the IMR consistency test applied to waveforms with multiple harmonics emitted by strongly precessing systems, depend on the choice of the cutoff frequency.
- (vi) The use of the Schwarzschild ISCO frequency as the cutoff [Eq. (24), Table VII] yields narrower posteriors for $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$, with some cases displaying false GR violations using the waveform model XPHM (SXS:0838, 0965, 1922, 2000, 2070 according to the criteria introduced in Sec. VA). On the other hand, using the Kerr ISCO frequency for the cutoff [Eq. (25), Table VII] leads to at most two GR violations (SXS:1921, 0716), albeit with broader posteriors for $\Delta M_f/\bar{M}_f$ and $\Delta a_f/\bar{a}_f$.
- (vii) Regardless of the cutoff, for systems with mass ratios ranging up to 4:1, TPHM is the most reliable model for passing the IMR consistency test with SEOB coming a close second.
- (viii) Model performance under the IMR consistency test degrades as we go from mass ratio 4:1 systems to

TABLE VII. Results from the IMR consistency test for two different choices of cutoff frequency, one based on the Schwarzschild (Sch), $f_{\text{ISCO}}^{\text{Sch}}$ [see Eq. (24)], and the other based on the Kerr ISCO frequency, $f_{\text{ISCO}}^{\text{Kerr}}$ [see Eq. (25)]. The probabilities $\mathcal{D}_{\text{Sch}}^{2\text{D}}$ and $\mathcal{D}_{\text{Kerr}}^{2\text{D}}$ denote the fraction of the 2D $p(\Delta M_f/\bar{M}_f, \Delta a_f/\bar{a}_f)$ posterior probability distribution enclosed by the isoprobability contour that passes through the GR value when using the Schwarzschild ISCO and Kerr ISCO as the cutoff frequencies respectively. In this table, smaller probabilities indicate better consistency with GR.

SXS:BBH	$f_{\text{ISCO}}^{\text{Sch}}$ (Hz)	$f_{\text{ISCO}}^{\text{Kerr}}$ (Hz)	$\mathcal{D}_{\text{Sch}}^{2\text{D}}$ (%)	$\mathcal{D}_{\text{Kerr}}^{2\text{D}}$ (%)
0764	85	200	94	70
0767	86	190	11	21
0838	85	199	92	76
0841	86	191	13	20
0925	85	205	31	7
0935	86	190	6	9
0965	85	205	85	30
0982	86	196	56	11
1205	86	194	1	12
1217	85	217	64	38
0716	82	187	53	89
0717	82	177	27	59
0812	82	179	8	44
0814	82	186	13	40
0926	82	183	66	39
0936	82	187	15	21
0966	83	194	43	23
0976	82	180	25	63
1197	84	192	75	53
1200	82	191	21	52
1916	75	166	73	3
1921	75	156	39	87
1922	75	151	90	49
1923	74	162	30	2
2000	74	150	84	76
2004	76	157	23	60
2070	75	149	87	24
2074	74	161	65	0
2075	75	155	18	22
2079	75	139	65	6

8:1 systems. This is especially evident for the results based on the Schwarzschild ISCO frequency cutoff.

- (ix) The recent phenomenological waveform model IMRPhenomXO4a performs better than XPHM in the IMR consistency test as shown in Fig. 10, but not as well as SEOB or TPHM. We also found instances where IMRPhenomXO4a produced more biased 1D posteriors compared to XPHM as exhibited in Fig. 12 in the Appendix.
- (x) Based on the results of our IMR consistency tests, we recommend that at least two waveform models be employed to mitigate against potential issues caused by waveform systematics. We also advocate the use of at least two frequency cutoffs.

- (xi) The NR-informed method of combining posteriors from multiple waveform models (by incorporating model accuracy [82]) produces intrinsic parameter posteriors most consistent with the injections for highly precessing systems with mass ratios from 1:1 to 4:1.
- (xi) The standard method of combining multimodel posteriors via equal weights leads to wide, and often multimodal distributions when individual waveform models disagree. This is most evident when inferring the precession spin χ_p for mass ratio 4:1 systems.

We also highlighted some specific simulations that may be useful for testing the next generation of waveform models: we suggest performing Bayesian inference on the SXS:1923 and 2074 simulations. For IMR consistency tests, we recommend analysing the eight simulations first highlighted in Fig. 5. For testing future methods of combining multiple model-dependent posterior distributions, or techniques that marginalize over waveform uncertainty [see, e.g., [160–165], we recommend analysing the injections SXS:0764, 0926, 1200, 1916, 1923, 2004. Given that the SXS catalog of numerical relativity simulations has recently been updated and expanded [18], the recommended list of simulations will inevitably change in the future.

Another obvious extension to our work here is to systematically document the biases caused by the coprecessing multipole symmetry (6). Though we looked at this using the model IMRPhenomXO4a for a few cases, a more thorough way is to employ the asymmetrized version of SEOBNRv5PHM à la Ref. [58] for all of our injections.

Our demonstration of the dependence of the IMR consistency test on the frequency cutoff brings up the question of whether a single universally meaningful cutoff exists. Though the ISCO frequency of a spinning (Kerr) black hole seems like a more appropriate choice than the Schwarzschild value, the former can sometimes be larger than the QNM frequencies of the final black hole. An alternative choice for the cutoff frequency could be obtained using the inverse adiabaticity parameter, $Q_\Omega := \Omega^2/\dot{\Omega}$ where Ω is the orbital frequency.⁹ When $Q_\Omega \gg 1$, the binary is in the inspiral regime. When $Q_\Omega \approx \mathcal{O}(1)$, the transition to plunge takes place. Our suggestion is to use $f(t_C)$ as the new cutoff where t_C is the root of $Q_\Omega(t) = \mathcal{C}$. Somewhat arbitrarily picking $\mathcal{C} = 10$ yields GW

frequencies between 80 Hz and 100 Hz for the 30 SXS simulations of Table I. These values are slightly higher than those given by the Schwarzschild ISCO frequency in Table VII so they may offer a reasonable alternative for the cutoff frequency. Of course, a detailed study into a suitable value of $\mathcal{C}$ is required.

Although we only focused on three state of the art waveform models that characterize gravitational waveforms from quasicircular binary black holes in this study, gravitational-wave astronomy has plethora of models available, which include different underlying physics, approaches and assumptions. This includes, but not limited to, NRSUR7DQ4 [59], TEOBResumS [35,46,87,167–175], TEOBResumS-Dali [176–181], SEOBNRv5EHM [182], IMRPhenomXODE [183], and IMRPHENOMTEHM [184]. We therefore encourage that a similar injection-recovery campaign is repeated to assess the performance of some (or all of) these additional (and future) models, and understand the level of bias that can be expected for astrophysically interesting binary systems.

ACKNOWLEDGMENTS

We thank Lorenzo Pompili for comments during the LIGO-Virgo-KAGRA internal review and the anonymous referee for the peer review. We especially thank Lorenzo Pompili for the suggestion of reperforming the IMR consistency test for simulations produced with the same model as for the recovery, and Nathan Johnson-McDaniel for the suggestion of increasing the cutoff frequency to the Kerr ISCO radius. We also thank Hector Estelles and Eleanor Hamilton for guidance in calculating the remnant properties with the IMRPhenomTPHM and IMRPhenomXO4a waveforms respectively. We are also grateful to Eleanor Hamilton for their expertise with rotating and plotting numerical relativity waveforms. We finally thank Laura Nuttall and Jonathan E. Thompson for discussions throughout this project. S. A. and J. M. U. acknowledge support from the University College Dublin Ad Astra Fellowship, and C. H. thanks the UKRI Future Leaders Fellowship for support through the Grant No. MR/T01881X/1. This work used the Sciama High Performance Compute (HPC) cluster, which is supported by the ICG, SEPNet and the University of Portsmouth.

DATA AVAILABILITY

The data that support the findings of this article are openly available [83].

⁹This quantity is related to the more ubiquitous Q_ω , first introduced in Ref. [166], via $Q_\Omega \approx Q_\omega/2 = \omega^2/(2\dot{\omega})$, where $\omega = 2\pi f$.

APPENDIX: ADDITIONAL TABLES AND FIGURES




FIG. 9. SEOB's performance under the IMR consistency test for the eight special cases mentioned in Sec. VA. Orange histograms are obtained using the Schwarzschild ISCO as the frequency cutoff [Eq. (24)] and the cyan ones using the Kerr ISCO frequency given in Eq. (25).




FIG. 10. Performance of the XPHM and IMRPhenomXO4a (XO4A) models under the IMR consistency test for eight cases in which XPHM does not recover the true value within two standard deviations. See Sec. VA for details. For these cases we consider a cutoff frequency based on the Schwarzschild ISCO.




FIG. 11. Same as Fig. 10, but we show results of the IMR consistency test for NR injections (orange) and XPHM injections with the same parameters as the NR (gray). For both cases, we analyse the injections with XPHM and consider a cutoff frequency based on the Schwarzschild ISCO.




FIG. 12. Same as Fig. 2, but we specifically compare the performance of the waveform models xPHM and xO4A for a selection of four SXS binary black hole simulations with mass ratios $Q = 2$ or $Q = 4$ from Sec. IV A. The red vertical lines indicate the true values.

- [1] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, *Phys. Rev. X* **13**, 041039 (2023).
- [2] A. H. Nitz, S. Kumar, Y.-F. Wang, S. Kasta, S. Wu, M. Schäfer, R. Dhurkunde, and C. D. Capano, 4-OGC: Catalog of gravitational waves from compact binary mergers, *Astrophys. J.* **946**, 59 (2023).
- [3] S. Olsen, T. Venumadhav, J. Mushkin, J. Roulet, B. Zackay, and M. Zaldarriaga, New binary black hole mergers in the LIGO-Virgo O3a data, *Phys. Rev. D* **106**, 043009 (2022).
- [4] A. K. Mehta, S. Olsen, D. Wadekar, J. Roulet, T. Venumadhav, J. Mushkin, B. Zackay, and M. Zaldarriaga, New binary black hole mergers in the LIGO-Virgo O3b data, *Phys. Rev. D* **111**, 024049 (2025).
- [5] D. Wadekar, J. Roulet, T. Venumadhav, A. K. Mehta, B. Zackay, J. Mushkin, S. Olsen, and M. Zaldarriaga, New black hole mergers in the LIGO-Virgo O3 data from a gravitational wave search including higher-order harmonics, [arXiv:2312.06631](https://arxiv.org/abs/2312.06631).
- [6] GraceDB—Gravitational-wave candidate event database, <https://gracedb.ligo.org/superevents/public/O3/>.
- [7] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
- [8] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin *et al.*, Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2014).
- [9] T. Akutsu *et al.* (KAGRA Collaboration), Overview of KAGRA: Detector design and construction history, *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2021).
- [10] LIGO Scientific and Virgo Collaborations, LIGO-india, proposal of the consortium for indian initiative in gravitational-wave observations (IndIGO), <https://dcc.ligo.org/LIGO-M1100296/public>.
- [11] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs, *Phys. Rev. X* **9**, 031040 (2019).
- [12] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, *Phys. Rev. X* **11**, 021053 (2021).
- [13] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, *Phys. Rev. D* **109**, 022001 (2024).
- [14] J. Veitch *et al.*, Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library, *Phys. Rev. D* **91**, 042003 (2015).
- [15] A. H. Mroue, M. A. Scheel, B. Szilagy, H. P. Pfeiffer, M. Boyle *et al.*, A catalog of 174 binary black-hole simulations for gravitational-wave astronomy, *Phys. Rev. Lett.* **111**, 241104 (2013).
- [16] M. Boyle *et al.*, The SXS Collaboration catalog of binary black hole simulations, *Classical Quantum Gravity* **36**, 195006 (2019).
- [17] SXS Gravitational Waveform Database, <https://data.black-holes.org/waveforms/index.html>.
- [18] M. A. Scheel *et al.*, The SXS Collaboration's third catalog of binary black hole simulations, [arXiv:2505.13378](https://arxiv.org/abs/2505.13378).
- [19] E. Hamilton *et al.*, Catalog of precessing black-hole-binary numerical-relativity simulations, *Phys. Rev. D* **109**, 044032 (2024).
- [20] E. Hamilton, E. Fauchon-Jones, M. Hannam, C. Hoy, C. Kalaghatgi, L. London, J. Thompson, D. Yeeles, S. Ghosh, S. Khan, P. Kolitsidou, and A. Vano-Vinuales, Precessing

- binary-black-hole numerical relativity catalogue (minimal data release), [10.5281/zenodo.7673796](https://doi.org/10.5281/zenodo.7673796) (2023).
- [21] E. Hamilton, E. Fauchon-Jones, M. Hannam, C. Hoy, C. Kalaghatgi, L. London, J. Thompson, D. Yeeles, S. Ghosh, S. Khan, P. Kolitsidou, and A. Vano-Vinuales, Precessing binary-black-hole numerical relativity catalogue (complete data release), [10.5281/zenodo.7677297](https://doi.org/10.5281/zenodo.7677297) (2023).
- [22] J. Healy, C. O. Lousto, Y. Zlochower, and M. Campanelli, The RIT binary black hole simulations catalog, *Classical Quantum Gravity* **34**, 224001 (2017).
- [23] J. Healy, C. O. Lousto, J. Lange, R. O’Shaughnessy, Y. Zlochower, and M. Campanelli, Second RIT binary black hole simulations catalog and its application to gravitational waves parameter estimation, *Phys. Rev. D* **100**, 024021 (2019).
- [24] J. Healy and C. O. Lousto, Third RIT binary black hole simulations catalog, *Phys. Rev. D* **102**, 104018 (2020).
- [25] J. Healy and C. O. Lousto, Fourth RIT binary black hole simulations catalog: Extension to eccentric orbits, *Phys. Rev. D* **105**, 124010 (2022).
- [26] P. Ajith, M. Boyle, D. A. Brown, B. Brügmann, L. T. Buchman *et al.*, The NINJA-2 catalog of hybrid post-Newtonian/numerical-relativity waveforms for nonprecessing black-hole binaries, *Classical Quantum Gravity* **29**, 124001 (2012).
- [27] I. Hinder, A. Buonanno, M. Boyle, Z. B. Etienne, J. Healy *et al.*, Error-analysis and comparison to analytical models of numerical waveforms produced by the NRAR Collaboration, *Classical Quantum Gravity* **31**, 025012 (2014).
- [28] K. Jani, J. Healy, J. A. Clark, L. London, P. Laguna, and D. Shoemaker, Georgia Tech Catalog of Gravitational Waveforms, *Classical Quantum Gravity* **33**, 204001 (2016).
- [29] D. Ferguson *et al.*, Second MAYA catalog of binary black hole numerical relativity waveforms, [arXiv:2309.00262](https://arxiv.org/abs/2309.00262).
- [30] A. Rashti, R. Gamba, K. Chandra, D. Radice, B. Daszuta, W. Cook, and S. Bernuzzi, Binary black hole waveforms from high-resolution GR-Athena++ simulations, *Phys. Rev. D* **111**, 104078 (2025).
- [31] E. A. Huerta *et al.*, Physics of eccentric binary black hole mergers: A numerical relativity perspective, *Phys. Rev. D* **100**, 064003 (2019).
- [32] J. Healy, C. O. Lousto, J. Lange, and R. O’Shaughnessy, Application of the third RIT binary black hole simulations catalog to parameter estimation of gravitational waves signals from the LIGO-Virgo O1/O2 observational runs, *Phys. Rev. D* **102**, 124053 (2020).
- [33] G. Pratten *et al.*, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, *Phys. Rev. D* **103**, 104056 (2021).
- [34] H. Estellés, M. Colleoni, C. García-Quirós, S. Husa, D. Keitel, M. Mateu-Lucena, M. L. Planas, and A. Ramos-Buades, New twists in compact binary waveform modeling: A fast time-domain model for precession, *Phys. Rev. D* **105**, 084040 (2022).
- [35] R. Gamba, S. Akçay, S. Bernuzzi, and J. Williams, Effective-one-body waveforms for precessing coalescing compact binaries with post-Newtonian twist, *Phys. Rev. D* **106**, 024020 (2022).
- [36] A. Ramos-Buades, A. Buonanno, H. Estellés, M. Khalil, D. P. Mihaylov, S. Ossokine, L. Pompili, and M. Shiferaw, Next generation of accurate and efficient multipolar precessing-spin effective-one-body waveforms for binary black holes, *Phys. Rev. D* **108**, 124037 (2023).
- [37] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Properties of the binary black hole merger GW150914, *Phys. Rev. Lett.* **116**, 241102 (2016).
- [38] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Effects of waveform model systematics on the interpretation of GW150914, *Classical Quantum Gravity* **34**, 104002 (2017).
- [39] R. Costa, A. Buonanno, A. Bohé, A. Taracchini, I. Hinder, and S. Ossokine, Enriching the symphony of gravitational waves from binary black holes by tuning higher harmonics, *Phys. Rev. D* **98**, 084028 (2018).
- [40] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, Surrogate model of hybridized numerical relativity binary black hole waveforms, *Phys. Rev. D* **99**, 064045 (2019).
- [41] H. Estellés, S. Husa, M. Colleoni, D. Keitel, M. Mateu-Lucena, C. García-Quirós, A. Ramos-Buades, and A. Borchers, Time-domain phenomenological model of gravitational-wave subdominant harmonics for quasicircular nonprecessing binary black hole coalescences, *Phys. Rev. D* **105**, 084039 (2022).
- [42] G. Pratten, S. Husa, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estelles, and R. Jaume, 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**, 064001 (2020).
- [43] C. García-Quirós, M. Colleoni, S. Husa, H. Estellés, G. Pratten, A. Ramos-Buades, M. Mateu-Lucena, and R. Jaume, Multimode frequency-domain model for the gravitational wave signal from nonprecessing black-hole binaries, *Phys. Rev. D* **102**, 064002 (2020).
- [44] G. Riemenschneider, P. Retegno, M. Breschi, A. Albertini, R. Gamba, S. Bernuzzi, and A. Nagar, Assessment of consistent next-to-quasicircular corrections and postadiabatic approximation in effective-one-body multipolar waveforms for binary black hole coalescences, *Phys. Rev. D* **104**, 104045 (2021).
- [45] L. Pompili *et al.*, Laying the foundation of the effective-one-body waveform models SEOBNRv5: Improved accuracy and efficiency for spinning non-precessing binary black holes, *Phys. Rev. D* **108**, 124035 (2023).
- [46] A. Nagar, P. Retegno, R. Gamba, S. Albanesi, A. Albertini, and S. Bernuzzi, Analytic systematics in next generation of effective-one-body gravitational waveform models for future observations, *Phys. Rev. D* **108**, 124018 (2023).
- [47] J. Mac Uilliam, S. Akçay, and J. E. Thompson, Survey of four precessing waveform models for binary black hole systems, *Phys. Rev. D* **109**, 084077 (2024).
- [48] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), The population of merging compact binaries inferred using gravitational waves through GWTC-3, *Phys. Rev. X* **13**, 011048 (2023).
- [49] C. Hoy, S. Fairhurst, and I. Mandel, Precession and higher order multipoles in binary black holes (and lack thereof), *Phys. Rev. D* **111**, 023037 (2025).

- [50] V. Kalogera, Spin orbit misalignment in close binaries with two compact objects, *Astrophys. J.* **541**, 319 (2000).
- [51] I. Mandel and R. O’Shaughnessy, Compact binary coalescences in the band of ground-based gravitational-wave detectors, *Classical Quantum Gravity* **27**, 114007 (2010).
- [52] D. Gerosa, E. Berti, R. O’Shaughnessy, K. Belczynski, M. Kesden, D. Wysocki, and W. Gladysz, Spin orientations of merging black holes formed from the evolution of stellar binaries, *Phys. Rev. D* **98**, 084036 (2018).
- [53] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, Illuminating black hole binary formation channels with spins in Advanced LIGO, *Astrophys. J. Lett.* **832**, L2 (2016).
- [54] K. Belczynski, A. Buonanno, M. Cantiello, C. L. Fryer, D. E. Holz, I. Mandel, M. C. Miller, and M. Walczak, The formation and gravitational-wave detection of massive stellar black-hole binaries, *Astrophys. J.* **789**, 120 (2014).
- [55] I. Mandel and A. Farmer, Merging stellar-mass binary black holes, *Phys. Rep.* **955**, 1 (2022).
- [56] M. Hannam, C. Hoy, J. E. Thompson, S. Fairhurst, V. Raymond *et al.*, General-relativistic precession in a black-hole binary, *Nature (London)* **610**, 652 (2022).
- [57] R. Macas, A. Lundgren, and G. Ashton, Revisiting the evidence for precession in GW200129 with machine learning noise mitigation, *Phys. Rev. D* **109**, 062006 (2024).
- [58] H. Estellés, A. Buonanno, R. Enficiaud, C. Foo, and L. Pompili, Adding equatorial-asymmetric effects for spin-precessing binaries into the SEOBNRv5PHM waveform model, [arXiv:2506.19911](https://arxiv.org/abs/2506.19911).
- [59] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, Surrogate models for precessing binary black hole simulations with unequal masses, *Phys. Rev. Res.* **1**, 033015 (2019).
- [60] P. Kolitsidou, J. E. Thompson, and M. Hannam, Impact of antisymmetric contributions to signal multipoles in the measurement of black-hole spins, *Phys. Rev. D* **111**, 024050 (2025).
- [61] E. Payne, S. Hourihane, J. Golomb, R. Udall, R. Udall, D. Davis, and K. Chatziioannou, Curious case of GW200129: Interplay between spin-precession inference and data-quality issues, *Phys. Rev. D* **106**, 104017 (2022).
- [62] N. Gupte *et al.*, Evidence for eccentricity in the population of binary black holes observed by LIGO-Virgo-KAGRA, [arXiv:2404.14286](https://arxiv.org/abs/2404.14286).
- [63] D. Fernando, R. O’Shaughnessy, and D. Williams, Efficient reanalysis of events from GWTC-3 with RIFT and asimov, [arXiv:2412.02999](https://arxiv.org/abs/2412.02999).
- [64] M. d. L. Planas, A. Ramos-Buades, C. García-Quirós, H. Estellés, S. Husa, and M. Haney, Eccentric or circular? A reanalysis of binary black hole gravitational wave events for orbital eccentricity signatures, [arXiv:2504.15833](https://arxiv.org/abs/2504.15833).
- [65] M. Colleoni, F. A. R. Vidal, C. García-Quirós, S. Akçay, and S. Bera, Fast frequency-domain gravitational waveforms for precessing binaries with a new twist, *Phys. Rev. D* **111**, 104019 (2025).
- [66] A. G. Abac *et al.* (LIGO Scientific, Virgo, KAGRA, and Virgo Collaborations), Observation of gravitational waves from the coalescence of a $2.5 - 4.5 M_{\odot}$ compact object and a neutron star, *Astrophys. J. Lett.* **970**, L34 (2024).
- [67] A. Dhani, S. Völkel, A. Buonanno, H. Estelles, J. Gair, H. P. Pfeiffer, L. Pompili, and A. Toubiana, Systematic biases in estimating the properties of black holes due to inaccurate gravitational-wave models, *Phys. Rev. X* **15**, 031036 (2025).
- [68] A. Bohé *et al.*, Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors, *Phys. Rev. D* **95**, 044028 (2017).
- [69] S. Babak, A. Taracchini, and A. Buonanno, Validating the effective-one-body model of spinning, precessing binary black holes against numerical relativity, *Phys. Rev. D* **95**, 024010 (2017).
- [70] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. Jiménez Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal, *Phys. Rev. D* **93**, 044006 (2016).
- [71] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jiménez Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, *Phys. Rev. D* **93**, 044007 (2016).
- [72] J. E. Thompson, E. Hamilton, L. London, S. Ghosh, P. Kolitsidou, C. Hoy, and M. Hannam, PhenomXO4a: A phenomenological gravitational-wave model for precessing black-hole binaries with higher multipoles and asymmetries, *Phys. Rev. D* **109**, 063012 (2024).
- [73] K. Chatziioannou, A. Klein, N. Yunes, and N. Cornish, Constructing gravitational waves from generic spin-precessing compact binary inspirals, *Phys. Rev. D* **95**, 104004 (2017).
- [74] LIGO Scientific and Virgo Collaborations, Noise curves for use in simulations pre-o4, DCC (2022) <https://dcc.ligo.org/LIGO-T2200043/public>.
- [75] S. Fairhurst, R. Green, C. Hoy, M. Hannam, and A. Muir, Two-harmonic approximation for gravitational waveforms from precessing binaries, *Phys. Rev. D* **102**, 024055 (2020).
- [76] A. Ghosh *et al.*, Testing general relativity using golden black-hole binaries, *Phys. Rev. D* **94**, 021101 (2016).
- [77] A. Ghosh, N. K. Johnson-McDaniel, A. Ghosh, C. K. Mishra, P. Ajith, W. Del Pozzo, C. P. L. Berry, A. B. Nielsen, and L. London, Testing general relativity using gravitational wave signals from the inspiral, merger and ringdown of binary black holes, *Classical Quantum Gravity* **35**, 014002 (2018).
- [78] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Tests of general relativity with the binary black hole signals from the LIGO-Virgo catalog GWTC-1, *Phys. Rev. D* **100**, 104036 (2019).
- [79] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog, *Phys. Rev. D* **103**, 122002 (2021).
- [80] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), Tests of general relativity with GWTC-3, [arXiv:2112.06861](https://arxiv.org/abs/2112.06861).

- [81] G. Ashton and S. Khan, Multiwaveform inference of gravitational waves, *Phys. Rev. D* **101**, 064037 (2020).
- [82] C. Hoy, S. Akcay, J. Mac Uilliam, and J. E. Thompson, Incorporating model accuracy into gravitational-wave Bayesian inference, *Nat. Astron.* **9**, 1256 (2025).
- [83] J. M. U. Sarp Akcay and Charlie Hoy, https://github.com/akcays2/Injection_Campaign, Injection_campaign.
- [84] P. Schmidt, I. W. Harry, and H. P. Pfeiffer, Numerical relativity injection infrastructure, [arXiv:1703.01076](https://arxiv.org/abs/1703.01076).
- [85] 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**, 6274 (1994).
- [86] S. Ossokine, M. Boyle, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, Comparing post-Newtonian and numerical-relativity precession dynamics, *Phys. Rev. D* **92**, 104028 (2015).
- [87] S. Akcay, R. Gamba, and S. Bernuzzi, A hybrid post-Newtonian—effective-one-body scheme for spin-precessing compact-binary waveforms, *Phys. Rev. D* **103**, 024014 (2021).
- [88] T. Damour, Coalescence of two spinning black holes: An effective one-body approach, *Phys. Rev. D* **64**, 124013 (2001).
- [89] E. Racine, Analysis of spin precession in binary black hole systems including quadrupole-monopole interaction, *Phys. Rev. D* **78**, 044021 (2008).
- [90] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Simple model of complete precessing black-hole-binary gravitational waveforms, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [91] M. Pürrer, M. Hannam, P. Ajith, and S. Husa, Testing the validity of the single-spin approximation in inspiral-merger-ringdown waveforms, *Phys. Rev. D* **88**, 064007 (2013).
- [92] P. Schmidt, F. Ohme, and M. Hannam, Towards models of gravitational waveforms from generic binaries II: Modeling precession effects with a single effective precession parameter, *Phys. Rev. D* **91**, 024043 (2015).
- [93] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, Tracking the precession of compact binaries from their gravitational-wave signal, *Phys. Rev. D* **84**, 024046 (2011).
- [94] P. Schmidt, M. Hannam, and S. Husa, Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and non-precessing inspiral signals, *Phys. Rev. D* **86**, 104063 (2012).
- [95] M. Boyle, R. Owen, and H. P. Pfeiffer, A geometric approach to the precession of compact binaries, *Phys. Rev. D* **84**, 124011 (2011).
- [96] M. Khalil, A. Buonanno, H. Estelles, D. P. Mihaylov, S. Ossokine, L. Pompili, and A. Ramos-Buades, Theoretical groundwork supporting the precessing-spin two-body dynamics of the effective-one-body waveform models SEOBNRv5, *Phys. Rev. D* **108**, 124036 (2023).
- [97] G. Ashton *et al.*, BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [98] O. Edy, A. Lundgren, and L. K. Nuttall, Issues of mis-modeling gravitational-wave data for parameter estimation, *Phys. Rev. D* **103**, 124061 (2021).
- [99] N. Metropolis and S. Ulam, The Monte Carlo method, *J. Am. Stat. Assoc.* **44**, 335 (1949).
- [100] J. Skilling, Nested sampling, *AIP Conf. Proc.* **735**, 395 (2004).
- [101] J. Skilling, Nested sampling for general Bayesian computation, *Bayesian Anal.* **1**, 833 (2006).
- [102] C. Pankow, P. Brady, E. Ochsner, and R. O’Shaughnessy, Novel scheme for rapid parallel parameter estimation of gravitational waves from compact binary coalescences, *Phys. Rev. D* **92**, 023002 (2015).
- [103] 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).
- [104] A. Delaunoy, A. Wehenkel, T. Hinderer, S. Nissanke, C. Weniger, A. R. Williamson, and G. Louppe, Lightning-fast gravitational wave parameter inference through neural amortization, [arXiv:2010.12931](https://arxiv.org/abs/2010.12931).
- [105] S. R. Green, C. Simpson, and J. Gair, Gravitational-wave parameter estimation with autoregressive neural network flows, *Phys. Rev. D* **102**, 104057 (2020).
- [106] A. J. K. Chua and M. Vallisneri, Learning Bayesian posteriors with neural networks for gravitational-wave inference, *Phys. Rev. Lett.* **124**, 041102 (2020).
- [107] S. R. Green and J. Gair, Complete parameter inference for GW150914 using deep learning, *Mach. Learn. Sci. Tech.* **2**, 03LT01 (2021).
- [108] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, Real-time gravitational wave science with neural posterior estimation, *Phys. Rev. Lett.* **127**, 241103 (2021).
- [109] H. Gabbard, C. Messenger, I. S. Heng, F. Tonolini, and R. Murray-Smith, Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy, *Nat. Phys.* **18**, 112 (2022).
- [110] V. Tiwari, C. Hoy, S. Fairhurst, and D. MacLeod, Fast non-Markovian sampler for estimating gravitational-wave posteriors, *Phys. Rev. D* **108**, 023001 (2023).
- [111] S. Fairhurst, C. Hoy, R. Green, C. Mills, and S. A. Usman, Simple parameter estimation using observable features of gravitational-wave signals, *Phys. Rev. D* **108**, 082006 (2023).
- [112] E. Baird, S. Fairhurst, M. Hannam, and P. Murphy, Degeneracy between mass and spin in black-hole-binary waveforms, *Phys. Rev. D* **87**, 024035 (2013).
- [113] J. E. Thompson, C. Hoy, E. Fauchon-Jones, and M. Hannam, On the use and interpretation of signal-model indistinguishability measures for gravitational-wave astronomy, [arXiv:2506.10530](https://arxiv.org/abs/2506.10530).
- [114] B. J. Owen, Search templates for gravitational waves from inspiraling binaries: Choice of template spacing, *Phys. Rev. D* **53**, 6749 (1996).
- [115] A. M. Knee, J. McIver, and M. Cabero, Prospects for measuring off-axis spins of binary black holes with plus-era gravitational-wave detectors, *Astrophys. J.* **928**, 21 (2022).
- [116] 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. R. Astron. Soc.* **499**, 3295 (2020).
- [117] C. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, PyCBC Inference:

- A Python-based parameter estimation toolkit for compact binary coalescence signals, *Publ. Astron. Soc. Pac.* **131**, 024503 (2019).
- [118] J. S. Speagle, DYNESTY: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [119] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XIII. Cosmological parameters, *Astron. Astrophys.* **594**, A13 (2016).
- [120] C. Mills and S. Fairhurst, Measuring gravitational-wave higher-order multipoles, *Phys. Rev. D* **103**, 024042 (2021).
- [121] J. A. Gonzalez, M. D. Hannam, U. Sperhake, B. Brügmann, and S. Husa, Supermassive recoil velocities for binary black-hole mergers with antialigned spins, *Phys. Rev. Lett.* **98**, 231101 (2007).
- [122] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt, Maximum gravitational recoil, *Phys. Rev. Lett.* **98**, 231102 (2007).
- [123] W. Tichy and P. Marronetti, Binary black hole mergers: Large kicks for generic spin orientations, *Phys. Rev. D* **76**, 061502 (2007).
- [124] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner, Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms, *Phys. Rev. D* **79**, 104023 (2009); *Phys. Rev. D* **84**, 049901(E) (2011).
- [125] J. Mielke, S. Ghosh, A. Borchers, and F. Ohme, Revisiting the relationship of black-hole kicks and multipole asymmetries, *Phys. Rev. D* **111**, 064009 (2025).
- [126] E. Poisson and C. M. Will, Gravitational waves from inspiraling compact binaries: Parameter estimation using second post-Newtonian wave forms, *Phys. Rev. D* **52**, 848 (1995).
- [127] L. Blanchet, Gravitational radiation from post-Newtonian sources and inspiralling compact binaries, *Living Rev. Relativity* **17**, 2 (2014).
- [128] X. Jiménez-Forteza, D. Keitel, S. Husa, M. Hannam, S. Khan, and M. Pürrer, Hierarchical data-driven approach to fitting numerical relativity data for nonprecessing binary black holes with an application to final spin and radiated energy, *Phys. Rev. D* **95**, 064024 (2017).
- [129] J. Healy and C. O. Lousto, Remnant of binary black-hole mergers: New simulations and peak luminosity studies, *Phys. Rev. D* **95**, 024037 (2017).
- [130] F. Hofmann, E. Barausse, and L. Rezzolla, The final spin from binary black holes in quasi-circular orbits, *Astrophys. J.* **825**, L19 (2016).
- [131] M. Breschi, R. O’Shaughnessy, J. Lange, and O. Birnholtz, Inspiral-merger-ringdown consistency tests with higher modes on gravitational signals from the second observing run of LIGO and Virgo, *Classical Quantum Gravity* **36**, 245019 (2019).
- [132] T. Regge and J. A. Wheeler, Stability of a Schwarzschild singularity, *Phys. Rev.* **108**, 1063 (1957).
- [133] C. V. Vishveshwara, Scattering of gravitational radiation by a Schwarzschild black-hole, *Nature (London)* **227**, 936 (1970).
- [134] F. J. Zerilli, Gravitational field of a particle falling in a Schwarzschild geometry analyzed in tensor harmonics, *Phys. Rev. D* **2**, 2141 (1970).
- [135] W. H. Press, Long wave trains of gravitational waves from a vibrating black hole, *Astrophys. J.* **170**, L105 (1971).
- [136] M. Davis, R. Ruffini, W. H. Press, and R. H. Price, Gravitational radiation from a particle falling radially into a Schwarzschild black hole, *Phys. Rev. Lett.* **27**, 1466 (1971).
- [137] S. Chandrasekhar and S. L. Detweiler, The quasi-normal modes of the Schwarzschild black hole, *Proc. R. Soc. A* **344**, 441 (1975).
- [138] S. L. Detweiler, Resonant oscillations of a rapidly rotating black hole, *Proc. R. Soc. A* **352**, 381 (1977).
- [139] K. D. Kokkotas and B. G. Schmidt, Quasi-normal modes of stars and black holes, *Living Rev. Relativity* **2**, 2 (1999).
- [140] E. Berti *et al.*, Black hole spectroscopy: From theory to experiment, [arXiv:2505.23895](https://arxiv.org/abs/2505.23895).
- [141] E. Berti, V. Cardoso, and A. O. Starinets, Quasinormal modes of black holes and black branes, *Classical Quantum Gravity* **26**, 163001 (2009).
- [142] N. Franchini and S. H. Völkel, Testing general relativity with black hole quasi-normal modes, [arXiv:2305.01696](https://arxiv.org/abs/2305.01696).
- [143] F. Echeverria, Gravitational wave measurements of the mass and angular momentum of a black hole, *Phys. Rev. D* **40**, 3194 (1989).
- [144] L. S. Finn, Detection, measurement and gravitational radiation, *Phys. Rev. D* **46**, 5236 (1992).
- [145] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Tests of general relativity with GW150914, *Phys. Rev. Lett.* **116**, 221101 (2016); *Phys. Rev. Lett.* **121**, 129902(E) (2018).
- [146] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Observation of gravitational waves from a binary black hole merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [147] E. Hamilton, L. London, J. E. Thompson, E. Fauchon-Jones, M. Hannam, C. Kalaghatgi, S. Khan, F. Pannarale, and A. Vano-Vinuales, Model of gravitational waves from precessing black-hole binaries through merger and ringdown, *Phys. Rev. D* **104**, 124027 (2021).
- [148] A. Gupta *et al.*, Possible causes of false general relativity violations in gravitational wave observations, *SciPost Phys. Commun. Rep.*, 5 (2025).
- [149] S. A. Bhat, P. Saini, M. Favata, and K. G. Arun, Systematic bias on the inspiral-merger-ringdown consistency test due to neglect of orbital eccentricity, *Phys. Rev. D* **107**, 024009 (2023).
- [150] Z. Carson and K. Yagi, Testing general relativity with gravitational waves, in *Handbook of Gravitational Wave Astronomy* (Springer, Singapore, 2021).
- [151] N. V. Krishnendu and F. Ohme, Testing general relativity with gravitational waves: An overview, *Universe* **7**, 497 (2021).
- [152] J. M. Bardeen, W. H. Press, and S. A. Teukolsky, Rotating black holes: Locally nonrotating frames, energy extraction, and scalar synchrotron radiation, *Astrophys. J.* **178**, 347 (1972).
- [153] W. Tichy and P. Marronetti, The Final mass and spin of black hole mergers, *Phys. Rev. D* **78**, 081501 (2008).

- [154] C. Hoy and V. Raymond, PESummary: The code agnostic parameter estimation summary page builder, *SoftwareX* **15**, 100765 (2021).
- [155] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Tests of general relativity with GW150914, *Phys. Rev. Lett.* **116**, 221101 (2016); *Phys. Rev. Lett.* **121**, 129902(E) (2018).
- [156] C. Foo and E. Hamilton, Systematic bias due to mis-modelling precessing binary black hole ringdown, *Phys. Rev. D* **110**, 104024 (2024).
- [157] E. Hamilton, L. London, and M. Hannam, Ringdown frequencies in black holes formed from precessing black-hole binaries, *Phys. Rev. D* **107**, 104035 (2023).
- [158] P. Narayan, N. K. Johnson-McDaniel, and A. Gupta, Effect of ignoring eccentricity in testing general relativity with gravitational waves, *Phys. Rev. D* **108**, 064003 (2023).
- [159] C. Hoy, Accelerating multimodel Bayesian inference, model selection, and systematic studies for gravitational wave astronomy, *Phys. Rev. D* **106**, 083003 (2022).
- [160] J. S. Read, Waveform uncertainty quantification and interpretation for gravitational-wave astronomy, *Classical Quantum Gravity* **40**, 135002 (2023).
- [161] L. Pompili, A. Buonanno, and M. Pürrer, Accounting for numerical-relativity calibration uncertainty in gravitational-wave modeling and inference, *arXiv:2410.16859*.
- [162] S. Khan, Probabilistic model for the gravitational wave signal from merging black holes, *Phys. Rev. D* **109**, 104045 (2024).
- [163] R. Bachhar, M. Pürrer, and S. R. Green, Incorporating waveform calibration error in gravitational-wave modeling and inference for SEOBNRv4, *Phys. Rev. D* **111**, 084050 (2025).
- [164] S. Mezzasoma, C.-J. Haster, C. B. Owen, N. J. Cornish, and N. Yunes, Uncertainty-aware waveform modeling for high-SNR gravitational-wave inference, *arXiv:2503.23304*.
- [165] S. Kumar, M. Melching, and F. Ohme, Accounting for the known unknown: A parametric framework to incorporate systematic waveform errors in gravitational-wave parameter estimation, *arXiv:2502.17400*.
- [166] L. Baiotti, T. Damour, B. Giacomazzo, A. Nagar, and L. Rezzolla, Analytic modelling of tidal effects in the relativistic inspiral of binary neutron stars, *Phys. Rev. Lett.* **105**, 261101 (2010).
- [167] T. Damour and A. Nagar, A new analytic representation of the ringdown waveform of coalescing spinning black hole binaries, *Phys. Rev. D* **90**, 024054 (2014).
- [168] S. Bernuzzi, A. Nagar, T. Dietrich, and T. Damour, Modeling the dynamics of tidally interacting binary neutron stars up to the merger, *Phys. Rev. Lett.* **114**, 161103 (2015).
- [169] A. Nagar, G. Riemenschneider, and G. Pratten, Impact of numerical relativity information on effective-one-body waveform models, *Phys. Rev. D* **96**, 084045 (2017).
- [170] A. Nagar *et al.*, Time-domain effective-one-body gravitational waveforms for coalescing compact binaries with nonprecessing spins, tides and self-spin effects, *Phys. Rev. D* **98**, 104052 (2018).
- [171] S. Akcay, S. Bernuzzi, F. Messina, A. Nagar, N. Ortiz, and P. Retegno, Effective-one-body multipolar waveform for tidally interacting binary neutron stars up to merger, *Phys. Rev. D* **99**, 044051 (2019).
- [172] A. Nagar, G. Riemenschneider, G. Pratten, P. Retegno, and F. Messina, Multipolar effective one body waveform model for spin-aligned black hole binaries, *Phys. Rev. D* **102**, 024077 (2020).
- [173] A. Albertini, A. Nagar, P. Retegno, S. Albanesi, and R. Gamba, Waveforms and fluxes: Towards a self-consistent effective one body waveform model for nonprecessing, coalescing black-hole binaries for third generation detectors, *Phys. Rev. D* **105**, 084025 (2022).
- [174] A. Nagar and S. Albanesi, Toward a gravitational self-force-informed effective-one-body waveform model for nonprecessing, eccentric, large-mass-ratio inspirals, *Phys. Rev. D* **106**, 064049 (2022).
- [175] A. Gonzalez, R. Gamba, M. Breschi, F. Zappa, G. Carullo, S. Bernuzzi, and A. Nagar, Numerical-relativity-informed effective-one-body model for black-hole–neutron-star mergers with higher modes and spin precession, *Phys. Rev. D* **107**, 084026 (2023).
- [176] D. Chiamello and A. Nagar, Faithful analytical effective-one-body waveform model for spin-aligned, moderately eccentric, coalescing black hole binaries, *Phys. Rev. D* **101**, 101501 (2020).
- [177] A. Nagar, A. Bonino, and P. Retegno, Effective one-body multipolar waveform model for spin-aligned, quasicircular, eccentric, hyperbolic black hole binaries, *Phys. Rev. D* **103**, 104021 (2021).
- [178] A. Nagar and P. Retegno, The next generation: Impact of high-order analytical information on effective one body waveform models for noncircularized, spin-aligned black hole binaries, *Phys. Rev. D* **104**, 104004 (2021).
- [179] T. Andrade *et al.*, Toward numerical-relativity informed effective-one-body waveforms for dynamical capture black hole binaries, *Phys. Rev. D* **109**, 084025 (2024).
- [180] A. Nagar, R. Gamba, P. Retegno, V. Fantini, and S. Bernuzzi, Effective-one-body waveform model for non-circularized, planar, coalescing black hole binaries: The importance of radiation reaction, *Phys. Rev. D* **110**, 084001 (2024).
- [181] S. Albanesi, R. Gamba, S. Bernuzzi, J. Fontbuté, A. Gonzalez, and A. Nagar, Effective-one-body modeling for generic compact binaries with arbitrary orbits, *arXiv:2503.14580*.
- [182] A. Gamboa *et al.*, Accurate waveforms for eccentric, aligned-spin binary black holes: The multipolar effective-one-body model SEOBNRv5EHM, *arXiv:2412.12823*.
- [183] H. Yu, J. Roulet, T. Venumadhav, B. Zackay, and M. Zaldarriaga, Accurate and efficient waveform model for precessing binary black holes, *Phys. Rev. D* **108**, 064059 (2023).
- [184] M. d. L. Planas, A. Ramos-Buades, C. García-Quirós, H. Estellés, S. Husa, and M. Haney, Time-domain phenomenological multipolar waveforms for aligned-spin binary black holes in elliptical orbits, *arXiv:2503.13062*.