



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Model of gravitational waves from precessing black-hole binaries through merger and ringdown

Eleanor Hamilton, Lionel London, Jonathan E. Thompson, Edward Fauchon-Jones, Mark Hannam, Chinmay Kalaghatgi, Sebastian Khan, Francesco Pannarale, and Alex Vano-Vinuales

Phys. Rev. D **104**, 124027 — Published 9 December 2021

DOI: [10.1103/PhysRevD.104.124027](https://doi.org/10.1103/PhysRevD.104.124027)

The final twist: A model of gravitational waves from precessing black-hole binaries through merger and ringdown

Eleanor Hamilton,^{1,2} Lionel London,³ Jonathan E. Thompson,¹ Edward Fauchon-Jones,¹ Mark Hannam,¹ Chinmay Kalaghatgi,^{4,5,6} Sebastian Khan,¹ Francesco Pannarale,^{7,8} and Alex Vano-Vinuales⁹

¹*School of Physics and Astronomy, Cardiff University, Cardiff, CF24 3AA, United Kingdom*

²*Physik-Institut, Universität Zürich, Winterthurerstrasse 190, 8057 Zürich, Switzerland*

³*MIT-Kavli Institute for Astrophysics and Space Research and LIGO*

Laboratory, 77 Massachusetts Avenue, 37-664H, Cambridge, MA 02139, USA

⁴*Nikhef – National Institute for Subatomic Physics, Science Park, 1098 XG Amsterdam, The Netherlands*

⁵*Institute for Gravitational and Subatomic Physics (GRASP), Utrecht*

University, Princetonplein 1, 3584 CC Utrecht, The Netherlands

⁶*Institute for High-Energy Physics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands*

⁷*Dipartimento di Fisica, Università di Roma “Sapienza”, Piazzale A. Moro 5, I-00185, Roma, Italy*

⁸*INFN, Sezione di Roma, Piazzale A. Moro 5, I-00185, Roma, Italy*

⁹*Centro de Astrofísica e Gravitação - CENTRA, Departamento de Física, Instituto Superior Técnico IST, Universidade de Lisboa UL, Avenida Rovisco Pais 1, 1049-001 Lisboa, Portugal*

We present PHENOMP NR, a frequency-domain phenomenological model of the gravitational-wave (GW) signal from binary-black-hole mergers that is tuned to numerical relativity (NR) simulations of precessing binaries. In many current waveform models, e.g., the “PHENOM” and “EOBNR” families that have been used extensively to analyse LIGO-Virgo GW observations, analytic approximations are used to add precession effects to models of non-precessing (aligned-spin) binaries, and it is only the aligned-spin models that are fully tuned to NR results. In PHENOMP NR we incorporate precessing-binary NR results in two ways: (i) we produce the first NR-tuned model of the signal-based precession dynamics through merger and ringdown, and (ii) we extend a previous aligned-spin model, PHENOMD, to include the effects of misaligned spins on the signal in the co-precessing frame. The NR calibration has been performed on 40 simulations of binaries with mass ratios between 1:1 and 1:8, where the larger black hole has a dimensionless spin magnitude of 0.4 or 0.8, and we choose five angles of spin misalignment with the orbital angular momentum. PHENOMP NR has a typical mismatch accuracy within 0.1% up to mass-ratio 1:4, and within 1% up to mass-ratio 1:8.

I. INTRODUCTION

Binary black hole (BBH) mergers are the primary source of gravitational waves observable with current ground-based detectors [1, 2]; of the 51 detections published by the LIGO-Virgo collaborations, 48 were confirmed as BBH [3–7]. Measurements of each binary’s properties — the black-hole (BH) masses and spins, and the location of the binary — rely in part on models of the signal predicted by general relativity. Model development is an active research area, with the aim that the measurement uncertainties due to model errors, approximations, and incomplete physics are smaller than statistical errors arising from the strength of the signal above the detector noise, or parameter degeneracies. Models are informed by analytic approximations for the inspiral of the two BHs and ringdown of the final BH, and NR solutions of Einstein’s equations for the late inspiral, merger and ringdown. One key physical effect is the precession of the binary’s orbital plane due predominantly to spin-orbit effects, but the two waveform families most commonly used for LIGO-Virgo parameter estimates, “PHENOM” [8–19] and “EOBNR” [20–26], have *not* been tuned to NR simulations of precessing binaries. Instead, precession effects during the strongest part of the signal have been estimated using simple approximations. These were likely

sufficient for observations to date, but, given that they do not capture several physical features of the merger signal (e.g., Ref. [27], plus other effects that we will describe in this paper) more accurate models will ultimately be required.

Here we present the first PHENOM model where merger-ringdown precession effects are explicitly tuned to NR simulations. We show that this model is in general significantly more accurate than previous models, particularly for binaries with large mass ratios, high spins, and a large spin misalignment.

A BBH system following non-eccentric inspiral is defined by the BH masses, m_1 and m_2 (we choose $m_1 > m_2$), and the BH spin-angular-momentum vectors \mathbf{S}_1 and \mathbf{S}_2 . As is standard, we choose the alternative parameterisation into total mass, $M = m_1 + m_2$, symmetric mass ratio $\eta = m_1 m_2 / M^2$, and the dimensionless spins $\chi_i = \mathbf{S}_i / m_i^2$, where $|\chi_i| \in [0, 1]$ respects the Kerr limit. It is also convenient to decompose the spins into their components parallel and perpendicular to the direction of the Newtonian orbital angular momentum, $\hat{\mathbf{L}}_N$, i.e., the magnitudes of the spins parallel to $\hat{\mathbf{L}}_N$ are $\chi_i^\parallel = \chi_i \cdot \hat{\mathbf{L}}_N$, and the components that lie in the orbital plane are $\chi_i^\perp = \chi_i - \chi_i^\parallel \hat{\mathbf{L}}_N$.

If the spins are parallel to the orbital angular momentum, i.e., $\chi_i^\perp = 0$, then the orientation of the binary’s

orbital plane, and the directions of the spin and orbital angular momenta, are all fixed. Waveforms from these aligned-spin, or non-precessing, binaries, have been modelled with a combination of post-Newtonian (PN) and effective-one-body (EOB) results to describe the inspiral, and NR results to model the late inspiral, merger and ringdown, to produce PHENOM and EOBNR waveform models [8, 9, 14, 15, 18–20, 23]. Surrogate models of non-precessing systems have also been constructed purely from NR waveforms, and also from PN-NR hybrids [28, 29].

When $\chi_i^\perp \neq 0$, the binary precesses. In most cases the binary undergoes simple precession [30, 31], where the orbital angular momentum and spins precess around the binary’s total angular momentum, which points in an approximately fixed direction. Precession modulates the amplitude and phase of the gravitational-wave signal, and leads to a significantly more complicated signal than in non-precessing configurations. However, if we transform to a non-inertial co-precessing frame that tracks the precession, then the signal recovers, to a good approximation, the simple form of a non-precessing signal [32], and, indeed, during the inspiral the co-precessing-frame waveform is approximately the signal from the corresponding non-precessing binary defined by setting $\chi_i^\perp = 0$ [33].

This observation has been used to construct current PHENOM and EOBNR waveform models, by using a non-precessing model as a proxy for the precessing-binary waveform in the co-precessing frame, and then transforming this to the inertial frame via an independent model for the precession dynamics [10, 12, 13, 16, 21, 22, 25]. Although some NR information from precessing-binary simulations has been used to model the final state [25], the precession effects have *not* been tuned to NR waveforms, and neither have in-plane-spin contributions to the co-precessing-frame signal. In addition to these models, surrogate models of precessing binaries have been constructed using NR waveforms that cover roughly 20 orbits before merger [28, 34, 35]. This puts an explicit limit on their applicability to comparatively short signals, i.e., from high-mass binaries with near-equal masses.

The current work extends the PHENOM approach, the development of which has proceeded in order of the most measurable physical effects. The most clearly measurable binary parameters are the chirp mass, $\mathcal{M} = M\eta^{3/5}$, for low-mass binaries where the detectable signal is dominated by the inspiral, and the total mass M for high-mass binaries where most of the detectable signal power is in the late inspiral, merger and ringdown. Hence the first PHENOM model considered non-spinning binaries [36, 37]. The next most significant effect is due to a mass-weighted combination of the aligned-spin components, and the next set of PHENOM models treated aligned-spin systems and were tuned to NR simulations that were parametrised by a single effective spin [8, 9, 38, 39]. All of these models considered only the dominant contribution to the signal, which is from the ($\ell = 2, |m| = 2$) multipole moments. Subdominant mul-

tipoles become stronger as the mass ratio is increased, and these were first included through an approximate mapping of the dominant multipole [11], and more recently with full tuning to NR simulations [15]. Individual black-hole spins are unlikely to be measurable for detections with a signal-to-noise ratio (SNR) of less than ~ 100 [40], but a handful of such detections are likely when the LIGO and Virgo detectors reach design sensitivity in the next few years [41], and the latest aligned-spin PHENOM models include NR tuning to unequal-spin NR simulations [14]. The PHENOM approach has been predominantly used to produce frequency-domain models, but has recently also been applied in the time domain [18, 19].

Precession effects are typically difficult to measure (as discussed in detail in Ref. [42], and illustrated with parameter-estimation examples in Ref. [43, 44]), and indeed have not yet been definitively observed in any single observation [3, 7]. The dominant precession effects follow the phenomenology of single-spin systems, and thus the first precessing PHENOM models [10] used a single-spin PN model to estimate the effects of precession. More recent models have included two-spin effects [12, 13, 16, 45], but, once again, individual spin measurements will require SNRs of at least 100, and in most cases likely much higher [13]. As such, the first priority for an NR-tuned precession model is the single-spin parameter space. Our new PHENOMPNR model is tuned to NR simulations that cover mass ratios from equal-mass to 1:8 ($\eta \sim 0.1$). The larger black hole has a spin magnitude up to $\chi_1 = 0.8$, and, as motivated by the preceding discussion, the smaller black hole has no spin. This is the widest systematic coverage of the mass-ratio-spin parameter space to date [46].

A. Model approximations, and motivation for a new model

Previous PHENOM and EOBNR models make use of several approximations. In this section we discuss each of these, and illustrate why we remove some of them in our new model, and the effect this has on the waveforms.

One set of approximations applies to the waveforms in the co-precessing frame.

First, as described above, during the inspiral the co-precessing-frame waveform is approximated by an equivalent non-precessing-binary waveform, h^{NP} . In the most recent EOBNR model, SEOBNRv4PHM [25], the EOB equations of motion are solved for the full precessing system from a chosen starting frequency, and then the approximate co-precessing-frame waveform is constructed by now solving the non-precessing PN equations of motion, but with time-varying $\chi_i^\parallel(t)$ taken from the earlier precessing-binary solution. In the PHENOM models, h^{NP} is defined by the aligned-spin components of the initial spin configuration, so χ_i^\parallel are constant. In both families of models, χ_i^\perp contributions to the waveform multipole

moment amplitudes are ignored.

Second, the mapping to an equivalent aligned-spin system breaks down at merger. This was already noted in the original presentation of the aligned-spin mapping [33], and is also discussed in Refs. [27, 47]. One reason is that the spin of the final black hole (and therefore the ringdown frequency and damping time) will be different to that in the non-precessing case; to first approximation, we must include the contribution from the in-plane spins, χ_i^\perp , to the spin of the final BH. In the PHENOM models, the merger-ringdown part of the aligned-spin waveform is modified by using this in-plane spin contribution to estimate a modified final spin, and hence complex ringdown frequency [10, 12, 13, 16]; the recent PHENOMXP model [16] provides a number of optional methods to achieve this. In the EOBNR models, the inspiral construction ends at the light ring [23], and ringdown modes are attached, and in the most recent SEOBNRv4PHM model [25] these are based on an NR-tuned final spin fit [48].

In PHENOMPNR, we retain the mapping to an equivalent aligned-spin system during the early inspiral, but we introduce the key improvement that in the late inspiral, merger and ringdown we explicitly tune the model to NR waveforms in the co-precessing frame. Rather than model the final mass and spin and use those to estimate the complex ringdown frequency via perturbation theory, we also explicitly model the ringdown frequencies from NR waveforms in the co-precessing frame. As discussed in Sec. IX, this is necessary because the ringdown frequency in the co-precessing frame is shifted with respect to that in the inertial frame.

This issue is illustrated in Fig. 1. The top panel shows the frequency-domain co-precessing-frame phase derivative $d\phi_{22}/df$ for one of our NR simulations, with mass-ratio $q = m_1/m_2 = 4$, large-black-hole spin $\chi_1 = 0.8$, and spin mis-aligned with the orbital angular momentum by $\theta_{\text{LS}} = 60^\circ$. The figure also shows the results from the earlier PHENOMPv3 model. In the inspiral we see a clear difference between the NR and PHENOMPv3 results that is largest at low frequencies. The middle panel shows a second case, this time with a larger misalignment angle of $\theta_{\text{LS}} = 150^\circ$. The location of the minimum can be approximately identified as the ringdown frequency, and we see that there is a clear shift between the ringdown frequency in the inertial frame (as used in PHENOMPv3), and the effective ringdown frequency of the NR waveform in the co-precessing frame. This shift is also apparent in the bottom panel, which shows the amplitude A_{22} in the co-precessing frame. PHENOMPNR fixes this problem; see, in particular, Sec. V.

A second set of assumptions apply to the precession.

In previous models the inertial-frame waveform was constructed via a time- or frequency-dependent rotation of h^{NP} , using the precession angles relative to the Newtonian orbital angular momentum, i.e., the normal to the binary's orbital plane. This produces the correct inertial-frame multipoles only in the quadrupole

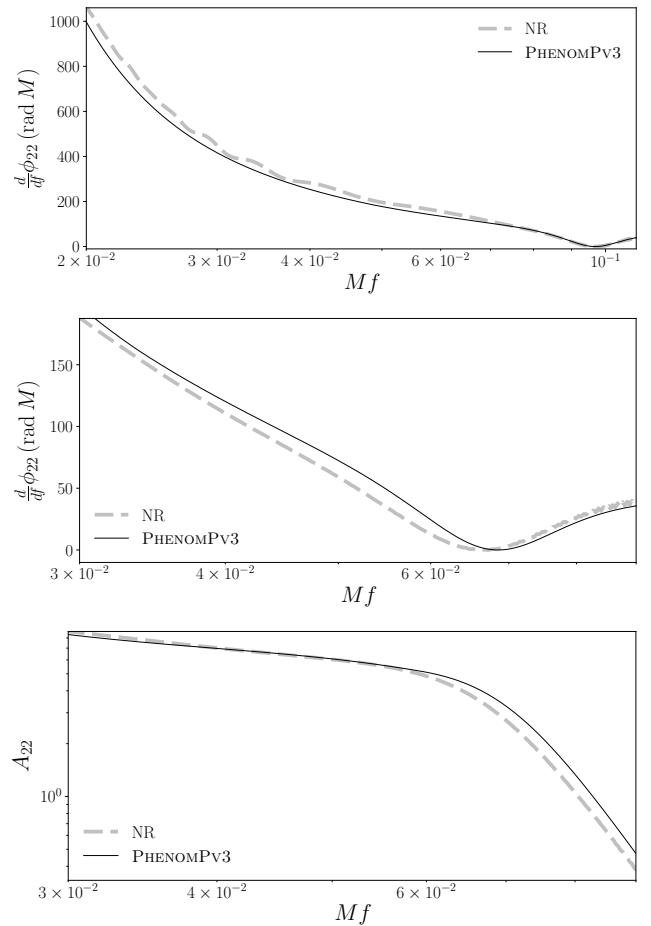


FIG. 1. Frequency domain comparison of NR and model waveforms in the co-precessing frame. (top) phase derivative for the $(q, \chi_1, \theta_{\text{LS}}) = (4, 0.8, 60^\circ)$ configuration, which illustrates the variation in the inspiral phase. (middle and bottom) phase derivative and amplitude for the $(q, \chi_1, \theta_{\text{LS}}) = (4, 0.8, 150^\circ)$ configuration, which demonstrate the shift in effective ringdown frequency.

approximation. In order to tune the precession angles to NR results, we need a consistent choice of co-precessing frame that can be applied both to PN and NR data. For PHENOMPNR we choose the quadrupole-aligned (QA) frame [32, 49, 50], which identifies the direction of maximum GW emission. In time-domain waveforms, the direction of maximum emission differs depending on whether it was defined using GW strain, h , the Bondi news function, \dot{h} , or the Weyl scalar, $\Psi_4 = \dot{h}$; and all three differ from the direction of the orbital angular momentum \mathbf{L} [32, 51–53]. (The direction of \mathbf{L} also depends on whether we use a Newtonian or post-Newtonian estimate.) However, we perform our modelling in the frequency domain, where the QA direction is independent of the choice of h or Ψ_4 . We explain this further in Sec. III, where we also describe in detail how we calculate the QA frame from the $\ell = 2$ multipoles of NR simulations, and in Sec. VIB we discuss the QA frame for PN waveforms.

We expect that the latter results would also allow the construction of more physically accurate EOBNR waveforms.

In most previous PHENOM models, the precession angles were estimated entirely from PN theory. These angles will not be valid through merger, but as a simple approximation, they were used throughout the entire waveform. This approximation was justified by the observation that the PN angles behave smoothly to arbitrarily high frequencies, and the model gives reasonable agreement to NR waveforms [10, 12, 13, 16]. However, in more extreme parts of parameter space (high mass ratios and large in-plane spins), the inaccuracy of this approximation will become more serious. In EOBNR models, the inspiral precession dynamics are provided from the solution of the EOB equations of motion, and in the SEOBNRv4PHM model the precession angles are extended through merger and ringdown using an approximation based on the quantitative behaviour of NR simulations; the time-domain PHENOM model, PHENOMTPHM, employs a similar approach [18].

Fig. 2 shows the precession angles (α, β, γ) for a configuration with $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 60^\circ)$. The figure shows both the NR results, and the multi-scale analysis (MSA) angles [54] used in the PHENOMPv3 and PHENOMXP models. We see that at high frequencies that correspond to the merger and ringdown, the MSA estimates fail to capture the phenomenology of the NR data. The angles α and γ both exhibit a “dip” or “bump”, reminiscent of the dip in the phase derivative in Fig. 1, which is absent in the MSA estimates. The NR opening angle β drops to close to zero at merger, as we might expect as the two-body inspiral motion terminates and we are left with only a single perturbed black hole. This feature cannot be captured by the MSA expressions, which simply extend the inspiral behaviour to higher frequencies. We also find that the NR β does not relax to zero, but to some non-zero value, which, if it does decay, typically does so very slowly. (There have been approximate estimates of this asymptotic β decay using a toy ring-down model [18, 55, 56], which we discuss and clarify in Sec. IX.) These features must also be modelled.

Finally, we see that at lower frequencies, the MSA α and γ agree well with the NR results. However, although we expect the MSA and NR β to also agree at sufficiently low frequencies, they do not agree over the frequency range of our NR data, and would likely require NR simulations that are many times longer. This discrepancy is due to the modelling inconsistency discussed earlier: the two estimates are of different quantities. The MSA β is the orientation of the orbital plane, while the NR β is the orientation of the QA direction of the signal, and these are not in general the same. We show how to significantly reduce this discrepancy in Sec. VIB. (The high-frequency oscillations in the NR β are due to a combination of numerical noise and Fourier-transform artifacts. All of our NR β results show similar oscillations, with varying amplitude and frequency, but in these single-spin cases we

will model only a smooth trend through the data, which we expect to represent their relevant physical features.)




FIG. 2. Comparison of the post-Newtonian expressions for each of the precession angles (blue dotted line) with the NR data (black solid line) for the case with $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 60^\circ)$. The grey vertical lines indicate the ISCO frequency ($Mf = 0.0287$) of the final black hole, which has final spin magnitude $\chi_f = 0.799$ and final mass $M_f = 0.981M$.

The bulk of the results in this paper present a merger-ringdown model for the co-precessing-frame waveforms (PHENOMDCP) and a separate model for the precession angles (PHENOMANGLES). Both modes are tuned to our NR data and capture all of the features described here. We then produce a complete inspiral-merger-ringdown model (PHENOMPvNR) by connecting our merger-ringdown models to inspiral results.

There are two remaining assumptions that were made in previous models, which we retain in our new model.

Non-precessing-binary waveforms satisfy a symmetry between the $m > 0$ and $m < 0$ multipoles that is broken in precessing binaries [27, 57, 58]. The “twisting-up” construction used by the PHENOM and EOBNR models neglects these asymmetries. Although asymmetries may need to be included in models to allow accurate spin measurements in some GW observations [58], in the current PHENOPNR model we retain the approximation that the asymmetries in the multipole moments are zero.

Current PHENOM and EOBNR models also assume that the direction of the total angular momentum remains fixed. Although the total angular momentum direction changes little through inspiral, there is *some* change due to the loss of angular momentum through GW emission. In PHENOPNR we explicitly transform the NR waveforms to a frame where $\hat{\mathbf{J}}$ remains fixed along the z -axis, and use those waveforms as the basis of the model. In this sense the fixed- $\hat{\mathbf{J}}$ approximation is retained in PHENOPNR and remains valid over the parameter space used to construct the model, which is further discussed in Sec. XI E.

This paper is organised as follows. In Sec. II we present our NR waveforms. In Sec. III we process the raw NR waveforms to produce the frequency-domain co-precessing-frame waveforms and precession angles that we wish to model. Since we limit the NR tuning to single-spin binaries, in Sec. IV we specify our procedure to map generic two-spin systems to approximately equivalent single-spin configurations. With all of these pieces in place, in Sec. V we present our co-precessing-frame model, PHENODCP, in Sec. VI our treatment of the precession angles during inspiral, and in Sec. VII our merger-ringdown angle model, PHENOMANGLES. All of these ingredients are put together into a full inspiral-merger-ringdown model in Sec. VIII. Having modelled precessing-binary waveforms, we discuss their physical features in more detail in Sec. IX, and evaluate their accuracy in Sec. XI.

In all of the discussion of NR and PN results, and in all modelling work, we use geometric units, $G = c = 1$. We also choose $M = 1$, although we retain “ M ” in plot labels, to make clear that we are dealing with dimensionless quantities. Physical masses will only be used in Sec. XI, where we study the performance of models with respect to a specific detector noise curve. All angles displayed in figures throughout the paper are given in radians unless otherwise specified. All of the earlier waveform models used to generate results in this work were called from the software package `LALSuite` [59]. The specific model names are `IMRPhenomD` for PHENOMD [8, 9], `IMRPhenomXAS` for PHENOMXAS [14], `IMRPhenomPv3` for PHENOMPv3 [12], `IMRPhenomXP` for PHENOMXP [16], `SEOBNRv4P` for SEOBNRv4P [25], and `NRSur7dq4` for NRSUR7DQ4 [35].

II. NUMERICAL RELATIVITY WAVEFORMS

In producing the first precessing-binary model tuned to NR waveforms, we wish to capture the dominant precession effects first. This can be achieved with single-spin systems, i.e., only one of the black holes is spinning, since two-spin effects typically produce only small modulations of the underlying simple precession [60, 61]. We therefore consider single-spin systems that obey simple precession, and the NR catalogue used to tune the model contains single-spin configurations where the spin is placed on the larger black hole and neglects two-spin configurations and the impact of the azimuthal spin angle. This reduces the binary parameter space from seven dimensions (mass ratio, plus the vector components of each black-hole spin), to three dimensions: the symmetric mass ratio, η , the magnitude of the spin on the larger black hole, $\chi \equiv \chi_1$, and the angle between the spin and the orbital angular momentum of the system, θ_{LS} . It is important to note that these are all defined as part of the initial data of the simulations, since θ_{LS} undergoes small oscillations about some mean value during the inspiral.

We wish our model to extend to the highest mass ratios feasible with current NR simulations. The earlier tuned non-precessing model PHENOMD [8, 9] was based on a catalogue containing systems up to mass ratio $q = m_1/m_2 = 18$, or $\eta \sim 0.05$. NR simulations at $q = 18$ are extremely computationally expensive, and since the mass-ratio of observations is heavily skewed towards comparable masses [3, 7], for the current model we restrict to $q = 8$. We note, however, that one recent GW observation, GW190814, was measured with a mass ratio of $q \sim 10$ [62], and therefore extending our model to higher mass ratios is an urgent requirement for future work.

In order to confidently capture the dependence of precession effects on mass ratio, we produced simulations at four different mass ratios, approximately equally spaced in symmetric mass ratio η . Similarly, we chose four equally spaced spin magnitudes χ . We already have aligned and anti-aligned waveforms in this range of mass ratios and spin magnitudes, and for non-aligned-spin configurations we chose five equally spaced values for the spin angle, θ_{LS} , excluding 0° and 180° .

The model is tuned to a subset of this catalogue of 80 waveforms, which was produced using the `BAM` code [63]. The complete catalogue contains simulations with $q \in [1, 2, 4, 8]$, (or $\eta \in [0.1, 0.16, 0.22, 0.25]$), $\chi \in [0.2, 0.4, 0.6, 0.8]$ and $\theta_{\text{LS}} (^\circ) \in [30, 60, 90, 120, 150]$. For tuning we used the 40 waveforms with $\chi = 0.4$ and 0.8 . We expect the dependence of the precession effects on spin magnitude to be approximately linear, so this is not anticipated to significantly degrade the accuracy of the tuned part of the model; and this is borne out in validation of the model against the remaining waveforms in the catalogue, plus 27 waveforms from the `SXS` and `Maya` catalogues [64–67].

Since our goal is a frequency-domain model, we would like NR waveforms that all cover a similar frequency

range. The majority of the waveforms start at a frequency of $M\Omega = 0.023$. However, some of the higher mass ratio configurations have a higher starting frequency in order to ensure the binary merged in a reasonable time to allow sufficient accuracy. The highest starting frequencies occur for configurations with a large spin magnitude where the spin is closest to being aligned with the orbital angular momentum, due to the hang-up effect [68]. The highest starting frequency is $M\Omega = 0.032$, for the $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 30^\circ)$ configuration. We find that these starting frequencies are in general sufficient to match smoothly to PN results. We will see in Sec. XID that there are a few cases for which we would prefer NR waveforms with lower starting frequencies, but these are actually configurations with large spins and large opening angle, e.g., $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 150^\circ)$. Having identified specific issues with these more challenging regions of parameter space, we will be able to focus on them in detail in future iterations of our model.

More details on the production of the NR catalogue, and error analysis of the waveforms, will be given in Ref. [46]. The greatest sources of error in these numerical waveforms are the finite resolution at which we performed the simulations and the finite distance from the source at which we extracted the GW data. We consider the mismatch (as defined in Sec. XIA) to be the most useful uncertainty estimate for our purposes. Using separate convergence analyses (consistent with fourth order convergence with respect to resolution and a $1/r$ fall off with respect to extraction radius) allows us to make a conservative estimate of the mismatch uncertainty between the waveforms in this NR catalogue and the theoretical ‘analytical’ solution of $\mathcal{O}(10^{-3})$. For the shorter waveforms in the catalogue, particularly the $q = 1$ and $q = 2$ cases, the mismatch was found to be $\mathcal{O}(10^{-4})$. As we will see when validating against independent NR data sets (e.g., those from the SXS catalogue, where the finite-extraction-radius error is minimal), the errors in our model are often an order of magnitude lower than our upper bound, and, where they are comparable or higher, the accuracy limits due to the modelling procedure are likely the dominant source of error.

For each NR simulation, spin weight -2 spherical harmonic multipole moment data are stored for the radiative Weyl scalar,

$$\psi_{\ell m}(t) = \int_{\Omega} r \Psi_4(t, r, \theta, \phi) {}_{-2}Y_{\ell m}^*(\theta, \phi) d\Omega, \quad (1)$$

where $*$ denotes complex conjugation. The $\psi_{\ell m}$ depend on the choice of decomposition frame, and we provide the details of our frame choice in Sec. III. Each $\psi_{\ell m}$ time series contains multipole moment data for inspiral, merger and ringdown.

In addition, spurious (“junk”) radiation, due to imperfect initial data [69], is windowed away, using a window function that increases from zero to one over the duration of three gravitational wavelengths. It is found that when windowing over more than two wavelengths the choice

of (smooth) window function has no significant effect on our modelling results. For simplicity, a standard Hann window is used [70]. The window starts at the first peak in the real part of ψ_{22} such that the following peak is less than or equal to the largest distance between peaks in the time series. This most often results in less than $200M$ of contaminated inspiral data being tapered away. The window is applied equally to the real and imaginary parts of Ψ_4 for all multipoles. Similarly, post-ringdown data are windowed such that the Hann window turns off to the right between the point where the exponential decay drops below the noise floor, as defined by fitting a constant value to the very end of the timeseries. The time domain data are also zero-padded to the right such that the frequency domain step size, in geometric units, is less than 5×10^{-4} .

The result of the inspiral and post-ringdown windows is the reduction of frequency-domain power that is broadband and unphysical. The result of zero-padding is to enforce that frequency domain features are consistently resolved.

III. WAVEFORM FRAMES, CONVENTIONS AND APPROXIMATIONS

We wish to model the dominant multipoles of the BBH signal. The multipoles depend on the choice of reference frame, and we attempt to choose a frame that simplifies the modelling. In this section we present the reference frame in which we construct our model, and several additional simplifications that we make to the data.

If we have a set of spin-weighted spherical-harmonic multipoles $q_{\ell m}^1$, and rotate the coordinate system through the Euler angles (α, β, γ) , then the multipoles in the new frame, $q_{\ell m}^2$, are given by,

$$q_{\ell m}^2 = \sum_{m'=-\ell}^{\ell} e^{im'\alpha} d_{m'm}^{\ell}(-\beta) e^{im\gamma} q_{\ell m'}^1, \quad (2)$$

where $d_{m'm}^{\ell}$ are the Wigner d-matrices [63, 71].

We apply these rotations twice to our data.

First, we retain the approximation that has been used in all PHENOM and EOBNR models to date, that the direction of the total angular momentum, $\hat{\mathbf{J}}$, is fixed. This convention amounts to a minor modification of the NR data, whose radiative $\mathbf{J}(t)$ varies by at most $\sim 6^\circ$ from its initial direction. To impose the fixed- $\hat{\mathbf{J}}$ convention we need to know $\mathbf{J}(t)$ at all times in the original simulation. At the beginning of the simulation $\mathbf{J}(0) = \mathbf{J}_{\text{ADM}}$, which can be calculated analytically from Bowen-York initial data [72]. The angular momentum flux can be calculated from the multipole moments, e.g., Ref. [73], and integrating this specifies the time evolution of $\mathbf{J}(t)$. As a consistency check, we compare \mathbf{J} at the end of the simulation with the estimate of the final black hole’s spin calculated on the apparent horizon [74], and find a disagreement of at most 5% in magnitude and 3% in direc-

tion. With $\mathbf{J}(t)$ now in hand, we use Eq. (2) to perform a time-dependent rotation to place the signal in a frame of reference where $\hat{\mathbf{J}}(t) = \hat{z}$ at all times. The impact of this frame convention is well below the total error budget of the final PHENOMPNR model, and is discussed in more detail in Sec. XI E.

Second, we make another time-dependent rotation into a co-precessing frame. We choose the QA frame, which was introduced in Ref. [32], and allows us to define a co-precessing frame using the gravitational-wave signal, which is the observable quantity we ultimately care about, rather than the orbital dynamics of the two black holes. The QA method was motivated by the observation that in the quadrupole approximation, if the orbital plane lies in the x - y plane, then the signal can be represented entirely by the $(\ell = 2, |m| = 2)$ multipoles. At any other orbital plane orientation, some signal power will be distributed to the $|m| = 1$ and $m = 0$ multipoles, therefore reducing the amplitude of the $(\ell = 2, |m| = 2)$ multipoles. It follows that we can always identify the orientation of the orbital plane by locating the direction with respect to which the $(\ell = 2, |m| = 2)$ multipoles are maximised. In a time-dependent co-precessing frame where this always holds, we can represent the entire signal using only the $|m| = 2$ multipoles, and, furthermore, precession modulations of the signal amplitude and phase will be significantly reduced. In general, i.e., beyond the quadrupole approximation, this direction is only approximately equal to the normal to the orbital plane, or to a PN estimate of the direction of the orbital angular momentum [32, 52, 53]. However, although it cannot be directly related to the dynamics, it does provide us with a convenient signal-based definition of a co-precessing frame that suppresses precession modulations.

In the following sections we use the method described in Appendix A to calculate the coprecessing frame [49, 50]. We use the Euler angles α , β and γ to describe the orientation of this direction. Equations (A5)-(A7) define the angles accordingly, and Fig. 3 illustrates their geometric meaning.

One potential ambiguity with the QA frame is that it differs depending on whether it is defined using the gravitational wave strain, or its time derivatives, the Bondi news \dot{h} or the Newman-Penrose scalar Ψ_4 . However, this ambiguity does not exist in the frequency domain.

To see this, consider the multipoles of the gravitational-wave strain, which can be written as,

$$h_{\ell m}(t) = A_{\ell m}(t)e^{-im\Phi(t)}. \quad (3)$$

Our NR data satisfy $\Psi_4 = \ddot{h}$, and so we can write,

$$\psi_{\ell m}(t) = A'_{\ell m}(t)e^{-im\Phi'(t)}, \quad (4)$$




FIG. 3. The Euler angles (α, β, γ) that make up the precession angles that describe the transformation from the fixed- $\hat{\mathbf{J}}$ frame into a co-precessing frame. There are different choices for the definition of \mathbf{V} ; the QA direction, the Newtonian orbital angular momentum and varying orders of the post-Newtonian orbital-angular momentum. These directions are all approximately the same, as discussed in the text.

where the new amplitude and phase are given by,

$$A'_{\ell m} = \sqrt{\left(\ddot{A} - m^2\dot{\Phi}^2 A\right)^2 + m^2 \left(2\dot{\Phi}\dot{A} + \ddot{\Phi}A\right)^2}, \quad (5)$$

$$\Phi' = \Phi + \frac{1}{m} \arctan \left(\frac{m \left(2\dot{\Phi}\dot{A} + \ddot{\Phi}A\right)}{\ddot{A} - m^2\dot{\Phi}^2 A} \right), \quad (6)$$

where we have dropped the (ℓ, m) subscripts for brevity. We see that the distribution of power between the multipoles will in general be different for h and for Ψ_4 in the time domain, and therefore the QA angles (α, β, γ) will differ.

By contrast, if we take a Fourier transform of the signal then in the frequency domain we have,

$$\tilde{\Psi}_4(\omega) = \text{F.T.}[\Psi_4(t)] = \text{F.T.}[\ddot{h}(t)] = -\omega^2 \tilde{h}(\omega), \quad (7)$$

where we are using the property of Fourier transforms that if $\text{F.T.}[g(t)] = \tilde{G}(\omega)$ then $\text{F.T.}[\dot{g}(t)] = i\omega\tilde{G}(\omega)$ where ω is the Fourier frequency. Since ω is an overall factor in front of all of the multipoles at a given frequency, the direction that maximises both $|\tilde{h}|^2$ and $\omega^4|\tilde{h}|^2$ will be the same. The QA precession angles will therefore be the same for h and for Ψ_4 . Given that the frequency-domain QA angles are independent of the choice of Ψ_4 or strain, we consider this to be the natural regime in which to work.

Finally, we also retain the standard PHENOM and EOBNR approximation that the co-precessing multipole moments of our model obey the same symmetry properties as their non-precessing counterparts. This means

that we neglect to model $\pm m$ asymmetries in the multipole moments. Although the asymmetric contributions are weak, there is some evidence that they are necessary for non-biased measurements of precessing systems [58], and they are certainly necessary for measurements of out-of-plane recoil of the binary [75], and we plan to model these contributions in future work.

Given $\psi_{\ell m}$ that have been transformed first to the fixed- $\hat{\mathbf{J}}$ and then QA frames in the time domain, we construct the symmetric combination,

$$\psi_{2,2}^{\text{sym}} = \frac{1}{2} (\psi_{2,2} + \psi_{2,-2}^*). \quad (8)$$

In Eq. (8), $\psi_{2,2}^{\text{sym}}$ effects an average of the co-precessing-frame mass-quadrupoles consistent with Ref. [52]. We then define a symmetrised ($\ell = 2, m = -2$) multipole according to the non-precessing symmetry relationship $\psi_{\ell, -m} = (-1)^\ell \psi_{\ell m}^*$, thus,

$$\psi_{2,-2}^{\text{sym}} = (\psi_{2,2}^{\text{sym}})^*. \quad (9)$$

Together, $\psi_{2,-2}^{\text{sym}}$ and $\psi_{2,2}^{\text{sym}}$ encapsulate all waveform information that will be retained at this stage. The QA-frame $\ell > 2$ multipoles are discarded, along with the ($\ell = 2, |m| < 2$) multipoles; we leave higher multipoles to future work.

The symmetrised multipoles are then rotated back into the fixed- $\hat{\mathbf{J}}$ frame. We then use these data as our starting point to transform the multipoles into the frequency domain, and then transform to the QA frame as defined in the frequency domain.

We separately produce a model (PHENOMDCP) of the co-precessing-frame multipole $h_{2,2}^{\text{CP}}(f)$, and another model (PHENOMANGLES) of the rotation angles ($\alpha(f), \beta(f), \gamma(f)$). Given these two models, our full inertial-frame model (PHENOMP NR) of the $\ell = 2$ multipoles, $h_{\ell m}^{\text{J}}(f; \boldsymbol{\lambda})$, is given via Eq. (2),

$$h_{\ell m}^{\text{J}}(f; \boldsymbol{\lambda}) = \sum_{m'=-\ell}^{\ell} e^{im'\alpha} d_{m'm}^{\ell}(-\beta) e^{im\gamma} h_{\ell m'}^{\text{CP}}(f; \boldsymbol{\lambda}). \quad (10)$$

IV. SPIN PARAMETRISATION

Our goal is to model generic non-eccentric black-hole binaries with any physically reasonable values of M , η , $\boldsymbol{\chi}_1$ and $\boldsymbol{\chi}_2$. Given NR waveforms that cover only the single-spin parameter space, we require a mapping between generic two-spin configurations and approximately equivalent configurations where $\boldsymbol{\chi}_2 = 0$. Reductions of the spin degrees of freedom have been used in earlier studies and are common in waveform modelling, e.g., Refs. [10, 38, 60, 61, 76–79]. In this section we summarise the spin parameterisation we have adopted. In Sec. XIE we demonstrate that the resulting model agrees well with a subset of the two-spin precessing-binary NR waveforms that are currently available.

Both our co-precessing-frame model PHENOMDCP and angle model PHENOMANGLES are tuned to the same 40 single-spin NR waveforms described in Sec. II.

In the inspiral region PHENOMD is based on PN expressions and so parameterised by the masses m_1 and m_2 and dimensionless spins χ_1^{\parallel} and χ_2^{\parallel} of the binary. The leading-order PN spin contribution to the phase is $\chi_{\text{PN}} = \chi_{\text{eff}} - \frac{38\eta}{113} (\chi_1^{\parallel} + \chi_2^{\parallel})$ [76, 77, 80], in which the main contribution is the symmetric spin combination [38, 39],

$$\chi_{\text{eff}} = \frac{m_1 \chi_1^{\parallel} + m_2 \chi_2^{\parallel}}{m_1 + m_2}. \quad (11)$$

As such, the NR calibrated merger-ringdown region of PHENOMD is parameterised by the normalised quantity,

$$\hat{\chi} = \left(1 - \frac{76\eta}{113}\right)^{-1} \chi_{\text{PN}}. \quad (12)$$

The final black hole is parameterised by the final mass M_f and spin a_f , which are estimated using independent fits to the NR data [8].

Although PHENOMD is tuned to equal-spin or single-spin NR waveforms, and is often described as a single-spin model, the use of both spins in the underlying inspiral PN phase expressions, and the two different single-spin parameterizations $\hat{\chi}$ and a_f in the merger-ringdown calibration, mean that the model also incorporates some two-spin effects, and indeed has been shown in some cases to describe two-spin configurations to high accuracy [81].

PHENOMDCP is constructed such that PHENOMD is explicitly recovered in the absence of precession. To this end, PHENOMD's phenomenological parameters, which we will generically refer to as λ_k , are modified according to,

$$\lambda'_k = \lambda_k + \chi_{\perp} \nu_k, \quad (13)$$

where ν_k is the new phenomenological parameter to be modelled across the intrinsic parameter space and χ_{\perp} quantifies the in-plane spin component and as such gives a measure of the degree of precession in the system. In Eq. (13) it is manifestly evident that, when $\chi_{\perp} = 0$, PHENOMDCP reduces to PHENOMD. The parameter χ_{\perp} is defined as part of our treatment of the precession angles, which we will now describe.

As with previous precessing-binary PHENOM models, we will also use PN results to describe the precession angles through inspiral. Ref. [12, 82] provide complete two-spin expressions, and as such are parameterised by the masses m_1 and m_2 and the dimensionless spins $\boldsymbol{\chi}_1$ and $\boldsymbol{\chi}_2$ of the binary.

Conversely, for the merger-ringdown we will construct phenomenological expressions for the angles, parameterised according to the parameters of the single-spin NR simulations, $(\eta, \chi, \theta_{\text{LS}})$. Although the NR-calibrated merger-ringdown angle model is a model of single-spin systems, we can estimate the angles for generic two-spin

systems by making an approximate mapping from two-spin systems to our single-spin angle model. Our mapping is defined as follows.

We first map the spin components to the two effective spin parameters used in previous PHENOM models. For the aligned-spin components we use the combination χ_{eff} , as defined in Eq. (11). Although χ_{PN} is the appropriate aligned-spin parameter from PN theory, in precessing systems χ_{eff} is a constant of the PN equations of motion without radiation reaction [83], and can be seen to vary less during inspiral than χ_{PN} .

Following Ref. [61], we also define the effective precession spin, χ_{p} , based on the leading-order PN precession dynamics,

$$\chi_{\text{p}} = \frac{S_{\text{p}}}{m_1^2}, \quad (14)$$

where $S_{\text{p}} = \frac{1}{A_1} \max(A_1 S_1^\perp, A_2 S_2^\perp)$, $A_1 = 2 + 3m_2/(2m_1)$, and $A_2 = 2 + 3m_1/(2m_2)$. χ_{eff} parameterises the spin parallel to the orbital angular momentum while χ_{p} parameterises the spin perpendicular to the orbital angular momentum, i.e., in the plane of the binary.

This definition was motivated by the observation that the vectors \mathbf{S}_1^\perp and \mathbf{S}_2^\perp rotate in the plane at different rates, and over the course of the inspiral the magnitude of their vector sum will oscillate between the sum and difference of their two magnitudes. As shown in Ref. [61], the average value of the in-plane spin contribution to the precession dynamics can be approximated well by χ_{p} for mass ratios $q \gtrsim 1.5$. However, at mass ratios very close to one the spins precess in the plane at approximately the same rate, and so add or cancel in the same way at all times, and χ_{p} does not provide an ideal single-spin mapping. (This is illustrated in more detail in Ref. [79].) Extreme examples are the ‘‘superkick’’ configurations [57], where the black holes are of equal mass, and $\chi_1^\parallel = \chi_2^\parallel = 0$ and $\chi_1^\perp = -\chi_2^\perp$. From the symmetry of the configuration, the two spins rotate at the same rate at all times, and therefore the total in-plane spin is zero, and the system does not precess. For a superkick configuration χ_{p} clearly does not provide the appropriate ‘‘single-spin’’ mapping, which in this case should be to a system with zero in-plane spin.

To deal with such cases, we also introduce χ_{s} , which is constructed from the vector sum of the in-plane spin vectors at a single reference time/frequency of the waveform. In our construction these are the in-plane components of the spin vectors input to the waveform generation. We define χ_{s} as,

$$\chi_{\text{s}} = \frac{|\mathbf{S}_1^\perp + \mathbf{S}_2^\perp|}{m_1^2}. \quad (15)$$

Given a two-spin system defined by \mathbf{S}_1 and \mathbf{S}_2 , we model the precession angles through the merger and ringdown by mapping to a corresponding single spin, which is placed on the larger black hole. This single spin has magnitude χ_{\parallel} in the direction parallel to the orbital angular

momentum and χ_{\perp} in the orbital plane, where,

$$\chi_{\parallel} = \frac{M\chi_{\text{eff}}}{m_1}, \quad (16)$$

$$\chi_{\perp} = \begin{cases} \cos^2(\theta_q)\chi_{\text{s}} + \sin^2(\theta_q)\chi_{\text{p}}, & 1 \leq q \leq 1.5 \\ \chi_{\text{p}}, & q > 1.5, \end{cases} \quad (17)$$

where $\theta_q = (q-1)\pi$. This combination of χ_{s} and χ_{p} given for $1 \leq q \leq 1.5$ is designed to provide a smooth transition between the regimes where χ_{s} and χ_{p} are most appropriate. The cut-off at $q = 1.5$ was motivated in part by the work in Ref. [61] and in part by studying the PN expressions for β while employing various two-spin mappings. From studying β for systems with $q < 1.5$ using both the PN expressions and NR data we observe that the magnitude of β for such systems is consistently small. We therefore conclude that the precession effects are weak for systems with $q < 1.5$, and so, since we mostly wish to capture the correct qualitative behaviour in this regime, we expect that different choices for χ_{s} , or for the transition to χ_{p} , would have an impact on GW measurements smaller than the other approximations used in our model. The mapping we employ here crucially captures the correct behaviour in the equal mass limit. (Alternative choices of single-spin mapping are suggested in Refs. [78, 79]; since we use a single-spin mapping only to connect our single-spin merger-ringdown model to a generic-spin inspiral model, we expect that there are many reasonable choices of mapping that would work equivalently well.) This expression for χ_{\perp} is also used to parameterise the in-plane spin effects in the coprecessing model, as described in Eq. (13).

The total spin magnitude χ and the angle between the orbital and spin angular momenta are given by

$$\chi = \sqrt{\chi_{\parallel}^2 + \chi_{\perp}^2}, \quad (18)$$

$$\cos\theta_{\text{LS}} = \frac{\chi_{\parallel}}{\chi}. \quad (19)$$

These reduce to the correct values for the cases to which we tuned the model and also correctly re-weight two-spin cases and cases where the spin is predominantly on the smaller black hole.

In the \mathbf{J} -aligned frame, in which we have constructed our model, the spin placed on the larger black hole has the components

$$\mathbf{S}' = \begin{pmatrix} \cos\alpha(\chi_{\perp}\cos\beta + \chi_{\parallel}\sin\beta) \\ \sin\alpha(\chi_{\perp}\cos\beta + \chi_{\parallel}\sin\beta) \\ -\chi_{\perp}\sin\beta + \chi_{\parallel}\cos\beta \end{pmatrix} \quad (20)$$

where α and β are the values of the precession angles introduced in Sec. III, here evaluated at the reference frequency.

V. CO-PRECESSING-FRAME MODEL

A key assumption of most precessing signal models has been that the coprecessing multipole moments are

largely devoid of precession related effects [10, 16, 25, 32]. This assumption is motivated by the PN description of inspiral, where in-plane spin components do not impact the coprecessing waveforms' phase, and so can be disregarded [31, 84]. In this sense, most precessing signal models have used un-modified non-precessing inspiral waveforms in the coprecessing frame. Because the PN motivation is only well suited for inspiral, for the waveforms' immediate pre-merger and merger, additional assumptions must be made [27, 33, 47]. For example, all previous precessing-binary PHENOM models use an estimate of the precessing system's final mass and spin to compute the remnant BH's Quasinormal Mode (QNM) frequencies. In turn, these QNM frequencies allow the frequency-domain waveforms' features at merger to be shifted such that they occur near physically appropriate values. In Sec. IA we illustrated deviations from the simplifying assumptions made in both the inspiral and merger-ringdown, and in this section we refine those assumptions by constructing a tuned coprecessing waveform model.

We introduce PHENOMDCP, a model for the $\ell = |m| = 2$ coprecessing gravitational wave multipole moment tuned to NR. PHENOMDCP is tuned to the 40 late inspiral, merger and ringdown NR simulations discussed in Sec. II. By construction, PHENOMDCP reduces to PHENOMD for non-precessing BBH systems. We could have instead adapted the more recent PHENOMXAS model [14], which is tuned also to two-spin systems, but since two-spin effects are unlikely to be measurable in most observations [13, 40], and we have tuned to NR results only from single-spin precessing systems, we will leave two-spin extensions of the co-precessing-frame model to future work.

We consider PHENOMDCP to be a first step towards a high accuracy coprecessing waveform model. Here we briefly review the structure of PHENOMD, and how this structure is extended by PHENOMDCP. Physical features of the NR waveforms and PHENOMDCP are provided and discussed in detail in Sec. IX. Plots showing fits of model parameters across the space of initial binary masses and spins are provided in Appendix C.

A. Briefly on the structure of PHENOMD

PHENOMD [8, 9] is a phenomenological model for the $\ell = |m| = 2$ frequency-domain multipole moments of gravitational waves from non-precessing BBHs. The morphology of each multipole moment is organized into three regimes: (1) inspiral, where PN theory applies, (2) intermediate, where the time domain evolution of the black holes is near merger, and (3) merger-ringdown, where the time domain evolution corresponds to the final coalescence and formation of a stationary remnant BH. PHENOMD models each of these regimes with different ansätze. The coefficients of each PHENOMD ansatz are functions of the initial binary's masses and aligned spins.

In PHENOMDCP these coefficients are modified to depend on information about the in-plane spins.

PHENOMD was calibrated to 19 NR waveforms between $q = 1$ and $q = 18$. For unequal-mass systems, PHENOMD is calibrated to $\chi_{\text{eff}} \in [-0.85, 0.85]$, and for equal-mass systems is it calibrated to $\chi_{\text{eff}} \in [-0.98, 0.98]$. In each NR simulation the black-hole spins were either equal, $\chi_1 = \chi_2$, or the smaller black hole was non-spinning. The calibration waveforms were hybrids of SEOBNRV2 (without NR tuning) and NR waveforms. Over the model's calibration region, its typical deviations (mismatches) from NR are less than 1% [9].

B. Construction of PHENOMDCP

In the PHENOMP models [10, 12, 13] PHENOMD is used as an approximate co-precessing-frame model, with the ringdown frequency modified according to an estimate of the final black hole's spin. In PHENOMDCP we instead use NR waveforms to tune in-plane-spin deviations to a subset of the model coefficients. Here we briefly overview the modifications of PHENOMD that result in PHENOMDCP.

As in previous models, PHENOMDCP assumes that in the coprecessing frame only the $(\ell, m) = (2, \pm 2)$ multipole moments are needed, and that the $m = 2$ and $m = -2$ strain moments are related by conjugation (Sec. III). Under these assumptions we only need model the amplitude and phase of h_{22}^{CP} ,

$$h_{22}^{\text{CP}}(f; \boldsymbol{\lambda}) = A(f; \boldsymbol{\lambda}) e^{-i\phi(f; \boldsymbol{\lambda})}. \quad (21)$$

In Eq. (21), $A(f; \boldsymbol{\lambda})$ is the frequency domain amplitude of h_{22}^{CP} , $\phi(f; \boldsymbol{\lambda})$ is its phase, $f = \omega/2\pi$ references a frequency bin in geometric units, and $\boldsymbol{\lambda}$ encapsulates the system's initial parameters (Sec. IV),

$$\boldsymbol{\lambda} \in (\eta, \chi, \theta_{\text{LS}}), \quad (22)$$

where, as described in Sec. IV, the total spin χ consists of the aligned-spin component χ_{eff} and the in-plane component χ_{\perp} , and for our single-spin calibration waveforms, $\chi_{\perp} = \chi_{\text{p}} = \chi_{\text{1}}^{\perp}$.

Given the system's initial parameters $\boldsymbol{\lambda}$, PHENOMDCP is defined by a series of polynomials between $\boldsymbol{\lambda}$ and phenomenological model parameters. PHENOMDCP's model parameters are based directly on those of PHENOMD (Eq. 13). Specifically, PHENOMDCP uses the PHENOMD amplitude and phase ansatz with model parameters offset by a term proportional to χ_{\perp} . Thus, when $\chi_{\perp} = 0$, PHENOMDCP reduces to PHENOMD.

Precession effects are known to be most relevant in the late inspiral and merger-ringdown [10, 13]. Thus PHENOMDCP is made to be equivalent to PHENOMD in the early inspiral. Modified versions of PHENOMD are used for the waveforms' late-inspiral phase, merger-ringdown phase, and merger-ringdown amplitude:

$$\phi_{\text{Int}} = \frac{1}{\eta} \left(\beta_0 + \beta_1 f + \beta_2' \ln(f) - \frac{\beta_3}{3} f^{-3} \right), \quad (23)$$

$$\phi_{\text{MR}} = \frac{1}{\eta} \left\{ \alpha_0 + \alpha_1 f - \alpha_2 f^{-1} + \frac{4}{3} \alpha_3 f^{3/4} + \alpha'_4 \tan^{-1} \left(\frac{f - \alpha_5 f_0^{(\phi)}}{f_1^{(\phi)}} \right) \right\}, \quad (24)$$

$$\frac{A_{\text{MR}}}{A_0} = \gamma_1 \frac{\gamma_3 f_1^{(A)}}{(f - f_0^{(A)})^2 + (\gamma_3 f_1^{(A)})^2} e^{-\frac{\gamma'_2 (f - f_0^{(A)})}{\gamma_3 f_1^{(A)}}}. \quad (25)$$

In Eqs. (23)-(25) Greek symbols denote model parameters defined in Ref. [9], and of those, primed symbols, such as α'_4 , denote parameters modified for PHENOMDCP. Please note that these Greek symbols should not be confused with the Euler angles that define the coprecessing frame. In Eq. (24), $f_0^{(\phi)}$ is an “effective ringdown frequency” that is particular to the phase. Similarly, $f_1^{(\phi)}$ corresponds to the ringdown decay rate. In the setting of PHENOMD, $f_0^{(\phi)}$ and $f_1^{(\phi)}$ are simply referred to as f_{RD} and f_{damp} . In Eq. (25), $f_0^{(A)}$ is an effective ringdown frequency particular to the amplitude, and $f_1^{(A)}$ is equivalent to the ringdown decay rate used in PHENOMD,

$$f_1^{(A)} = f_{\text{damp}}. \quad (26)$$

Our notation for the effective ringdown frequencies signals that we will not assume a direct relationship between the ringdown frequencies predicted by BH perturbation theory, and those relevant for coprecessing waveforms. This point is discussed further in Sec. IX.

In constructing PHENOMDCP it was found that only a subset of PHENOMD’s parameters needed to be modified. These parameters are those needed to address the disconnect between PHENOMD and the coprecessing frame NR data discussed in Sec. II. The modified parameters correspond to the late inspiral behavior of the frequency domain phase,

$$\beta'_2 = \beta_2 + \chi_{\perp} \zeta_2, \quad (27)$$

the merger-ringdown phase,

$$\alpha'_4 = \alpha_4 + \chi_{\perp} \nu_4 \quad (28)$$

$$f_0^{(\phi)} = f_0 + \chi_{\perp} \nu_5 \quad (29)$$

$$f_1^{(\phi)} = f_1 + \chi_{\perp} \nu_6, \quad (30)$$

and the merger-ringdown amplitude,

$$\gamma'_2 = \gamma_2 + \chi_{\perp} \mu_2 \quad (31)$$

$$f_0^{(A)} = f_0 + \chi_{\perp} \mu_4. \quad (32)$$

In Eqs. (23)-(26), all parameters not defined in Eqs. (27)-(32) are defined in Ref. [13]. Similarly, in Eqs. (27)-(32), $\{\alpha_4, f_0, f_1, \gamma_2\}$ are defined in Ref. [13].

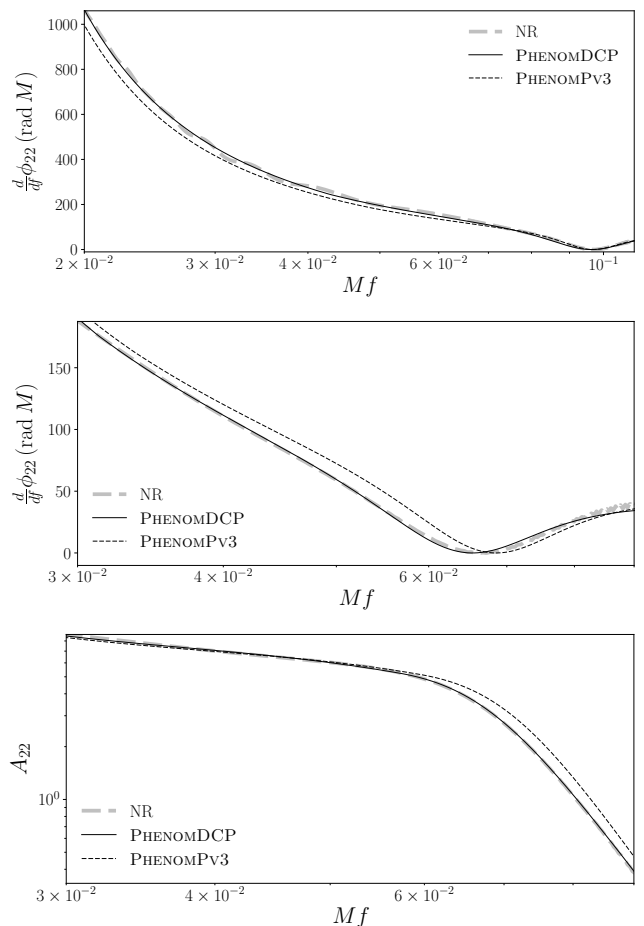


FIG. 4. Frequency domain comparison of NR and model waveforms in the co-precessing frame. (top) phase derivative for the $(q, \chi, \theta_{\text{LS}}) = (4, 0.8, 60^\circ)$ configuration, which illustrates the variation in the inspiral phase. (middle and bottom) phase derivative and amplitude for the $(q, \chi, \theta_{\text{LS}}) = (4, 0.8, 150^\circ)$ configuration, which demonstrates the shift in effective ringdown frequency.

The calibration of PHENOMDCP has been performed by fitting Eqs. (23)-(26) to each NR waveform in our calibration set. This yields a collection of calibration points for each model parameter. For each of PHENOMDCP’s model parameters, these points were modeled as polynomials in λ using `gmvpfit`, which uses multidimensional least-squares regression driven by a greedy algorithm [85, 86].

Figures 25-26 show the behavior of the PHENOMDCP model parameters as functions of symmetric mass-ratio and θ_{LS} over the calibration space. The parameter surfaces shown in Figs. 25-26 correspond to percent root-mean-square errors of 3.42% in amplitude and 2.53% in phase.

Figure 4 compares evaluations of PHENOMDCP to NR and PHENOMPV3 for the cases discussed in Sec. IA. The top panel of Fig. 4 highlights the effect of modifying the phase. The middle and bottom panels highlight the effect

of modifying the effective ringdown frequency and damping times. We see that PHENOMDCP successfully corrects for the discrepancies in the modified-PHENOMD co-processing-frame model used in PHENOMPv3; see Sec. XI for quantitative accuracy results.

VI. PRECESSION ANGLE MODEL: INSPIRAL

Our model of the precession angles consists of two parts. The first describes the precession during inspiral, and is based on the MSA angles presented in Ref. [54], and used in previous PHENOM models [12, 13, 16]. The second part is a phenomenological model of the precession angles during merger and ringdown, tuned to the NR waveforms presented in Sec. II. We discuss the inspiral angles in this section, the merger-ringdown angles in Sec. VII, and the combined inspiral-merger-ringdown (IMR) angle model in Sec. VIII.

A. MSA angles

The precession angles in the inspiral regime are calculated using PN theory. In Ref. [54, 82] the authors derived a closed-form analytic approximation to the inspiral precession dynamics. To achieve this GW driven radiation-reaction was introduced into an analytic solution to the conservative precession dynamics [87] by exploiting the hierarchy of timescales in the binary inspiral problem using a mathematical technique called multiple scale analysis [88, 89]. The hierarchy of timescales are $t_{\text{orb}} \ll t_{\text{prec}} \ll t_{\text{rr}}$, where t_{orb} , t_{prec} and t_{rr} are the orbital, precession and radiation-reaction timescales respectively. This model is a function of all 6 spin components (two 3-vectors for each BH) and incorporates spin-orbit and spin-spin effects to leading order in the conservative dynamics and up to 3.5PN order in the dissipative dynamics, ignoring spin-spin terms. The MSA angles are shown for an example configuration in Fig. 2. We can see that the agreement is poor for all three angles at high frequencies, which correspond to the merger and ringdown. At lower frequencies, the PN and NR values for α and γ agree well, but for β do not. As noted earlier, this is because the PN β describes the inclination of the orbital plane with respect to $\hat{\mathbf{J}}$, which differs from the inclination of the QA direction.

In the next section we apply higher-order PN information to improve the PN estimate of β .

B. Higher-order PN corrections to β

As discussed in Sec. III, in the quadrupole approximation the maximum GW signal power is emitted perpendicular to the orbital plane, and therefore the angles that describe the precession dynamics of the orbital plane are the same as those associated with the QA frame of the

GW signal [32, 49, 50]; this motivated the original QA procedure presented in Ref. [32]. For the full signal, this identification is only approximate [32, 51–53], and we expect the approximation to be less accurate at higher frequencies. Our modelling approach is based on applying a frequency-dependent rotation to a model of the waveform in the co-processing QA frame, and as such the rotation angles should be those associated with the *signal*. However, all current models [10, 12, 21, 22] use the angles associated with the *dynamics*.

As we saw in Fig. 2, the MSA dynamics α and γ provide a good approximation to the corresponding NR signal angles at low frequencies, but the MSA β does not. Fortunately, we have access to PN signal amplitudes beyond the quadrupole approximation, and can use these to calculate a more accurate estimate of the signal β . One way to do this would be to calculate a full PN waveform, e.g., from the model in Ref. [82], and apply the quadrupole-alignment procedure to calculate β . However, this will be much more computationally expensive than the current MSA approximant, and it is possible to obtain a sufficiently accurate result with a simpler approach.

In this calculation we will refer to the opening angle of the orbital plane with respect to \mathbf{J} as ι , and continue to denote the opening angle of the QA frame by β . These angles are therefore given by

$$\cos \iota = \hat{\mathbf{L}} \cdot \hat{\mathbf{J}}, \quad (33)$$

$$\cos \beta = \hat{\mathbf{V}} \cdot \hat{\mathbf{J}}, \quad (34)$$

where $\hat{\mathbf{L}}$ is the orbital angular momentum and $\hat{\mathbf{V}}$ is the direction of maximum emission.

To illustrate our approach, consider the rotation from a co-processing signal that contains only the ($\ell = 2, |m| = 2$) multipoles, $h_{2,\pm 2}^{\text{NP}}$, to produce a precessing-binary signal in the inertial frame. We begin in the quadrupole approximation, where the inertial frame is identified with the precession of the orbital plane, and so we use the opening angle ι . We will focus on only the resulting (2, 2) and (2, 1) multipoles in the inertial frame, and only the angles ι, α (since the additional phase rotation γ will not affect our argument). The precessing-binary signal in the inertial frame, h^{P} , is now,

$$h_{2,2}^{\text{P}} = e^{-2i\alpha} \left(\cos^4 \left(\frac{\iota}{2} \right) h_{2,2}^{\text{NP}} + \sin^4 \left(\frac{\iota}{2} \right) h_{2,-2}^{\text{NP}} \right), \quad (35)$$

$$h_{2,1}^{\text{P}} = -2e^{-i\alpha} \left(\cos^3 \left(\frac{\iota}{2} \right) \sin \left(\frac{\iota}{2} \right) h_{2,2}^{\text{NP}} - \cos \left(\frac{\iota}{2} \right) \sin^3 \left(\frac{\iota}{2} \right) h_{2,-2}^{\text{NP}} \right). \quad (36)$$

The non-precessing multipoles can be written as,

$$h_{2,\pm 2}^{\text{NP}} = A e^{\mp 2i\Phi}, \quad (37)$$

where A and Φ are the time/frequency-dependent amplitude and orbital phase. When ι is small, $h_{2,2}^{\text{NP}}$ makes the strongest contribution to the precessing-waveform multipoles, and we see that ι determines the relative amplitude

of $h_{2,2}^P$ and $h_{2,1}^P$. We can isolate the $e^{-2i\Phi}$ term as follows,

$$\bar{h}_{2,2}^P = \frac{1}{2\pi} \int_0^{2\pi} h_{2,2}^P e^{2i\Phi} d\Phi \quad (38)$$

$$= A e^{-2i\alpha} \cos^4\left(\frac{\iota}{2}\right), \quad (39)$$

$$\bar{h}_{2,1}^P = -2A e^{-i\alpha} \cos^3\left(\frac{\iota}{2}\right) \sin\left(\frac{\iota}{2}\right). \quad (40)$$

From these we can readily calculate that the inclination ι is

$$\iota = 2 \tan^{-1} \left(\frac{|\bar{h}_{2,1}^P|}{2|\bar{h}_{2,2}^P|} \right). \quad (41)$$

At leading (quadrupole) order, ι is the precession angle β .

If we now use higher-order PN amplitude expressions [84], then the angle β that identifies the frame in which the ($\ell = 2, |m| = 2$) multipoles are maximised will not necessarily be the same as the inclination angle ι , but the expression above *will* still give us an estimate of the orbit-averaged β . Note that the MSA angles in Ref. [82] are also orbit-averaged (i.e., nutation effects are absent), so this is a consistent treatment.

The multipole expressions in Ref. [84] are given in terms of the orbital phase Φ , the precession angles α and ι , and the spin components. For the spin components, we make an approximate reduction to our single-spin systems as follows. The inclination of the spin from the z -axis is the spin's inclination from the orbital angular momentum vector, θ_{LS} , minus the inclination of the orbital angular momentum from the z -axis, ι . The azimuthal angle of the spin vector is $(\alpha + \pi)$, because, since $\mathbf{L} = \mathbf{J} - \mathbf{S}$, the x - y -plane components of \mathbf{L} and \mathbf{S} will be in opposite directions, and so their azimuthal angles will differ by π . The final result, for a given configuration, depends only on the dynamics inclination ι as a function of frequency; we use the MSA expression for $\iota(f)$.

In Ref. [84] the amplitudes are expanded in powers of $v = (\pi f)^{1/3}$. We define $\delta = m_1 - m_2$, where $m_1 > m_2$, and so $\delta > 0$; $\eta = m_1 m_2 / (m_1 + m_2)^2$, $\chi_s = (\chi_1 + \chi_2)/2$, $\chi_a = (\chi_1 - \chi_2)/2$, and so,

$$\begin{aligned} \chi_{s/a,x} &= \chi \sin(\theta_{LS} - \iota) \cos(\alpha + \pi)/2, \\ \chi_{s/a,y} &= \chi \sin(\theta_{LS} - \iota) \sin(\alpha + \pi)/2, \\ \chi_{s/a,z} &= \chi \cos(\theta_{LS} - \iota)/2. \end{aligned} \quad (42)$$

If we substitute these into the PN multipole expressions for $h_{2,2}^P$ and $h_{2,1}^P$, and then apply Eq. (41), we obtain the relatively simple expression,

$$\beta = 2 \tan^{-1} \left(\frac{\sec(\iota/2) (c_0 + c_2 v^2 + c_3 v^3)}{d_0 + d_2 v^2 + d_3 v^3} \right), \quad (43)$$




FIG. 5. Opening angles for the $(q, \chi, \theta_{LS}) = (8, 0.8, 60^\circ)$ configuration. Solid black: the NR opening angle of the QA frame, β . Dotted blue: the PN opening angle of the orbital plane, ι . Dashed magenta: Approximate QA angle β as a function of ι ; see text for details.

where

$$\begin{aligned} c_0 &= 84 \sin \iota, \\ c_2 &= (110\eta - 214) \sin \iota, \\ c_3 &= -7(6 + 6\delta + 5\eta)(2 \cos \iota - 1)\chi \sin \theta_{LS}, \\ &\quad + 56(3\pi - (1 + \delta - \eta)\chi \cos \theta_{LS}) \sin \iota, \\ d_0 &= 84 \cos\left(\frac{\iota}{2}\right), \\ d_2 &= (110\eta - 214) \cos\left(\frac{\iota}{2}\right), \\ d_3 &= 14(6 + 6\delta + 5\eta)\chi \sin \theta_{LS} \sin\left(\frac{\iota}{2}\right) \\ &\quad + 56 \cos\left(\frac{\iota}{2}\right) (3\pi - (1 + \delta - \eta)\chi \cos \theta_{LS}). \end{aligned} \quad (44)$$

Fig. 5 also shows the modified $\beta(\iota)$ for the $(q, \chi, \theta_{LS}) = (8, 0.8, 60^\circ)$ configuration. We see the PN inspiral $\beta(\iota)$ now shows much better agreement with the NR result at low frequencies. We find similar results across the parameter space that we have considered, and therefore to calculate β in our model, we use Eq. (43) in conjunction with the MSA ι as calculated in Refs. [12, 82], to construct β through the inspiral. The features of the NR (α, β, γ) at higher frequencies, which are not captured at all by the PN expressions, will be explicitly modelled in Sec. VII.

C. Two-spin β

The MSA ι for a two spin system shows oscillations that become unphysically large through late inspiral and towards merger and which are not seen in the precession angles calculated for two-spin NR systems, as can be seen in Fig. 6. These oscillations also complicate connecting the inspiral expression to the single-spin-tuned merger-ringdown ansatz. We therefore taper these oscillations to recover the value and gradient of β for an




FIG. 6. Various options for the PN expression for the opening angle. The left-hand panel shows the PN value of ι for a two-spin system (blue) and for the equivalent single-spin system (green) calculated using the expressions used in PHENOMPv3. In light blue is shown the effect of tapering the two-spin oscillations to the single-spin value at the connection frequency f_c , shown as a grey vertical line. In the right-hand panel the value for β used in the model (pink) is compared with the NR value of β found for this case. We only show ι and β up to f_c , since the merger-ringdown model is used at higher frequencies. The configuration shown is SXS1397 in Table II.

equivalent single-spin system at the point at which we wish to connect the inspiral and merger-ringdown parts of the model.

For a system described by two spins \mathbf{S}_1 and \mathbf{S}_2 (with total spin $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$) we use the mapping to the appropriate single spin system defined in Sec. IV: \mathbf{S}'_1 is given by Eq. (20) and $\mathbf{S}'_2 = (0, 0, 0)$. We evaluate the PhenomPv3 expression for ι for both of these configurations and identify the oscillations introduced by the two-spin effects as,

$$\iota_{\text{osc}} = \iota(\mathbf{S}_1, \mathbf{S}_2) - \iota(\mathbf{S}'_1, \mathbf{S}'_2). \quad (46)$$

We then apply a taper to these oscillations that ensures ι will tend to the single spin value and gradient at a given frequency f_c and add the oscillations back to the single-spin function. The final two-spin expression for ι is then given by

$$\iota = \begin{cases} \iota(\mathbf{S}'_1, \mathbf{S}'_2) + \cos^2\left(\frac{2\pi f}{4f_c}\right) \times \iota_{\text{osc}} & f \leq f_c \\ \iota(\mathbf{S}'_1, \mathbf{S}'_2) & f > f_c, \end{cases} \quad (47)$$

where f_c is the frequency at which the inspiral expression for β is connected to the merger-ringdown expression defined below in Eq. (59).

Given an estimate for the dynamics ι , we now wish to rescale it to produce an estimate for the signal β , as described in Sec. VI B. To do this we also need an estimate of the frequency-dependent in-plane spin component, and therefore χ and θ_{LS} , as required in Eqs. (42). We assume that the component of the spin \mathbf{S} parallel to the orbital angular momentum, S_{\parallel} , remains fixed. We further approximate that the frequency dependence of the magnitude of \mathbf{J} is dominated by changes to the magnitude of

\mathbf{L} ,

$$J(f) = J_0 + L(f) - L_0, \quad (48)$$

where the magnitude L is given by the 3PN expression for the orbital angular momentum used by PHENOMPv3 to calculate ι and the 0-subscript denotes quantities specified at the reference frequency. As such, we may write the frequency-dependent magnitude of the in-plane component S_{\perp} of spin \mathbf{S} as

$$S_{\perp}(f) = J(f) \sin \iota \quad (49)$$

Substituting this expression for S_p in Eq. (14) we get a value for χ_p . The quantities χ and $\cos \theta_{\text{LS}}$ are then calculated as described in Eqs. (11)–(19) and these values are used to rescale ι to produce β , according to Eq. (43).

The effect of this treatment can be seen in Fig. 6, which shows β for SXS1397 (the intrinsic properties of which are given in Tab. II). The PN expression for the angle captures the oscillations seen at low frequency very well. However, these oscillations do not continue to high frequency and are greatly over-estimated by the full two-spin PN expression. Tapering the oscillations to the single spin value at the connection frequency resolves this issue well. For $f > f_c$ the PN expression is replaced by the merger-ringdown expression described in the following section, so the behaviour of the PN angles here are not an issue. In the rare event where the merger-ringdown contributions are not attached (see Sec. VIII D), only the effective single-spin beta is used beyond $f > f_c$.

VII. PRECESSION ANGLE MODEL: MERGER-RINGDOWN

The PN expressions for the precession angles cannot be reliably extended through merger and ringdown and when compared with the NR angles do not capture the features present at high frequency, as was clear in Fig. 2. We therefore present a phenomenological description of the precession angles α and β in the merger-ringdown regime; the remaining angle γ can then be calculated via Eq. (A7). We describe the functional form of the angles and produce a global fit for each of the co-efficients of the ansatz. This provides a frequency domain description of the precession angles across the parameter space.

A. Functional forms of α and β

The morphology of the merger-ringdown part of α is qualitatively very similar to that of the phase derivative, seen in Ref. [8, 9]. α shows a $1/f$ fall-off with a Lorentzian dip centred around what is approximately the ringdown frequency of the BBH system. This prompts the ansatz,

$$\alpha(f) - \langle \alpha(f) \rangle = \frac{A_1}{f} + \frac{A_2 \sqrt{A_3}}{A_3 + (f - A_4)^2}, \quad (50)$$

where $\langle \alpha(f) \rangle$ is the mean value of α calculated over the frequency range of the fitting region defined below and A_1 , A_2 , A_3 and A_4 are free co-efficients.

The fitting region is based around the Lorentzian dip; it is defined to be the range $f_{\text{dip}} - 0.0225 \leq f \leq f_{\text{dip}} + 0.0075$, where f_{dip} is the frequency at which α reaches its minimum, and recall that we have chosen $M = 1$. The global fit for α within this fitting region has a root mean square error of 4.80×10^{-5} , averaged across the 40 waveforms. Some example comparisons of the result of these fits with the NR value for α are shown in Fig. 7.

During merger and ringdown, β drops rapidly as the dominant emission direction relaxes to its final direction, as discussed in more detail in Sec. IX. The ansatz used to describe β is therefore chosen to grow at low frequencies (as seen in the PN expressions), turnover at the correct frequency, capture the drop and finally tend asymptotically towards the constant value to which the dominant emission direction relaxes. The ansatz we chose to describe this behaviour is,

$$\beta(f) - \langle \beta(f) \rangle = \frac{B_1 + B_2 f + B_3 f^2}{1 + B_4 (f + B_5)^2}, \quad (51)$$

where $\langle \beta(f) \rangle$ is the mean value of β calculated over the frequency range of the fitting region defined below and B_1 , B_2 , B_3 , B_4 and B_5 are free co-efficients.

The fitting region for β is centred around the inflection point in the turnover f_{inf} ; $f \in f_{\text{inf}} \pm 0.03$. The global fit for β within this fitting region has a root mean square error of 7.47×10^{-6} , averaged across the 40 waveforms.

Some example comparisons of the result of these fits with the NR value for β are shown in Fig. 8.

It should be noted that a key feature of the above ansatz is that it does *not* fall to zero after merger. This feature can be seen in both the time and frequency domain values of β , as shown in Figs. 8 and 12. We discuss this in more detail in Sec. IX.

B. The phenomenological co-efficients

The two ansätze given above, which describe the merger-ringdown behaviour of α , Eq. (50), and β , Eq. (51), have 10 free co-efficients between them. Each of these co-efficients was fit across the three-dimensional parameter space described by the symmetric mass ratio, η , the dimensionless spin magnitude, χ , and the cosine of the angle between the orbital angular momentum and the spin angular momentum, $\cos \theta_{\text{LS}}$.

The optimum value of each of the co-efficients for each waveform in the calibration set was found by fitting the relevant ansatz to the NR data using the non-linear least-squares fitting function `curve_fit` from the python package `Scipy` [90]. This function uses the Levenberg-Marquardt algorithm to perform the least-squares fitting. We then performed a three-dimensional fit of each of the co-efficients using the fitting algorithm `mvpolyfit` [85, 86]. This gives each of the co-efficients as a polynomial expansion in η , χ , $\cos \theta_{\text{LS}}$. We specify the terms that appear in the expansion and the algorithm finds the co-efficients of these terms that optimise the fit as well as a measure of how good the fit is. Since we have 40 calibration waveforms, the maximum possible number of terms that can appear in these expressions is 39 in order to avoid over fitting. The fits are restricted so that the highest order term in each dimension is one less than the total number of data points in that dimension. Since the value of each of the co-efficients in the ansatz is to some extent dependent on the value of each of the other co-efficients, we found a global fit for each co-efficient in turn, re-fitting the ansatz to the data while keeping fixed the co-efficients that had already been fit. We first fitted the co-efficients that varied most smoothly across the parameter space and those for which the general behaviour across the parameter space was already understood. For α this meant we first fitted the location of the dip, A_4 , followed by the other co-efficients in the order A_1 , A_2 and A_3 . For β we fitted the value of $\langle \beta(f) \rangle$ separately as this had a clear parameter space trend. We then fitted the co-efficients in the order B_1 , B_2 , B_3 , B_5 and B_4 since the co-efficients in the numerator were generally better behaved than those in the denominator.

The general expression for each co-efficient is

$$\Lambda^i = \sum_{p=0}^3 \sum_{q=0}^1 \sum_{r=0}^4 \lambda_{pqr}^i \eta^p \chi^q \cos^r \theta_{\text{LS}}, \quad (52)$$

where $\Lambda \in [A, B]$ are the co-efficients in the ansatz de-




FIG. 7. Comparison of the phenomenological ansatz presented in Eq. (50) (solid lines) with the NR data (translucent lines) over the frequency range to which the co-efficients in the ansatz were tuned for a selection cases in the NR catalogue with $\theta_{\text{LS}} = 90^\circ$ at varying mass ratios. We have made use of the freedom to choose a constant offset in α in order to offset the curves shown here to make them easier to distinguish.




FIG. 8. Comparison of the phenomenological ansatz presented in Eq. (51) (solid lines) with the NR data (translucent lines) over the frequency range to which the co-efficients in the ansatz were tuned for a selection of cases in the NR catalogue with $\theta_{\text{LS}} = 90^\circ$ at varying mass ratios.

cribing α and β respectively and $i \in [1, 2, 3, 4]$ and $[0, 1, 2, 3, 4, 5]$ respectively. The λ_{pqr}^i give the co-efficients of the polynomial expansion of the multi-dimensional fits of Λ_i . This expression has a maximum of 40 terms. Not all of these terms are used in the expressions for each of the co-efficients; the co-efficient with the fewest number of terms has only 25 while that with the greatest number of terms contains 39.

The co-efficients for α and β vary smoothly across the parameter space, as can be seen in Figs. 27 and 28 in Appendix C respectively. The residual plots above the fit surfaces show that the global fits agree closely with the values of the co-efficients found from fitting the ansatz to each individual simulation.

VIII. FULL INSPIRAL-MERGER-RINGDOWN ANGLE MODEL

The expressions for the precession angles for the two distinct inspiral and merger-ringdown regions are connected so that the connection is smooth and the full IMR expression for the angles agrees with the NR data over the entirety of the region for which it is available. The method used to connect the two regions was different for each angle.

A. Connection method for α

For α , the regions are connected using an interpolating function of the form

$$\alpha_{\text{interp}}(f) = a_0 f^2 + a_1 f + a_2 + \frac{a_3}{f}, \quad (53)$$

defined over the frequency range $[f_1, f_2]$. This range was chosen to be as small as possible. The lower frequency limit was chosen to be the highest frequency for which the inspiral expressions agreed with the NR data while the upper frequency limit was chosen to be the lower limit for which the fitted merger-ringdown expressions still agreed well with the NR data. Since the MSA PN expressions for the angles agree well with the NR data over most of the waveform, there is a wide range of frequency values over which the interpolation could be performed. We choose the frequency range to be defined in terms of the location of the Lorentzian dip, A_4 : $f_1 = 2A_4/7$ and $f_2 = A_4/3$.

The co-efficients of Eq. (53) are chosen so that

1. $\alpha_{\text{interp}}(f_1) = \alpha_{\text{PN}}(f_1)$ and $\alpha_{\text{interp}}(f_2) = \alpha_{\text{MR}}(f_2)$, since there is freedom in an overall constant offset in α ,
2. $\alpha'_{\text{interp}}(f_1) = \alpha'_{\text{PN}}(f_1)$ and $\alpha'_{\text{interp}}(f_2) = \alpha'_{\text{MR}}(f_2)$ in order to ensure the two parts are connected continuously.

α_{PN} is the MSA PN expression used for α in the inspiral regime. α_{MR} is the merger-ringdown ansatz given in Eq. (51). The co-efficients are given by

$$\begin{aligned} a_0 &= \frac{1}{D} [2(f_1\alpha_1 - f_2\alpha_2) - (f_1 - f_2)((f_1\alpha'_1 + f_2\alpha'_2) + (\alpha_1 - \alpha_2))], \\ a_1 &= \frac{1}{D} [3(f_1 + f_2)(f_1\alpha_2 - f_2\alpha_1) + (f_1 - f_2)((f_1 + 2f_2)(f_1\alpha'_1 + \alpha_1) + (2f_1 + f_2)(f_2\alpha'_2 + \alpha_2))], \\ a_2 &= \frac{1}{D} [6f_1f_2(f_1\alpha_1 - f_2\alpha_2) + (f_1 - f_2)(f_2(2f_1 + f_2)(f_1\alpha'_1 + \alpha_1) + f_1(f_1 + 2f_2)(f_2\alpha'_2 + \alpha_2))], \\ a_3 &= \frac{1}{D} [f_1f_2^2(f_2 - 3f_1)\alpha_1 - f_1^2f_2(f_1 - 3f_2)\alpha_2 + f_1f_2(f_1 - f_2)(f_2(f_1\alpha'_1 + \alpha_1) + f_1(f_2\alpha'_2 + \alpha_2))], \end{aligned} \quad (54)$$

where α_i and α'_i , $i = 1, 2$, are the value of α and its derivative at the limits of the frequency range and $D = (f_2 - f_1)^3$.

B. Connection method for β

For β , the agreement between the PN expression and the NR data is insufficient to employ the interpolation method described above. Even including the higher order amplitude corrections described in Sec. VIB, the starting frequency of the NR simulations is not low enough

in order to cover the region in which the PN expression closely matches the data for all cases. Instead, we employ a rescaling function that leaves the PN expression invariant at low frequencies but ensures it smoothly connects with the merger-ringdown value of β at the connection frequency f_c . This rescaling function is given by

$$k(f) = 1 + b_1 f + b_2 f^2, \quad (55)$$

which tends to one at low frequencies thus leaving the PN expression unchanged. In order to ensure the value of β and its derivative match at the connection frequency, the

co-efficients b_1 and b_2 are given by

$$b_1 = -\frac{1}{\beta_1^2 f_c} [-2\beta_1 (\beta_2 - \beta_1) + (\beta_1 \beta_2' - \beta_2 \beta_1') f_c], \quad (56)$$

$$b_2 = -\frac{1}{(\beta_1 f_c)^2} [\beta_1 (\beta_2 - \beta_1) - (\beta_1 \beta_2' - \beta_2 \beta_1') f_c], \quad (57)$$

where the β_i and β_i' are the value of β and its derivative evaluated at the connection frequency. The subscript 1 indicates that this is the value of β given by the original PN expressions while 2 indicates the values from the merger-ringdown expression.

The definition of the connection frequency depends on the morphology of the merger-ringdown ansatz for β for a particular case. As can be seen in Fig. 8, in some parts of the parameter space β rises gently until just before merger then turns over and drops rapidly. However, in other parts of the parameter space this turnover is much more gradual and begins at much lower frequencies. In general, we note that the turnover shifts to lower frequencies as the mass ratio increases and that this trend is most marked for cases with larger θ_{LS} and larger spin magnitude. Our ansatz for β captures both of these morphologies well. In cases where the turnover occurs within the fitting region, we define the connection frequency f_c as the frequency at which the merger-ringdown part has a particular gradient $d\beta_c$. The value of this gradient varies across the parameter space. We define it to be

$$d\beta_c = 2.5 \times 10^{-4} \times d\beta_{\text{inf}}^2, \quad (58)$$

where $d\beta_{\text{inf}}$ is the gradient at the inflection point. The connection frequency is then found by expanding the gradient of the curve about the maximum as a Taylor series. We find the connection frequency is given by

$$f_c = f_{\text{max}} + \frac{1}{\beta''' } \left[-\beta'' + \sqrt{\beta''^2 + 2\beta''' d\beta_c} \right], \quad (59)$$

where f_{max} is the frequency at which the maximum occurs and β'' and β''' are the second and third derivatives of β evaluated at f_{max} , respectively.

In cases where the turnover is not present within the fitting region we instead define the connection frequency to be the lower frequency limit of the fitting region, thus ensuring β is still falling at this frequency. In this case,

$$f_c = \begin{cases} f_{\text{inf}} - 0.03, & f_{\text{inf}} \geq 0.06 \\ 3f_{\text{inf}}/5, & f_{\text{inf}} < 0.06, \end{cases} \quad (60)$$

where f_{inf} is the inflection point.

C. Full IMR expressions

The expressions describing the precession angles in each of the different regions are connected using piecewise C^1 -continuous functions.

The full IMR expression for α is

$$\alpha_{\text{IMR}}(f) = \begin{cases} \alpha_{\text{PN}} & 0 \leq f < f_1 \\ \alpha_{\text{interp}} & f_1 \leq f < f_2 \\ \alpha_{\text{MR}} & f_2 \leq f \end{cases} \quad (61)$$

where α_{PN} , α_{interp} and α_{MR} are the PN expression used to describe α during inspiral, the interpolating function used to describe the late inspiral angles in the region f_1 to f_2 and the phenomenological ansatz used which has been tuned to NR to describe the merger-ringdown angles respectively.

Across the majority of the parameter space, the merger-ringdown ansatz for β has a minimum immediately following the inflection point (as shown in the central panel of Fig. 10). In these cases, the full IMR expression for β is

$$\beta_{\text{IMR}}(f) = \begin{cases} k\beta_{\text{PN}} & 0 \leq f < f_c \\ \beta_{\text{MR}} & f_c \leq f < f_t \\ \beta_{\text{RD}} & f_t \leq f \end{cases}, \quad (62)$$

where β_{PN} is the PN expression for β including the higher-order amplitude corrections discussed in Sec. VIB, k is the rescaling function applied to these expressions as outlined above, β_{MR} is the phenomenological ansatz which has been tuned to NR in the merger-ringdown regime, and β_{RD} is the constant value of β to which the system settles down after merger, as discussed in Sec. IX B. We model this quantity by the minimum value of β in the merger-ringdown expression. f_t is correspondingly given by the frequency at which the minimum occurs.

In cases where β tends towards an asymptote immediately following the inflection point (which occur in some regions of parameter space beyond the fitting region), the full IMR expression for β is

$$\beta_{\text{IMR}}(f) = \begin{cases} k\beta_{\text{PN}} & 0 \leq f < f_c \\ \beta_{\text{MR}} & f_c \leq f \end{cases}. \quad (63)$$

We would physically expect β to be bounded by 0 and π across the parameter space. In order to enforce this requirement, we pass the resulting β_{IMR} through a windowing function $w(\beta)$ given by

$$w(\beta) = \text{sgn}\left(\beta - \frac{\pi}{2}\right) \left(\frac{\pi}{2}\right)^{1-p} \arctan^p \left[\left(\frac{\beta - \frac{\pi}{2}}{\left(\frac{\pi}{2}\right)^{1-p}} \right)^{\frac{1}{p}} \right] + \frac{\pi}{2}, \quad (64)$$

where $p = 0.002$. This function is linear with $w(\beta) = \beta$ over the range $\beta \in [0.01, \pi - 0.01]$ to within 0.045%. This ensures that the fits for β are unaffected within the calibration but that β is bounded by 0 and π across the whole of parameter space.

The precession angle γ is then calculated over the entirety of the frequency range for which the waveform is produced by enforcing the minimal rotation condition given in Eq. (A7). The decision to do this rather




FIG. 9. Comparison of the complete model for each of the precession angles (thick red line) with the NR data (thin black line). The MSA angles (blue dotted line) are shown for reference. The left hand column shows the case with $(q, \chi, \theta) = (1, 0.4, 30^\circ)$. The right hand column shows the case with $(q, \chi, \theta) = (8, 0.8, 60^\circ)$. The vertical black lines show the connection frequencies for α and β .

than produce a separate model for γ was made as it was found that γ must be very accurate in order to consistently transform between an inertial frame and the co-precessing frame. The very small discrepancy between the expression for γ presented in [82] and the numerically calculated value is sufficient to seriously degrade the model. This discrepancy is exacerbated here since we are no longer using the dynamical expression for β presented in [82]. (We note that independently integrating Eq. (A7) was also found to be more accurate in the SEOBN-RPV4HM and PHENOMTPHM models [18, 25].)

The full model of these angles is shown for two examples in very different parts of the parameter space in Fig. 9.

D. Behaviour beyond calibration region

As with any tuned model, beyond the calibration region there is no guarantee of the accuracy of the model for the angles. However, we want to ensure that they do not display pathological or obviously physically incorrect




FIG. 10. Possible morphologies of the ansatz given by Eq. (51) depending on the values taken by the co-efficients in different regions of the parameter space. From left to right the panels show systems with $(q, \chi, \theta_{\text{LS}}) = (8, 0.2, 155^\circ)$, $(2.5, 0.4, 90^\circ)$ and $(5, 0.8, 160^\circ)$. The red dots mark the extrema, the green crosses show the inflection points and the blue dot indicates the inflection point chosen as described in Sec. VIII D. The points of maximum curvature around this inflection point are shown by the black lines, which give a measure of the width of the turnover. The solid black line in the shaded region indicates the frequency region that will be used as the merger-ringdown portion of the full angle model. All cases within our calibration region will have the morphology shown in the middle panel; the outer panels show that a reasonable choice is made outside the calibration region.

behaviour.

For α there are a number of possibilities inherent in the ansatz to see either pathological or physically incorrect behaviour. We have implemented restrictions on the values taken by the co-efficients to ensure this does not occur and a visual inspection of the waveforms shows that we do not see any pathological features. We would see pathological behaviour for $A_3 < 0$ and physically incorrect behaviour for $A_1 < 0$ (α would decrease as a function of frequency) or $A_2 > 0$ (the dip in α would have the wrong sign). As it is only a small region of parameter space in which this might happen, we enforce the conditions that $A_1, A_3 > 0$ and $A_2 < 0$ by taking the absolute value of the co-efficients with the appropriate sign. For A_2 we replace any positive values with zero. A_1 and A_2 take the wrong sign for systems with $q < 10$ only at very small spins ($\chi < 0.1$) or large anti-aligned spins ($\chi\sqrt{-\cos\theta} \sim 0.7$). For A_2 there is an additional region for $q > 7$ around $\chi = 0.4$ for anti-aligned spins ($\cos\theta > 0.75$). A_3 does not go negative within the calibration region, though this does start to occur for $q > 10$.

We see pathological behaviour for $B_4 \lesssim 0$. Physically incorrect behaviour starts to emerge when B_4 drops below $\mathcal{O}(10^2)$. In order to avoid such behaviour we require $B_4 \geq 175$ and replace the fitted value of B_4 by 175 where it falls below this value. Since $B_4 > 175$ throughout the calibration region this concern only arises for systems where the accuracy of the model cannot be guaranteed anyway (such as $q > 10$).

The morphology of the merger-ringdown ansatz of β also changes in some parts of the parameter space outside the calibration region, as shown in Fig. 10. We can ensure we always employ the correct part of the expression (for which β displays a drop at merger) in our model by selecting the correct inflection point. The inflection points of an expression occur at the roots of the second derivative of the expression. The second derivative of

Eq. (51) takes the form

$$\beta''(f) = \frac{af^3 + bf^2 + cf + d}{(1 + B_4(B_5 + f)^2)^3}, \quad (65)$$

where a, b, c and d are functions of the fitting co-efficients B_1, B_2, B_3, B_4 and B_5 . In order to find the roots of this cubic we re-write it in the form of a depressed cubic

$$x'^3 + px' + q = 0, \quad (66)$$

where

$$x' = x + \frac{b}{3a}, \quad (67)$$

$$p = \frac{3ac - b^2}{3a^2}, \quad (68)$$

$$q = \frac{2b^3 - 9abc + 27a^2d}{27a^3}. \quad (69)$$

In the case where this expression has three real roots, these are given by

$$x' = 2\sqrt{-\frac{p}{3}} \cos \left[\frac{1}{3} \arccos \left(\frac{3q}{2p} \sqrt{-\frac{3}{p}} \right) - \frac{2n\pi}{3} \right], \quad (70)$$

where $n = 0, 1, 2$.

We want to be able to define a single, smoothly varying inflection point that tracks the location of the turnover in β during merger across the parameter space. As the co-efficients of the cubic vary, the morphology of Eq. (51) changes, as shown in Fig. 10. For $a < 0$ we have the morphology shown in the central panel of the figure. We therefore select the central root, which is the only one with a negative gradient. For $a > 0$, we have the morphology shown in the outer panels. For this morphology we need to distinguish between the two outer roots, which both have a negative gradient. This is determined by the

“shift” of the roots, $b/3a$. In cases where

$$\frac{b}{3a} > \frac{B_5}{2} + \frac{\lambda_{004}^{B_2}}{4\lambda_{004}^{B_3}}, \quad (71)$$

where the λ_{pqr}^i are the co-efficients given in Eq. (52), we choose the first root (as seen in the left-hand panel), otherwise we choose the final root (as seen in the right-hand panel). This condition was found to select the correct root across the entire calibration region for the model as well as most of the extended regions encompassing the validation waveforms.

In the case where we have complex roots, two of the roots will be in the complex plane while one will be on the real axis. In this case we select the only real root.

We also consider the case where $a = 0$ and the second derivative is a quadratic. In this case we have only one root with a negative gradient, which is the desired root. Finally, we consider the case where both $a = 0$ and $b = 0$. Here we have only one root which gives us the desired inflection point.

Enforcing these conditions gives us a smoothly varying value of the inflection point across the parameter space and ensures our expression for β always has the correct morphology, dropping off at merger.

IX. PHYSICAL FEATURES OF THE WAVEFORMS

In motivating, constructing and presenting the PHENOMPNR model, we have observed several features of precessing-binary waveforms that deserve more detailed discussion.

A. Ringdown frequency

As discussed in Sec. V, in previous PHENOM models, the co-precessing-frame model consists of an aligned-spin model, with ringdown frequency and damping time adjusted according to the values predicted for the full precessing configuration. This prediction was made by using approximate NR fits for the final mass and spin, which then imply, via perturbation theory, the ringdown frequencies [8, 85, 91, 92]. This prediction of the ringdown frequency was then used in the co-precessing-frame model [9, 85, 86].

One interesting feature of this approach is that in some parts of the parameter space it leads to a discontinuity in the ringdown-frequency estimate. This arises as follows. There are two possibilities of ringdown frequency for a given BH spin, depending on whether the BH perturbations were generated by orbits that were prograde or retrograde with respect to the final BH spin; this can be represented as choosing either a positive or negative final spin [8, 91]. As an example, consider configurations with mass ratio $q = 8$ and a spin on the larger BH of $\chi = 0.8$.

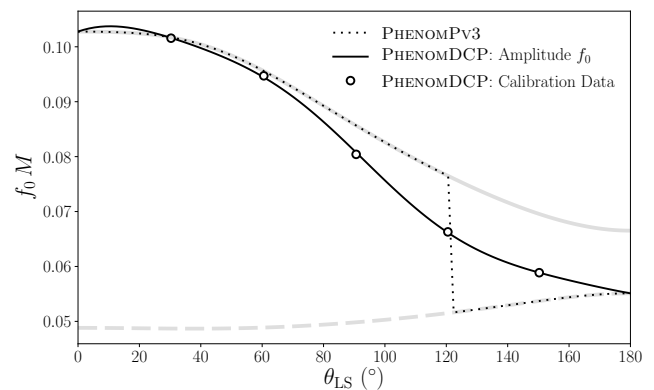


FIG. 11. Effective frequency-domain ringdown frequencies for $(q, \chi) = (8, 0.8)$, as modelled by PHENOMPv3 and PHENOMDCP. Additional lines show QNM frequencies predicted from standard perturbation theory methods using the remnant BH’s mass and spin [91]. The solid thick grey line traces prograde QNM frequencies, and the dashed thick grey line traces the retrograde QNM frequencies. All curves are bound between pro- and retrograde QNM frequencies. PHENOMPv3 displays a discontinuity near $\theta_{LS} = 120^\circ$, while NR data and PHENOMPv3 do not.

If the spin is aligned with the orbital angular momentum, we predict that after merger the final BH will have a spin of 0.86, and a ringdown frequency of ~ 0.1 . If the large BH spin is anti-aligned to the orbital angular momentum, i.e., $\theta_{LS} = 180^\circ$, then the final BH spin is -0.275 , and the ringdown frequency is ~ 0.06 . (These values were calculated using the PHENOMD final-spin fit [8], which is also what is used in our new model; at this level of accuracy any more recent fit will give similar results.)

We can now ask, what happens for other values of θ_{LS} ? One way to estimate the final spin is as follows. We first estimate the final spin for an equivalent aligned-spin binary, χ_{AS} , and then calculate the vector sum of this aligned spin with the in-plane spin contribution χ_P , which, in our single-spin example above, would take the value $\chi \sin(\theta_{LS})$. The final spin can then be estimated as,

$$\chi_f = \pm \sqrt{\chi_{AS}^2 + (m_1/M_f)^2 \chi_P^2}. \quad (72)$$

This approach was first used in the PHENOMP model described in Ref. [10], with more explicit details provided in the presentation of PHENOMPv3 in Ref. [12]; a similar approach is also followed in all current precessing-binary PHENOM models [13, 16, 45].

When we use this final-spin estimate to calculate the ringdown frequency, we must choose a sign. In previous PHENOM models, the same sign was chosen as χ_{AS} , but in some cases (as in the example above), this means that χ_f swaps sign at some value of θ_{LS} , and the resulting estimate of the ringdown frequency is discontinuous. This is illustrated by the dashed line in Fig. 11 for our $q = 8$, $\chi = 0.8$ series of configurations. As an estimate of the

ringdown frequency in the ($\ell = 2, |m| = 2$) multipoles in the J -aligned frame, this approach appears to be quite accurate, including the sharp transition from prograde to retrograde branches.

One issue with this approach is that the transformation from the co-precessing to inertial frame will introduce a shift in the GW frequency, and therefore a change in the ringdown frequency. If we apply the correct inertial-frame ringdown frequency to our co-precessing-frame model, it will be changed when the angle model is applied, and the final model will have the *wrong* ringdown frequency. This is what happens in previous PHENOM models. We could take this shift into account when we prescribe the ringdown frequency in the co-precessing frame, but instead we simply produce a phenomenological fit to the ringdown frequency in the construction of the co-precessing-frame model PHENOMDCP. This is also shown in Fig. 11, in comparison with the effective co-precessing-frame ringdown frequency that we find from the NR data.

B. The collapse of β through merger

During the inspiral, the angle β is related to the opening angle between the total and orbital angular momenta, i.e., the opening angle of the precession cone. At merger the orbital motion ceases, and we are left with a ringing black hole, and would expect that the corresponding optimal emission direction would relax to the $\hat{\mathbf{J}}$ direction of the final black hole. However, we may also consider an alternative picture. A stationary BH does not radiate. We may perturb a non-spinning BH such that we completely determine the dominant emission direction as the perturbation rings down. Adding spin to the BH, either small or large, does not change this freedom. Thus the optimal emission direction after merger, and in particular, the final values of α and β , may encode information about how the remnant BH was perturbed through merger, and the relationship to $\hat{\mathbf{J}}$ is not so clear.

In Ref. [55] an attempt was made to describe the late-time precession behaviour using results from perturbation theory. We know the general form of the ringdown signal,

$$h_{\ell m}(t) \approx A_{\ell m} e^{i\omega_{\ell m} t} e^{-t/\tau_{\ell m}}, \quad (73)$$

where the $A_{\ell m}$ are unknown constants, and $\{\omega_{\ell m}, \tau_{\ell m}\}$ are determined by the mass and spin of the final BH through perturbation theory. Given this general form, we can predict the general behaviour of the precession angles α and β in the ringdown regime, similar to the approximate approach followed in Sec. VI B. Ref. [55] note that, if we consider only the dominant $\ell = 2$ modes, the QA direction precesses around $\hat{\mathbf{J}}$ with a frequency $\omega_{22} - \omega_{21}$, and β either falls exponentially to zero at a rate given by $\tau_{22} - \tau_{21}$, or grows exponentially to π , depending on the relative magnitude of the two damping times. A similar calculation was later discussed in Refs. [18, 25, 56].




FIG. 12. Comparison of analytic ringdown estimate [Eq. (57b) of Ref. [18]] and NR for (top) $(q, \chi, \theta_{\text{LS}}) = (4, 0.4, 60^\circ)$ and (bottom) $(q, \chi, \theta_{\text{LS}}) = (8, 0.4, 30^\circ)$.

Several points are worth noting. (1) For much of the parameter space, although the decay of $\beta(t)$ is exponential, it is nonetheless extremely slow, and on a much longer timescale than the decay of the signal amplitude. (2) We can consider non-zero ($\ell = 2, |m| = 2$) and ($\ell = 2, |m| = 1$) multipoles where the QA direction does not precess at all, for example aligned-spin binaries. (3) Just as BH perturbation theory cannot tell us how much each QNM is excited [93, 94], this analysis cannot tell us the magnitude of β at whatever point we wish to designate as the beginning of the ringdown regime.

We now turn to our NR data to address these points. Fig. 12 shows the late-time behaviour of NR β for the $(q, \chi, \theta_{\text{LS}}) = (4, 0.4, 60^\circ)$ and $(8, 0.4, 30^\circ)$ configurations, as well as approximate fits to the β decay rate predicted by the ringdown toy model discussed above. In particular, the analytic ringdown estimate shown in Fig. 12 is derived in Ref. [18]. In these fits the decay rate is prescribed by the toy model and only the overall amplitude is fit to the numerical data. The data are not clean enough to conclusively show that late-time β follows the decay rate predicted by the toy model, but the data are certainly consistent with that model. What is worth highlighting is that the decay rate is indeed very slow; we expect β to be greater than, say, 10% of its peak value, for several hundred M after merger, at which point the total signal amplitude will have decayed by several orders of magnitude. In this context, our simple approximation in PHENOMP NR, that late-time β is constant, appears to be justified.

The other important observation is that this late-time ringdown behaviour begins *after* β has dropped significantly through merger. This strongly suggests that ring-

down begins significantly after the peak in both strain and ψ_4 , which is possibly at tension with recent efforts to apply BH perturbation theory at those points [95].

Although a PN treatment can approximately describe β during the inspiral, and a simple ringdown analysis can describe the decay rate of β during ringdown, neither can capture the rapid drop in β through merger, or predict the value of β at the point where the ringdown behaviour takes over. This feature, which is included in PHENOMP NR, was not explicitly modelled in previous PHENOM and EOBNR models; PHENOMP/Pv2/Pv3/XP used the MSA angles at all frequencies, and both SEOBNRv4PHM and PHENOMTP use a constant late-time value of β determined by its value near merger.

C. Hierarchy in the turnover frequency of the $\ell = 2$ multipoles

The rapid drop in β described in the previous section results in a key feature of precessing waveforms: a hierarchy in the turnover frequency of the $\ell = 2$ multipoles. From Eq. (41) we can see that β is approximately given by the ratio of the amplitude of the (2,2) and (2,1) multipoles. The drop in β therefore implies that the amplitude of the (2,1) multipole must have decreased relative to the (2,2) multipole and so the (2,1) multipole will begin to experience ringdown decay before the (2,2) multipole. Once both multipoles are decaying exponentially (at roughly the same rate) β levels off. This trend continues for all of the $\ell = 2$ multipoles.

By capturing the drop in β in our model, we successfully model this hierarchy in the turnover frequency of the $\ell = 2$ multipoles, as seen in the top panel of Fig. 13. This feature has not been modelled in previous precessing PHENOM models, and the central panel of Fig. 13 shows the multipole hierarchy for PHENOMXP, which is also the behaviour for PHENOMPv3, since both use the same MSA angle model. We see that in these models each of the $\ell = 2$ multipoles turn over at the same frequency. SEOBNRv4P, shown in the bottom panel, does capture this hierarchy but the amplitude of the higher order multipoles is not well modelled. This is due to modelling ι rather than β , which typically overestimates the amplitude as discussed in Sec. VI B.

X. TIME DOMAIN VALIDATION

A. Time domain waveform

The improvements made in modelling precessing systems presented here — both to the underlying co-precessing model and the precession angles — can also be clearly seen when inspecting the waveforms in the time domain. As can be seen in Fig. 14, PHENOMP NR correctly captures the precession envelope and the phasing of the waveform through inspiral, merger and ringdown.




FIG. 13. Amplitudes of the $\ell = 2$ multipoles for the $(q, \chi, \theta_{\text{LS}}) = (4, 0.4, 90^\circ)$ configuration at $100 M_\odot$. The NR data are shown in black on all panels, with PHENOMP NR (top panel in blue), PHENOMPv3 and PHENOMXP (central panel in purple) and SEOBNRv4P (bottom panel in red). The vertical line indicates the frequency of the peak (2,2) amplitude. The small vertical line segments indicate the turnover frequency of each of the subdominant modes.

This figure also clearly shows that the frequency-domain modelling presented here does not introduce any strange artefacts in the time domain. For comparison we also show SEOBNRv4P, a naturally time domain precessing model. The configurations shown here are for a binary with the intrinsic properties $(q, \chi, \theta_{\text{LS}}) = (4, 0.8, 60^\circ)$ (top panel) and $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 60^\circ)$ (bottom panel). We have plotted both waveform models for a binary taken to be face-on ($\theta_{\text{LN}} = 0^\circ$) at the reference frequency, which is given in Tab. I. We particularly note the good agreement between PHENOMP NR and the NR waveform after merger (the shaded region) due to accurately modelling the merger-ringdown precession angles and the effective ringdown frequency of the co-precessing waveform.

The time and phase alignment of the waveforms plotted in Fig. 14 has been performed over the same range of frequencies as were used in calculating the matches detailed in Sec. XI and quoted in the figure legend. This range is much greater than that shown in the plot so the deviations between the models and the NR seen here do not contribute as much as might naïvely be expected. We have plotted the waveform for the in-plane spin configu-




FIG. 14. A comparison of the time domain obtained from PHENOMP NR with the NR data. The top panel shows the case $(q, \chi, \theta_{\text{LS}}) = (4, 0.8, 60^\circ)$ while the bottom panel shows the case $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 60^\circ)$. Both are for a face on ($\theta_{\text{LN}} = 0^\circ$) binary with a total mass of $100M_\odot$. For comparison, we also show the waveform produced using SEOBNRv4P. The match values for the specific configuration for each of the waveforms plotted are given in the legend.

ration and polarisation of the signal for which we get the maximum match. PHENOMP NR agrees well with the NR data from inspiral through merger and ringdown, capturing both the precession envelope and the phasing of the waveform correctly.

B. Time domain angles

Accurately modelling the merger-ringdown features of the angles in the frequency domain has also enabled us to reproduce key features of the angles in the time domain after merger. We compared the time domain angles for four models (PHENOMPv3, PHENOMP XP, SEOBNRv4P and PHENOMP NR) with the NR angles. In order to avoid the introduction of artefacts due to unnecessary processing of the NR data, we compare against the time domain angles calculated using the cleaned and

symmetrised Ψ_4 data rather than h .

In order to calculate the precession angles for each of the models we require the $\ell = 2$ modes in the \mathbf{J} -frame. For each of the models investigated we used the \mathbf{J} -frame defined by the developers of the model.

The time domain angles for the frequency domain models (those belonging to the PHENOM family) are calculated as follows. First we compute the $\psi_{4,\ell m}$ $\ell = 2$ multipoles from the strain multipoles in the frequency domain using $\tilde{\psi}_{4,\ell m}(f) = (2\pi f)^2 \tilde{h}_{\ell m}(f)$. We then compute the time domain multipoles by performing the inverse Fourier transform each of the $\ell = 2$ frequency domain multipoles. Finally we calculate the precession angles from the set of time domain $\ell = 2$ multipoles.

For SEOBNRv4P, a time domain model, we differentiated each of the $\ell = 2$ time domain multipoles twice to get $\psi_{4,\ell m}$ from $h_{\ell m}$. We then calculated the precession angles using these multipoles. Since the connection be-




FIG. 15. Comparison of the time domain precession angles for the PHENOMPv3, PHENOMXP, SEOBNRv4P and PHENOMPvNR models with the NR data. These angles are for the case with $(q, \chi, \theta_{\text{LS}}) = (4, 0.4, 60^\circ)$.

tween the inspiral and ringdown parts of the models for the multipole moments and the precession angles used in SEOBNRv4P is C^1 -continuous, we see a discontinuity in the time domain angles presented here as a result of the double differentiation.

The results of this comparison are shown in Figs. 15 and 16. Since PHENOMPv3 and PHENOMXP use the same model for the precession angles with a different co-processing model, the time domain angles presented here agree very closely. The two most notable features in the time domain angles are the continued rise in α after the merger and the rapid drop in the value of β . If α takes a constant value it implies the precession of the optimum emission direction has stopped. As has been noted previously [55], this is clearly not seen in the NR data. This feature of the precessional motion is captured by SEOBNRv4P and PHENOMPvNR but not by PHENOMPv3 and PHENOMXP. The rapid drop in the value of β is captured accurately only by PHENOMPvNR, although SEOBNRv4P does show some evidence of a drop in the value of β . This shows we have managed to capture the closing up of the opening angle as the angular momentum is radiated away through gravitational wave emission. The final feature to note is the amplitude of β throughout inspiral is captured reasonably well by PHENOMPvNR whereas the other models all show a slight offset since (as previously discussed) they use the angles that describe the precessional dynamics rather than the precession of the direction of optimal emission. This can be seen more clearly at earlier times than are shown in Figs. 15 and 16 since here we chose to focus on the merger-ringdown region where data processing artefacts from the Fourier transform are stronger.




FIG. 16. Comparison of the time domain precession angles for the PHENOMPv3, PHENOMXP, SEOBNRv4P and PHENOMPvNR models with the NR data. These angles are for the case with $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 60^\circ)$.

XI. MODEL VALIDATION: MATCHES

We now wish to test the accuracy of our new precessing model in the context of gravitational wave signal analysis. To do this we calculate the match (using the method detailed below in Sec. XI A) between the NR waveform and our model for a given configuration. We performed three sets of matches in order to inspect each of the components of our model individually as well as the complete final model. To assess the accuracy of the underlying co-processing model we calculated the standard non-precessing match for a waveform containing only the $(2,2)$ -multipole between the co-processing model PHENOMDCP and the co-processing NR waveform. In order to assess the accuracy of the angle model itself, we model the precessing waveform by twisting up the co-processing NR waveform with the PHENOMANGLES angles and match it against the corresponding \mathbf{J} -aligned NR waveform. Finally, we assessed the accuracy of the complete tuned precessing model PHENOMPvNR by performing the SNR-weighted match between the model and the \mathbf{J} -aligned NR waveforms, containing the $\ell = 2$ multipoles. The match calculated between the NR waveform and the complete model will contain errors introduced by inaccuracies in both PHENOMDCP and PHENOMANGLES. Since we do not aim to model asymmetries in the multipole moments in this work, our model does not capture them. We therefore perform matches testing the angle model using the symmetrised NR waveform (in both the \mathbf{J} -aligned and co-processing frames).

A. Match Definitions

The disagreement between two waveforms, a model template h_t and an NR signal h_s , is quantified using the standard inner product weighted by the power spectral

density of the detector $S_n(f)$ [77], chosen for this work to be the noise spectrum of advanced LIGO at design sensitivity [96]:

$$\langle h_s | h_t \rangle = 4\text{Re} \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_s(f) \tilde{h}_t^*(f)}{S_n(f)} df. \quad (74)$$

The *match* is then given by the inner product between two normalised waveforms,

$$\mathcal{M}(h_s, h_t) = \max_{\Xi_t} \frac{\langle h_s | h_t \rangle}{\sqrt{\langle h_s | h_s \rangle \langle h_t | h_t \rangle}}, \quad (75)$$

maximised over a set of template parameters Ξ_t described below.

Time shifts and reference phase shifts have no physical effect on the signal; a time shift corresponds only to a change in the merger time of the binary, while a change in the phase corresponds to a change in the initial orientation of the binary's orbit. For non-precessing waveforms containing only the (2,2)-multipole, the resulting match value is independent of the inclination and polarisation of the signal, as changes to the inclination simply re-scale the overall amplitude of both the signal and template, and the polarisation is degenerate with the reference phase and therefore optimised away. When computing the match for non-precessing signals, as is done in Sec. XIC, the maximisation done in Eq. (75) is done over time and phase shifts, $\Xi_t = \{t_0, \phi_0\}$.

For precessing waveforms, both the inclination and polarisation must be taken into account. First, we compute the match outlined in Eq. (75) whilst keeping the signal phase and polarisation fixed, and maximise over time shifts, reference phase and template polarisation following Ref. [97]. We further optimise over rotations to the in-plane spin components of the template at the reference frequency as in Ref. [16], which effectively optimises the match over the initial precession phase α_0 , *i.e.*, $\Xi_t = \{t_0, \phi_0, \psi_0, \alpha_0\}$. We then follow previous efforts to quantify precessing models [12, 16, 61] and introduce an *SNR-weighted match*.

The SNR-weighted match is computed by averaging the match computed at each given signal phase and polarisation whilst volume-weighting with the SNR of the signal,

$$\mathcal{M}_w = \left(\frac{\sum_{\psi_s, \phi_s} \mathcal{M}^3 \langle h_s | h_s \rangle^{\frac{3}{2}}}{\sum_{\psi_s, \phi_s} \langle h_s | h_s \rangle^{\frac{3}{2}}} \right)^{\frac{1}{3}}, \quad (76)$$

where we have summed over the values of signal phase and polarisation, ϕ_s and ψ_s , respectively. This is done to better account for the large variation in detectability and signal strength with sky location that occurs in precessing signals.

Finally, we compute the *mismatch* between the signal and template for non-precessing signals as,

$$\mathfrak{M} = 1 - \mathcal{M}, \quad (77)$$

and similarly for precessing signals the SNR-weighted mismatch,

$$\mathfrak{M}_w = 1 - \mathcal{M}_w. \quad (78)$$

B. Verification waveforms

We performed matches against 76 of the waveforms taken from the BAM catalogue described in Sec. II. We also considered an additional set of waveforms taken from the SXS [64, 98] and Maya catalogues [66]. This enabled us to test the accuracy of the model for configurations for which it was not tuned, including two-spin configurations. The additional verification waveforms from the SXS and Maya catalogues were chosen to include those cases that gave particularly poor matches for PHENOMPv3 (since that is the model we aim to directly improve upon) with a few cases of additional interest such as SXS1397 (a long two-spin waveform) and precessing systems with $q > 4$. The additional BAM waveforms provide a systematic coverage of the parameter space over which we have tuned our model.

A summary of the waveforms taken from the BAM catalogue are given in Table I, while the details of those taken from the SXS and Maya catalogues are in Table II. Only the subset of waveforms taken from the BAM catalogue were used to study the accuracy of the individual components of the model; the underlying co-precessing model and the model for the precession angles. The complete set of waveforms, taken from all three catalogues, was used to test the accuracy of the full model over a range of total masses for the system.

C. Matches: Accuracy of the co-precessing model

We computed the match between various models for the co-precessing waveform and the co-precessing NR waveform. We considered a system of total mass $100M_\odot$ and performed the match over the frequency range for which the NR data was available; from $(f_{\text{ref}} + 5)\text{Hz}$ to 244Hz . The value of the reference frequency f_{ref} for each simulation is given in Table I. The co-precessing-frame models we consider are PHENOMD, PHENOMXAS and PHENOMDCP. For both PHENOMD and PHENOMXAS we considered the version where the final spin has been modified to take into account the in-plane spin component for precessing systems.

As can be seen from Fig. 17, the assumptions that go into producing the aligned-spin mapping used in the production of modified PHENOMD and modified PHENOMXAS become less accurate as both mass ratio and spin are increased. PHENOMDCP performs better than both modified PHENOMD and modified PHENOMXAS for almost all cases, with the most noticeable improvement for the higher mass ratio, high-spin cases where we are in greatest need of a tuned co-precessing model. In the cases




FIG. 17. Mismatches for each of the BAM calibration and verification waveforms, at a total mass of $100M_{\odot}$. Mismatches are between the symmetrised co-precessing NR waveforms and PHENOMDCP (purple diamonds), modified PHENOMD (blue circles) and modified PHENOMXAS (red squares). The configuration mass ratio increases from left to right (with $q \in \{1, 2, 4, 8\}$). Solid black lines separate cases mass ratios and dotted lines separate spin magnitudes.




FIG. 18. SNR-weighted mismatches for the same configurations as in Fig. 17, averaged over inclination. These mismatches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (purple diamonds) and twisted up with the angle model used by PHENOMPv3 (steel blue triangles).

where PHENOMDCP has a similar or slightly worse performance than either of the other two models the match is generally already comparable to the accuracy level of our input NR waveforms.

D. Matches: Accuracy of the angle model

In order to test the accuracy of the angle model we constructed a set of precessing waveforms by calculating the symmetrised frequency-domain co-precessing NR waveform containing only the $\ell = 2$ multipoles and “twisting” this waveform up with the modelled precession angles. We have previously seen the angle model performs well qualitatively in Sec. VII. We choose to study the match here as it provides the most physically meaningful way to assess the significance of the angle model as a whole. We constructed two sets of precessing waveforms in this fashion; one using the model for the angles presented in this paper, and the other using the MSA angles, in order to quantify the effect of modelling the merger-ringdown

behaviour of the angles. We then calculated the SNR-weighted match between these waveforms and the symmetrised NR waveforms in the \mathbf{J} -aligned frame comprising only the $\ell = 2$ multipoles. As with the co-precessing matches described above, these matches were calculated at a fixed total mass $M = 100M_{\odot}$ and performed over a frequency range from $(f_{\text{ref}} + 5)\text{Hz}$ to 244Hz (the frequency range for which the NR data was available).

In Fig. 18, we have shown the inclination average of the full precessing match for ease of presentation. We can see that the matches using the improved angle model are above 0.99 across the majority of the parameter space. The only cases for which this is not true are in the most extreme corner of the parameter space we modelled; cases with $q = 8$, $\chi = 0.6$ and $\theta_{\text{LS}} \geq 90^\circ$. In these cases we find the PN expressions used for α during inspiral deviate from those calculated from the NR waveform at reasonably low frequencies. In the case of $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 120^\circ)$ this is before the start of the NR waveform, as shown in Fig. 19. Improving the model for these cases would require a model for the interme-




FIG. 19. Comparison MSA α (blue dashed line) with the value calculated from the NR waveform (black solid line) for the $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 120^\circ)$ configuration. In order to see a region over which the two values agree well we would need a longer NR waveform; see text for more details.

diate region between where the PN expression ceases to be accurate and where the current model begins, which may require longer NR waveforms to be produced. Additionally, we expect that modelling this intermediate region will improve matches for several other cases as well, where the PN expressions for the angles deviate from what we see in the NR data at lower frequencies than are covered by our current merger-ringdown model for the angles. Nonetheless, in all cases we see significant improvement over the previous model.

The best matches are seen in the least extreme part of parameter space; namely for low mass ratio systems. This is the region of parameter space where existing models for the angles already perform reasonably well. The biggest improvement in the matches as a result of the improved model for the angles is seen at higher mass ratios, particularly for larger θ_{LS} .

For a selection of these cases we show the mismatch as a function of θ_{JN} in Fig. 20. The figure shows both the SNR-weighted average, and the range of mismatches with respect to signal polarisation and phase. We see that the mismatches against symmetrised NR waveforms are approximately symmetric about $\theta_{\text{JN}} = \pi/2$. The MSA angles generally give the worst SNR-weighted average mismatch for systems with $\theta_{\text{JN}} = 0, \pi$, although this is not always the case, and the variations with respect to different choices of polarisation and phase are often larger than those with respect to inclination. This mismatch then typically improves as it approaches $\theta_{\text{JN}} = \pi/2$ systems, with a slight increase for systems at exactly $\pi/2$ in most cases. In contrast, the SNR-weighted average mismatches involving the new angle model show one of two main behaviours with respect to inclination: the first gives the lowest mismatches for systems with $\theta_{\text{JN}} = 0, \pi$ with a marked degradation towards $\theta_{\text{JN}} = \pi/2$, while the second shows approximately constant values for the mismatch with respect to inclination, with a possible slight improvement for systems with $\theta_{\text{JN}} = \pi/2$. However, we

do not observe any clear pattern in how these two trends manifest themselves across the parameter space. The most important result to note is that in comparing the new angle model with the MSA angles in this figure, for the new angle model the lowest mismatch is always better (the lower edge of the envelopes), the highest mismatch is always better (the upper edge of the envelopes), and the SNR-weighted mismatch is always better.

In general we might expect errors in the angle models to lead to worse mismatches for edge-on configurations, since at these orientations the contributions of the subdominant $\ell = 2$ multipoles are largest, and the strength of those multipoles in our model is directly related to the precession angles, in particular β . However, $\theta_{\text{JN}} = \pi/2$ doesn't necessarily correspond to the binary being edge-on to the detector, unless β is close to zero; in general, a system viewed from $\theta_{\text{JN}} = \pi/2$ is *never* edge-on. Because of this, and because of the large variation in matches across the cases shown in Fig. 20, we revisit this question in the full-model mismatches in the next section, where we specify the binary orientation at the beginning of the waveform (so $\theta_{\text{LN}} = \pi/2$ corresponds to edge-on at least at one point in the inspiral), and consider an exhaustive set of masses, orientations and polarisations for every NR waveform.

E. Matches: Accuracy of PHENOMP NR

In this section we compare the accuracy of the complete PHENOMP NR model to existing precessing waveform models by computing SNR-weighted mismatches between these approximants and the various NR waveforms detailed in Sec. XIB. Each SNR-weighted mismatch is computed over a range of total masses $M_{\text{total}} \in [100, 120, 140, 160, 180, 200, 220, 240]M_{\odot}$ and at four inclination values, $\theta_{\text{LN}} \in [0, \pi/6, \pi/3, \pi/2]$, specified at the reference frequencies given in the waveform tables. The choice to sample in θ_{LN} , rather than θ_{JN} as done above, was motivated partially by the frame convention of LALSuite [59], which specifies that the LAL inertial frame [99] in which the waveforms are generated be instantaneously $\hat{\mathbf{L}}$ -aligned at the given reference frequency. This choice allows for comparisons with match results already present in the literature. As was also noted in the previous section, the conventional wisdom gleaned from non-precessing signals regarding model performance and the importance of higher multipoles for configurations with $\theta_{\text{LN}} \sim \pi/2$ also holds for precessing cases where β remains small throughout most of the inspiral, as is the case with most of the NR waveforms we consider. The matches were performed starting at a frequency of 20Hz or $f_{\text{ref}} + 5\text{Hz}$, whichever was higher, with f_{ref} listed for each NR waveform in Tables I and II. The nominal starting frequency of 20Hz was chosen to match the approximate low-frequency cut-off of typical signal analysis, and dictated our choice of $100 M_{\odot}$ as the lowest total mass we consider.

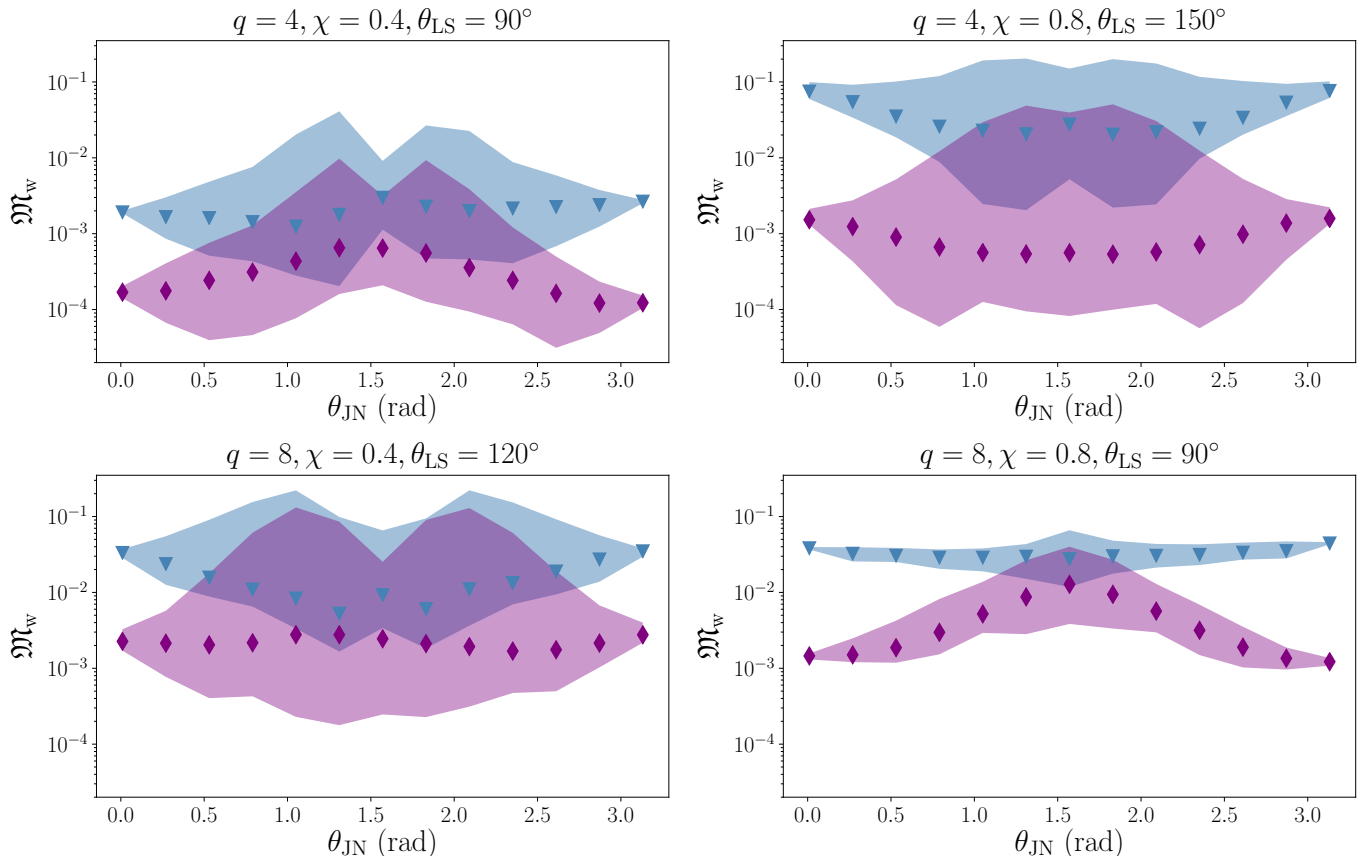


FIG. 20. Mismatch as a function of the inclination of the binary, quantified by the angle between the line of sight and the total angular momentum θ_{JN} , for four cases, at $100M_{\odot}$. These mismatches consider the co-precessing NR waveform twisted up with the angle model used by PHENOMPv3 (steel blue) and PHENOMPnr (purple). The solid markers show the SNR-weighted average mismatch while the shaded regions show the variation with respect to signal polarisation and phase.

1. Match variation with inclination

As discussed in the previous section, we would expect improvement in PHENOMPnr to be most apparent when trying to replicate highly-precessing signals at high inclination, $\theta_{\text{LN}} \sim \pi/2$, where the modulations in the signal due to precession grow stronger as more power is distributed across the $\ell = 2$ multipoles. We therefore compare the performance of PHENOMPnr with the earlier precessing model PHENOMPv3, plotting the SNR-weighted mismatches between these two models and the NR waveforms for each inclination value used. We choose PHENOMPv3, since it was the base model that we modified to produce PHENOMPnr, and this comparison provides the most direct measure of the level of improvement achieved by including NR-tuned precession effects in both the co-precessing-frame and angle models. The overall distribution of SNR-weighted mismatches is shown in Fig. 21. For PHENOMPv3, which uses the uncalibrated MSA angles, the performance noticeably degrades as the signal inclination increases, whereas the mismatches for PHENOMPnr remain relatively unchanged with respect to changes in inclination. This is largely consistent with

Fig. 20, where the average match shows little variation with respect to inclination for three out of the four configurations shown, and suggests that the behaviour of the model for the $(q, \chi, \theta_{\text{LS}}) = (8, 0.8, 90^\circ)$ case is atypical.

2. General match results

We compare the performance of PHENOMPnr against the precessing waveform approximants PHENOMPxp, SEOBNRv4P, and NRSUR7DQ4. The full results are shown for all inclinations and total masses in Fig. 22, and the mass- and inclination-averaged SNR-weighted mismatches are shown per waveform in Fig. 23. The model NRSUR7DQ4 was calibrated only up to $q = 4$, and while its implementation in LALSuite allows for extrapolation beyond this, we choose to limit the comparison with this model to the subset of the available NR waveforms with $q \leq 4$ to ensure accuracy is maintained.

Overall we see an improvement in the mismatches between PHENOMPnr and the NR waveforms compared to PHENOMPxp and SEOBNRv4P. The mismatch results show comparable performance between PHENOMPnr

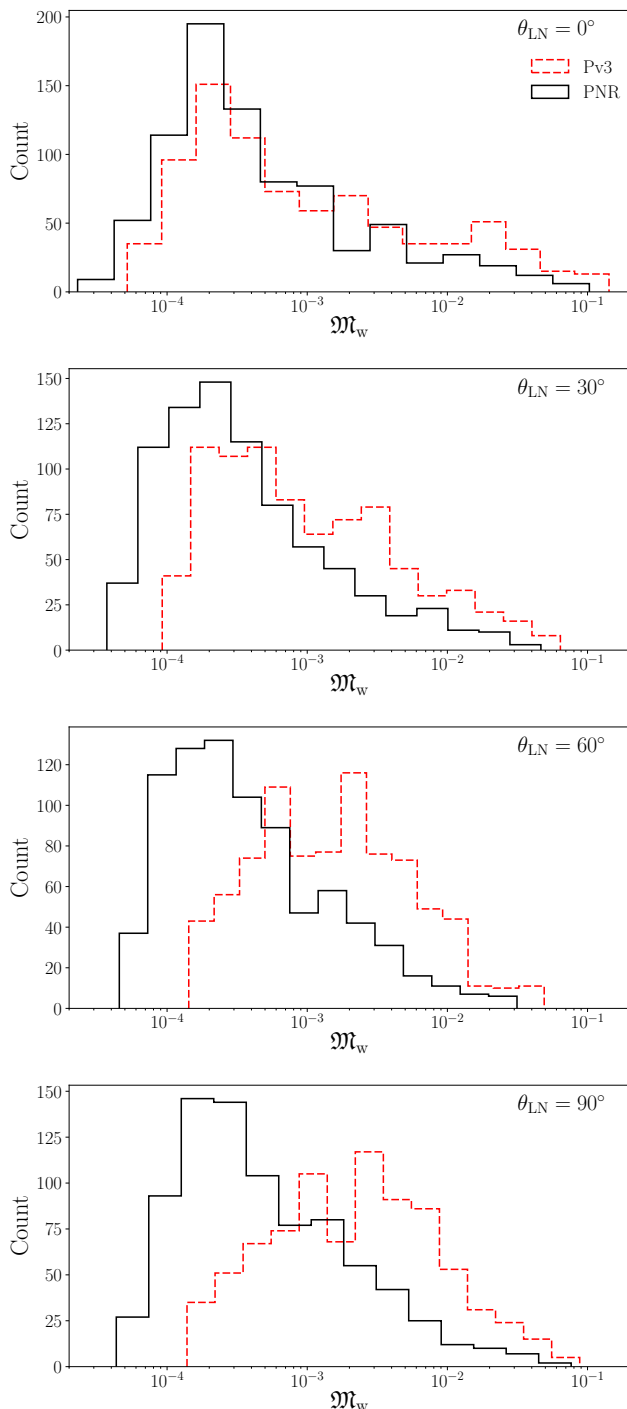


FIG. 21. Histograms of the SNR-weighted mismatches between the NR waveforms listed in Tables I-II and the waveform models PHENOMPNR and PHENOMPv3. Each subplot contains the SNR-weighted mismatches for all total masses separated by inclination descending as $\theta_{LN} \in [0^\circ, 30^\circ, 60^\circ, 90^\circ]$. The mismatches for all total mass values listed in Sec. XI E are included at each inclination. The results for PHENOMPNR are present in solid black while the results for PHENOMPv3 are given in dashed red.

and NRSUR7DQ4, but we caution a reminder that NRSUR7DQ4 is a model that does not make the simplifying assumptions outlined in Sec. III, and while the effects of these additional physical features are generally small, we expect that their presence in NRSUR7DQ4 compared to the NR data used for this comparison would bias the results toward slightly higher mismatches. Nonetheless, it is encouraging to observe that the PHENOMPNR model, while tuned to a comparatively small number of waveforms over a large configuration parameter space, and using a simple set of model ansätze, and several simplifying assumptions, in general has comparable mismatches to the NRSUR7DQ4 model.

From Fig. 23 it is apparent that the overall mismatch increases with mass-ratio, and for each mass ratio the mismatch generally worsens with increasing spin magnitude. This behaviour is largely responsible for the long tails toward high mismatch value present in Fig. 22. Such a trend is also visible in Figs. 17-18. A simple explanation for this observation arises from the PN scaling of the opening angle ι with symmetric mass ratio and spin magnitude in quasi-circular binaries with simple precession [30],

$$\sin \iota = \frac{S_{\perp}}{\sqrt{(\eta\sqrt{MR} + S_{\parallel})^2 + S_{\perp}^2}}, \quad (79)$$

where M is the system's total mass and R its orbital separation. A larger opening angle increases the impact of precession modulations on the signal, and these are where model inaccuracies will be most apparent. One would similarly expect to see worsening mismatches as S_{\perp} is maximised, *i.e.*, $\theta_{LS} = 90^\circ$ for single-spin cases; however this trend is not as apparent in the results. The results in Fig. 23 show that PHENOMPNR is an improvement over PHENOMXP in the most extreme region of parameter space for $q = 8$, $\chi \in [0.6, 0.8]$, while SEOBNRv4P yields better results in this region when $\theta_{LS} > 90^\circ$.

Regarding the performance of PHENOMPNR for the two-spin NR cases listed in Table II, specifically cases 8-19 and cases 23-27, we observe that PHENOMPNR and PHENOMXP perform surprisingly similarly for these cases, both for the SXS and MAYA cases, whereas PHENOMPNR provides a general improvement over SEOBNRv4P for the two-spin cases. These results provide a reassuring validation of the single-spin mapping detailed in Sec. IV.

Finally we remark on the impact of the fixed- $\hat{\mathbf{J}}$ assumption used in the modeling of PHENOMPNR and outlined in Sec. IA. We computed the SNR-weighted mismatches between the raw NR signals in an initially $\hat{\mathbf{J}}$ -aligned frame and those in the fixed- $\hat{\mathbf{J}}$ frame and find that the resulting mismatches are more than an order of magnitude lower than the mismatches between PHENOMPNR and the fixed- $\hat{\mathbf{J}}$ frame NR signals presented in this section, and in all cases lower than 5.1×10^{-4} at $100M_{\odot}$. The full comparison is displayed in Fig. 24, and shows that the fixed- $\hat{\mathbf{J}}$ approximation remains valid over




FIG. 22. Histograms of the SNR-weighted mismatches between various models in comparison and the NR waveforms listed in Tables I-II. The mismatches for all inclination and total mass values listed in Sec. XI E are included. In all three subplots, the results for PHENOPNR (“PNR”) are presented with a solid black outline, with the other model results given with dashed outlines from left to right as PHENOMP (“XP”) in red, SEOBNRv4P (“EOB”) in blue, and NRSUR7DQ4 (“SUR”) in green. For the comparison plot between PHENOPNR and NRSUR7DQ4 we only include results of NR waveforms for which both models are run.




FIG. 23. SNR-weighted mismatches averaged over total mass and inclination between the precessing waveform models PHENOPNR (“PNR”), PHENOMP (“XP”), and SEOBNRv4P (“EOB”), and the NR waveforms listed in Tables I-II, shown in order of the table listings. For the BAM cases, the solid vertical lines separate cases by mass ratio, and the dashed vertical lines separate spin magnitude. For the SXS and Maya cases, the solid vertical line splits by NR catalogue, and the dashed vertical line indicates a transition from single-spin to two-spin cases.




FIG. 24. SNR-weighted mismatches computed at $100M_{\odot}$ and averaged over inclination between the precessing waveform model PHENOMP NR (“PNR”) and the NR waveforms listed in Table I, shown in order of the table listings. Alongside these results are plotted the SNR-weighted mismatches computed between the NR waveforms in the initially $\hat{\mathbf{J}}$ -aligned frame and the fixed- $\hat{\mathbf{J}}$ frame (“NR”). The solid vertical lines separate cases by mass ratio, and the dashed vertical lines separate spin magnitude.

a broad range of parameter space where $\theta_{\text{LS}} < 90^{\circ}$ but begins to break down for systems with higher mass ratio and opening angle, implying that future modelling efforts should take care to re-evaluate the validity of this approximation in more extreme regions of parameter space.

XII. CONCLUSION

We have presented a new model of the GW signal from the inspiral, merger and ringdown of precessing non-eccentric black-hole binaries, PHENOMP NR. This is the first model to explicitly calibrate precession effects through merger and ringdown to NR simulations, and to use higher-order PN amplitude terms to consistently define a signal-based co-precessing frame (the “quadrupole aligned” (QA), or “optimal emission direction”) throughout the model.

The model is calibrated to 40 NR simulations of binaries where only the larger black hole is spinning; the simulations cover four mass ratios ($q = 1, 2, 4, 8$), two spin magnitudes ($\chi = 0.4, 0.8$), and five values of the spin misalignment angle ($\theta_{\text{LS}} = 30^{\circ}, 60^{\circ}, 90^{\circ}, 120^{\circ}, 150^{\circ}$). In the frequency domain we separately model the co-precessing-frame signal, $h_{2,2}^{\text{CP}}(f)$, and the precession angles, (α, β, γ) . We model only the dominant ($\ell = 2, |m| = 2$) multipoles in the co-precessing frame, and neglect $\pm m$ asymmetries in the multipoles.

The co-precessing-frame model, PHENOMDCP, is an extension of the earlier aligned-spin model PHENOMD, which was calibrated to 19 NR simulations of either

single-spin or equal-spin binaries, up to $q = 18$ and spins of $|\chi| \leq 0.85$. Our extension captures the effect of in-plane spin on the amplitude and phase of the co-precessing-frame signal in the late inspiral and merger-ringdown. We note for the first time that the final black hole’s ringdown frequency is shifted to an *effective ringdown frequency* in the co-precessing frame. For this reason we explicitly model the effective ringdown frequency across the single-spin parameter space, and do not make use of estimates of the final black hole’s mass and spin.

The angle model, PHENOMANGLES, uses during inspiral the MSA PN angles used in previous PHENOM models. These angles describe the dynamics of the orbital plane of the binary, which is only approximately equal to the QA direction that we require to correctly model the signal. We find that this approximation holds well throughout the inspiral for the angles α and γ , but is not sufficiently accurate for β . However, it is possible to use higher-order PN amplitude expressions, and reduction to a single-spin subspace, to transform the MSA binary inclination into a good approximation of the QA β . This is discussed in Sec. VI B. Note that current EOBNR models also use the orbital precession dynamics as an approximation to the signal precession dynamics, and so an approach like the one used here is likely to also improve the accuracy of those models.

The PHENOMANGLES precession angles through merger and ringdown are where our model differs most significantly from previous PHENOM and EOBNR models. We observe and model a “dip” in $\alpha(f)$ (and therefore $\gamma(f)$) around the effective ringdown frequency, similar to that found in the phase derivative when constructing PHENOMD [8, 9]. Most importantly, we also model the steep collapse of β through merger. As we discuss in Sec. IX, this feature is quite distinct from the asymptotic ringdown behaviour of β , and results in a shift in the frequency location of the peak amplitude for each of the $\ell = 2$ multipoles; see Fig. 13.

Our precession model is tuned to single-spin NR simulations, but we make use of a non-bijective mapping between the six spin components required to describe a two-spin system, and the two components required in our single-spin fits to NR data; see Sec. IV. In some parts of our model, this is equivalent to using the χ_{p} parameter from earlier PHENOM models, but we also introduce modifications to produce a mapping with greater physical fidelity near $q = 1$, and in the precession dynamics we taper away two-spin oscillations as the system approaches merger. The result is an approximate IMR model for two-spin systems.

In Sec. XI we demonstrate the accuracy of PHENOMDCP, PHENOMANGLES, and the complete IMR model PHENOMP NR, by calculating matches against NR waveforms. The matches are calculated against not only the 40 calibration waveforms, but an additional 36 BAM verification waveforms from across the same single-spin parameter space, plus 27 SXS and Maya waveforms, which include two-spin systems. We find that our model

in general improves significantly over previous PHENOM and EOBNR models, as illustrated in Fig. 23.

There are several immediate directions for future work. PHENOMP NR does not model subdominant multipoles in the co-precessing frame, but these will be essential for measuring the properties of observations at larger mass ratios, which is the very region of parameter space where PHENOMP NR shows the greatest improvement over previous models. This could be achieved through directly modelling each of the multipoles, and including mode-mixing effects as in Ref. [15]. Alternatively, one could estimate the subdominant multipoles through the approximation used in Ref. [11].

Beyond this, the model needs to be extended to include explicit NR calibration to two-spin systems, and to model $\pm m$ multipole asymmetries. Our results also suggest that the angle modelling needs to be improved at lower frequencies for cases with large mass ratios, large spins, and large values of θ_{LS} ; it is possible that this will require longer NR simulations.

PHENOMP NR models the signal in a frame where the direction of the total angular momentum is constant, by first transforming the calibration NR waveforms to a frame that tracks the evolution of $\hat{\mathbf{J}}(t)$. The error incurred by this approximation is evaluated in Fig. 24 for the calibration cases used to tune the model, and we note that this error is in general at least an order of magnitude below the other sources of modelling error as measured from model mismatches except in two cases: $(q, \chi, \theta_{\text{LS}}) = (4, 0.8, 150^\circ)$ and $(q, \chi, \theta_{\text{LS}}) = (8, 0.4, 150^\circ)$. While these results largely validate the use of this assumption in constructing PHENOMP NR, it is clear from Fig. 24 that current modelling efforts are encroaching regions of parameter space where this approximation is no longer valid. In the future if we wish to construct models with mismatch errors below 10^{-4} or extend the current model to even higher mass-ratios and spin magnitudes, this approximation will need to be removed.

Finally, although most GW observations to date have been of systems with comparable masses, there has been one observation (GW190814 [62]) where the mass ratio is likely outside the calibration region of this model, and so it is necessary that the calibration region be extended to higher mass ratios. All of these areas are the subject of ongoing work.

XIII. ACKNOWLEDGEMENTS

We would like to thank other members of the Cardiff Gravity Exploration Institute who performed simulations that were used in this project: Shrobana Ghosh, Charlie Hoy, Panagiota Kolitsidou and David Yeeles.

The authors were supported in part by Science and Technology Facilities Council (STFC) grant ST/V00154X/1 and European Research Council (ERC) Consolidator Grant 647839. E. Hamilton was supported in part through the COST Action CA18108, supported

by COST (European Cooperation in Science and Technology), by Swiss National Science Foundation (SNSF) grant IZCOZ0-189876. L. London was supported at Massachusetts Institute of Technology (MIT) by National Science Foundation Grant No. PHY-1707549 as well as support from MITs School of Science and Department of Physics. A. Vano-Vinuales also thanks the PhD researcher Decree-Law no. 57/2016 of August 29 (Portugal) and Project No. UIDB/00099/2020 for support.

M. Hannam thanks Università di Roma ‘‘Sapienza’’ for hospitality while part of this work was completed.

Simulations used in this work were performed on the DiRAC@Durham facility, managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/P002293/1 and ST/R002371/1, Durham University and STFC operations grant ST/R000832/1. In addition, several of the simulations used in this work were performed as part of an allocation graciously provided by Oracle to explore the use of our code on the Oracle Cloud Infrastructure.

This research also used the supercomputing facilities at Cardiff University operated by Advanced Research Computing at Cardiff (ARCCA) on behalf of the Cardiff Supercomputing Facility and the HPC Wales and Supercomputing Wales (SCW) projects. We acknowledge the support of the latter, which is part-funded by the European Regional Development Fund (ERDF) via the Welsh Government. In part the computational resources at Cardiff University were also supported by STFC grant ST/I006285/1.

Various plots and analyses in this paper were made using the Python software packages LALSuite [59], Matplotlib [100], Numpy [101], PyCBC [102], and Scipy [90].

Appendix A: Calculation of Precession Angles

Here we outline the calculation of the coprecessing frame quantified by $\alpha(f)$, $\beta(f)$ and $\gamma(f)$; see Fig. 3. To calculate each angle we use the rotationally invariant eigenvalue method [49, 50]. As shown in Ref. [50], when multipole moments are limited to cases where $\ell = 2$ this is equivalent to the original QA method [32]. The result is independent of the initial inertial frame when the minimum rotation condition is imposed [50].

The eigenvalue method and minimal rotation convention are described in Refs. [47, 49] and [50]. Ref. [49] introduces the eigenvalue method. Ref. [47] details the practical structure of this method in its Appendix A, and Ref. [50] adds the minimal rotation convention which defines the optimal emission direction in a frame invariant way. Here we provide a self-contained description of the algorithm to calculate $\alpha(f)$, $\beta(f)$ and $\gamma(f)$.

Starting with the discrete Fourier transform of Ψ_4 decomposed into spin weight -2 spherical harmonics, $\tilde{\psi}_{\ell m}$,

we compute the effect of all pair-wise angular momentum generators averaged about the binary's centre of mass. This is $\langle \mathcal{L}_{(a)\mathcal{L}_b} \rangle$, where

$$\begin{aligned} \langle \mathcal{L}_{(a)\mathcal{L}_b} \rangle &= \frac{1}{2} \langle \mathcal{L}_a \mathcal{L}_b + \mathcal{L}_b \mathcal{L}_a \rangle \\ &= \frac{\int_{\Omega} \tilde{\Psi}_4^*(f) \mathcal{L}_{(a)\mathcal{L}_b} \tilde{\Psi}_4(f) d\Omega}{\int_{\Omega} |\tilde{\Psi}_4(f)|^2 d\Omega}, \end{aligned} \quad (\text{A1})$$

with a and b over $\{x, y, z\}$, and where

$$\begin{aligned} \mathcal{L}_x &= \frac{1}{2}(\mathcal{L}_+ + \mathcal{L}_-), \quad \mathcal{L}_y = -i\frac{1}{2}(\mathcal{L}_+ - \mathcal{L}_-), \\ \mathcal{L}_{\pm} &= e^{\pm i\varphi} [\pm i\partial_{\theta} - \cot\theta\partial_{\varphi} - is \csc\theta], \end{aligned} \quad (\text{A2})$$

and

$$\mathcal{L}_z = \partial_{\varphi}. \quad (\text{A3})$$

In Eq. (A2), s is the spin weight of the object being acted upon [103, 104]. As we are only interested in outgoing gravitational wave radiation, $s = -2$.

In practice, evaluation of Eq. (A1) need not require direct integration when $\tilde{\Psi}_4$ is written in terms of its multiple moments, $\tilde{\psi}_{\ell m}$; see Eq. (1). That is, as the operation of \mathcal{L}_{\pm} and \mathcal{L}_z on ${}_2Y_{\ell m}$ are known [73, 103], one finds that

$$\langle \mathcal{L}_{(a)\mathcal{L}_b} \rangle = \frac{1}{\sum_{\ell, m} |\tilde{\psi}_{\ell m}|^2} \begin{bmatrix} I_0 + \text{Re}(I_2) & \text{Im}I_2 & \text{Re}I_1 \\ I_0 - \text{Re}(I_2) & \text{Im}I_1 & I_{zz} \end{bmatrix} \quad (\text{A4a})$$

where

$$\begin{aligned} I_2 &\equiv \frac{1}{2}(\tilde{\Psi}, \mathcal{L}_+ \mathcal{L}_+ \tilde{\Psi}) \\ &= \frac{1}{2} \sum_{\ell, m} c_{\ell m} c_{\ell m+1} \tilde{\psi}_{\ell m+2}^* \tilde{\psi}_{\ell m} \end{aligned} \quad (\text{A4b})$$

$$\begin{aligned} I_1 &\equiv (\tilde{\Psi}, \mathcal{L}_+ (\mathcal{L}_z + 1/2) \tilde{\Psi}) \\ &= \sum_{\ell m} c_{\ell m} (m + 1/2) \tilde{\psi}_{\ell m+1}^* \tilde{\psi}_{\ell m} \end{aligned} \quad (\text{A4c})$$

$$\begin{aligned} I_0 &\equiv \frac{1}{2}(\tilde{\Psi} | \ell(\ell+1) - \mathcal{L}_z^2 | \tilde{\Psi}) \\ &= \frac{1}{2} \sum_{\ell m} [\ell(\ell+1) - m^2] |\tilde{\psi}_{\ell m}|^2 \end{aligned} \quad (\text{A4d})$$

$$I_{zz} \equiv (\tilde{\Psi}, \mathcal{L}_z \mathcal{L}_z \tilde{\Psi}) = \sum_{\ell m} m^2 |\tilde{\psi}_{\ell m}|^2 \quad (\text{A4e})$$

and where $c_{\ell m} = \sqrt{\ell(\ell+1) - m(m+1)}$.

The resulting tensor, $\langle \mathcal{L}_{(a)\mathcal{L}_b} \rangle$, is analogous to the Cauchy stress tensor in continuum mechanics, and describes infinitesimal changes in momenta (linear and angular) associated with $\tilde{\Psi}_4(f)$ averaged about the source.

From the discussion in Sec. III, we see that $\langle \mathcal{L}_{(a)\mathcal{L}_b} \rangle$ is unchanged when considering $\tilde{h}(f)$ rather than $\tilde{\Psi}_4(f)$, as

the factor of $1/2\pi f$ amounts to a simple overall rescaling that does not affect normalized eigenvectors. From these points, it follows that the eigenvector of $\langle \mathcal{L}_{(a)\mathcal{L}_b} \rangle$ with the largest eigenvalue describes the direction about the source that experiences the largest strain (or curvature) and strain-rate (or curvature-rate). This is the coprocessing frame.

If we label $\langle \mathcal{L}_{(a)\mathcal{L}_b} \rangle$'s dominant normalised eigenvector as $\hat{V} = (v_x, v_y, v_z)$, then the angles associated with the coprocessing frame are given by

$$\alpha(f) = \arctan\left(\frac{v_y(f)}{v_x(f)}\right), \quad (\text{A5})$$

$$\beta(f) = \arccos(v_z(f)), \quad (\text{A6})$$

$$\gamma(f) = -\int^f (\partial_{f'} \alpha(f')) \cos \beta(f') df'. \quad (\text{A7})$$

Eqs. (A5)-(A6) follow from the use of a source centered spherical polar coordinate system in the asymptotically flat decomposition frame. Equivalently, this is related to the frame of a distant observer. Eq. (A7) is the minimum rotation condition presented in Ref. [50], which removes secular changes in phase due to the evolution of α and β .

Appendix B: Waveforms used in analysis

The NR waveforms used in the analysis of the model are listed in Tables I and II. Table I contains the 80 waveforms which comprise the BAM catalogue [46] of single spin precessing systems up to mass ratio $q = 8$ and single-spin magnitude $\chi = 0.8$. A subset of 40 of these waveforms were also used in tuning the model. Table II lists the additional waveforms taken from the SXS [64, 98] and Maya catalogues [66] used in assessing the accuracy of the model and ensuring it was not over-fitted. This selection of waveforms includes two-spin cases.

Appendix C: Parameter-space fits

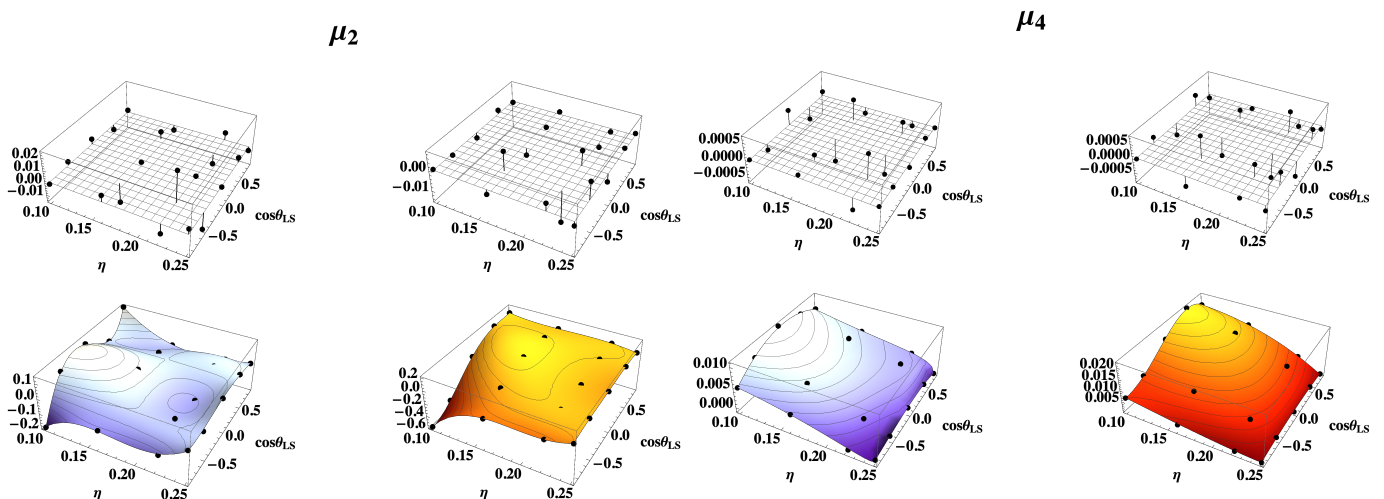
Here we show how each of the co-efficients that appear in PHENOMP NR vary across the parameter space. Figs. 25 and 26 show the variation of the co-efficients which appear in PHENOMDCP, as described in Sec. V. Fig. 27 shows the co-efficients in the ansatz for α and Fig. 28 shows those in the ansatz for β , which are presented in Sec. VII. As can be seen from these figures, the co-efficients vary smoothly across the parameter space.

Simulation ID	$\frac{100M_{\odot}}{M} f_{\text{ref}}$ (Hz)	q	χ	$\theta_{\text{LS}}(^{\circ})$	Simulation ID	$\frac{100M_{\odot}}{M} f_{\text{ref}}$ (Hz)	q	χ	$\theta_{\text{LS}}(^{\circ})$
CF21-1	14.8	1	0.2	30	CF21-41	–	–	–	–
CF21-2	14.8	1	0.2	60	CF21-42	16.0	4	0.2	60
CF21-3	–	–	–	–	CF21-43	16.8	4	0.2	90
CF21-4	–	–	–	–	CF21-44	15.3	4	0.2	120
CF21-5	14.7	1	0.2	150	CF21-45	15.2	4	0.2	150
CF21-6	14.8	1	0.4	30	CF21-46	16.6	4	0.4	30
CF21-7	14.8	1	0.4	60	CF21-47	16.3	4	0.4	60
CF21-8	14.9	1	0.4	90	CF21-48	14.7	4	0.4	90
CF21-9	14.8	1	0.4	120	CF21-49	14.8	4	0.4	120
CF21-10	14.8	1	0.4	150	CF21-50	15.0	4	0.4	150
CF21-11	18.2	1	0.6	30	CF21-51	17.0	4	0.6	30
CF21-12	14.8	1	0.6	60	CF21-52	16.2	4	0.6	60
CF21-13	14.9	1	0.6	90	CF21-53	15.8	4	0.6	90
CF21-14	14.8	1	0.6	120	CF21-54	15.1	4	0.6	120
CF21-15	14.8	1	0.6	150	CF21-55	14.0	4	0.6	150
CF21-16	14.8	1	0.8	30	CF21-56	17.5	4	0.8	30
CF21-17	14.7	1	0.8	60	CF21-57	16.8	4	0.8	60
CF21-18	14.9	1	0.8	90	CF21-58	14.9	4	0.8	90
CF21-19	14.9	1	0.8	120	CF21-59	14.8	4	0.8	120
CF21-20	15.2	1	0.8	150	CF21-60	14.9	4	0.8	150
CF21-21	14.7	2	0.2	30	CF21-61	18.4	8	0.2	30
CF21-22	14.7	2	0.2	60	CF21-62	18.1	8	0.2	60
CF21-23	14.8	2	0.2	90	CF21-63	17.8	8	0.2	90
CF21-24	15.2	2	0.2	120	CF21-64	17.3	8	0.2	120
CF21-25	14.8	2	0.2	150	CF21-65	17.2	8	0.2	150
CF21-26	14.8	2	0.4	30	CF21-66	19.0	8	0.4	30
CF21-27	14.6	2	0.4	60	CF21-67	18.6	8	0.4	60
CF21-28	14.7	2	0.4	90	CF21-68	17.8	8	0.4	90
CF21-29	14.8	2	0.4	120	CF21-69	17.0	8	0.4	120
CF21-30	14.8	2	0.4	150	CF21-70	16.5	8	0.4	150
CF21-31	14.7	2	0.6	30	CF21-71	19.7	8	0.6	30
CF21-32	14.9	2	0.6	60	CF21-72	–	–	–	–
CF21-33	14.5	2	0.6	90	CF21-73	17.9	8	0.6	90
CF21-34	14.9	2	0.6	120	CF21-74	16.7	8	0.6	120
CF21-35	14.4	2	0.6	150	CF21-75	17.0	8	0.6	150
CF21-36	14.9	2	0.8	30	CF21-76	20.5	8	0.8	30
CF21-37	14.9	2	0.8	60	CF21-77	19.5	8	0.8	60
CF21-38	14.7	2	0.8	90	CF21-78	18.0	8	0.8	90
CF21-39	15.0	2	0.8	120	CF21-79	16.0	8	0.8	120
CF21-40	15.0	2	0.8	150	CF21-80	15.2	8	0.8	150

TABLE I. BAM single-spin configurations used in tuning the co-precessing and angle models as well as in the assessment of the accuracy of the model.

Simulation ID	$\frac{100M_{\odot}}{M} f_{\text{ref}}$ (Hz)	q	χ	$\theta_{\text{LS}}(^{\circ})$	χ_1	χ_2
SXS0097	9.2	1.5	0.5	90	(-0.493, 0, 0.083)	(0, 0, 0)
SXS0018	7.9	1.5	0.5	90	(-0.494, 0, 0.078)	(0, 0, 0)
SXS0092	9.4	1.5	0.5	150	(-0.29, 0, -0.407)	(0, 0, 0)
SXS0033	11.2	3.0	0.5	30	(-0.19, 0, 0.463)	(0, 0, 0)
SXS0035	8.5	3.0	0.5	90	(-0.476, 0, 0.154)	(0, 0, 0)
SXS1109	10.2	5.0	0.5	90	(-0.435, 0, 0.246)	(0, 0, 0)
SXS0062	14.1	5.0	0.5	116	(-0.492, 0, 0.088)	(0, 0, 0)
SXS0161	9.2	1.0	1.199	120	(-0.579, 0, -0.158)	(-0.579, 0, -0.158)
SXS0115	10.2	1.07	0.246	74	(-0.027, -0.016, -0.203)	(-0.236, 0.018, 0.304)
SXS0116	10.2	1.08	0.167	40	(0.022, -0.099, 0.032)	(-0.143, 0.115, 0.106)
SXS0124	10.2	1.26	0.412	44	(-0.247, -0.041, 0.091)	(-0.079, 0.065, 0.294)
SXS0102	9.3	1.5	0.5	90	(-0.486, 0, 0.116)	(-0.486, 0, 0.116)
SXS1397	5.1	1.56	0.299	111	(-0.242, 0.037, -0.172)	(0.458, -0.089, 0.102)
SXS0135	10.3	1.64	0.186	128	(-0.059, 0.095, 0.025)	(0.003, -0.257, -0.228)
SXS0143	10.2	1.92	0.441	28	(-0.072, 0.041, 0.443)	(-0.413, -0.15, -0.056)
SXS0144	10.2	1.94	0.214	146	(-0.135, 0.011, -0.281)	(0.05, -0.04, 0.213)
SXS0049	11.2	3.0	0.527	72	(-0.474, 0, 0.159)	(0.159, 0, 0.474)
SXS1160	11.1	3.0	0.658	63	(-0.455, -0.024, 0.531)	(0.41, 0.212, -0.384)
SXS0165	18.1	6.0	0.93	125	(-0.648, 0.003, 0.639)	(-0.186, -0.094, -0.216)
GT0745	22.3	6.0	0.6	91	(-0.437, 0, 0.411)	(0, 0, 0)
GT0742	21.9	7.0	0.6	91	(-0.404, 0, 0.444)	(0, 0, 0)
GT0834	20.9	7.0	0.8	168	(-0.375, 0, 0.706)	(0, 0, 0)
GT0880	22.0	4.5	0.537	52	(-0.269, 0, 0.536)	(0.269, 0, -0.536)
GT0887	21.8	5.0	0.543	51	(-0.256, 0, 0.542)	(0.256, 0, -0.542)
GT0889	21.6	6.0	0.552	50	(-0.234, 0, 0.552)	(0.234, 0, -0.552)
GT0888	21.4	7.0	0.559	49	(-0.216, 0, 0.56)	(0.216, 0, -0.56)
GT0886	21.4	8.0	0.564	49	(-0.2, 0, 0.566)	(0.2, 0, -0.566)

TABLE II. Additional configurations from the SXS and MAYA catalogues used in the assessment of the accuracy of the model.

FIG. 25. Amplitude parameters for tuned co-precessing waveform model, PHENOMDCP. The fits are shown as two-dimensional surfaces covering the parameter space described by η and $\cos\theta_{\text{LS}}$. On the left in blue are the fits for the simulations with $\chi = 0.4$ and on the right in red are the fits for $\chi = 0.8$. Above each of these surfaces are shown the residuals.

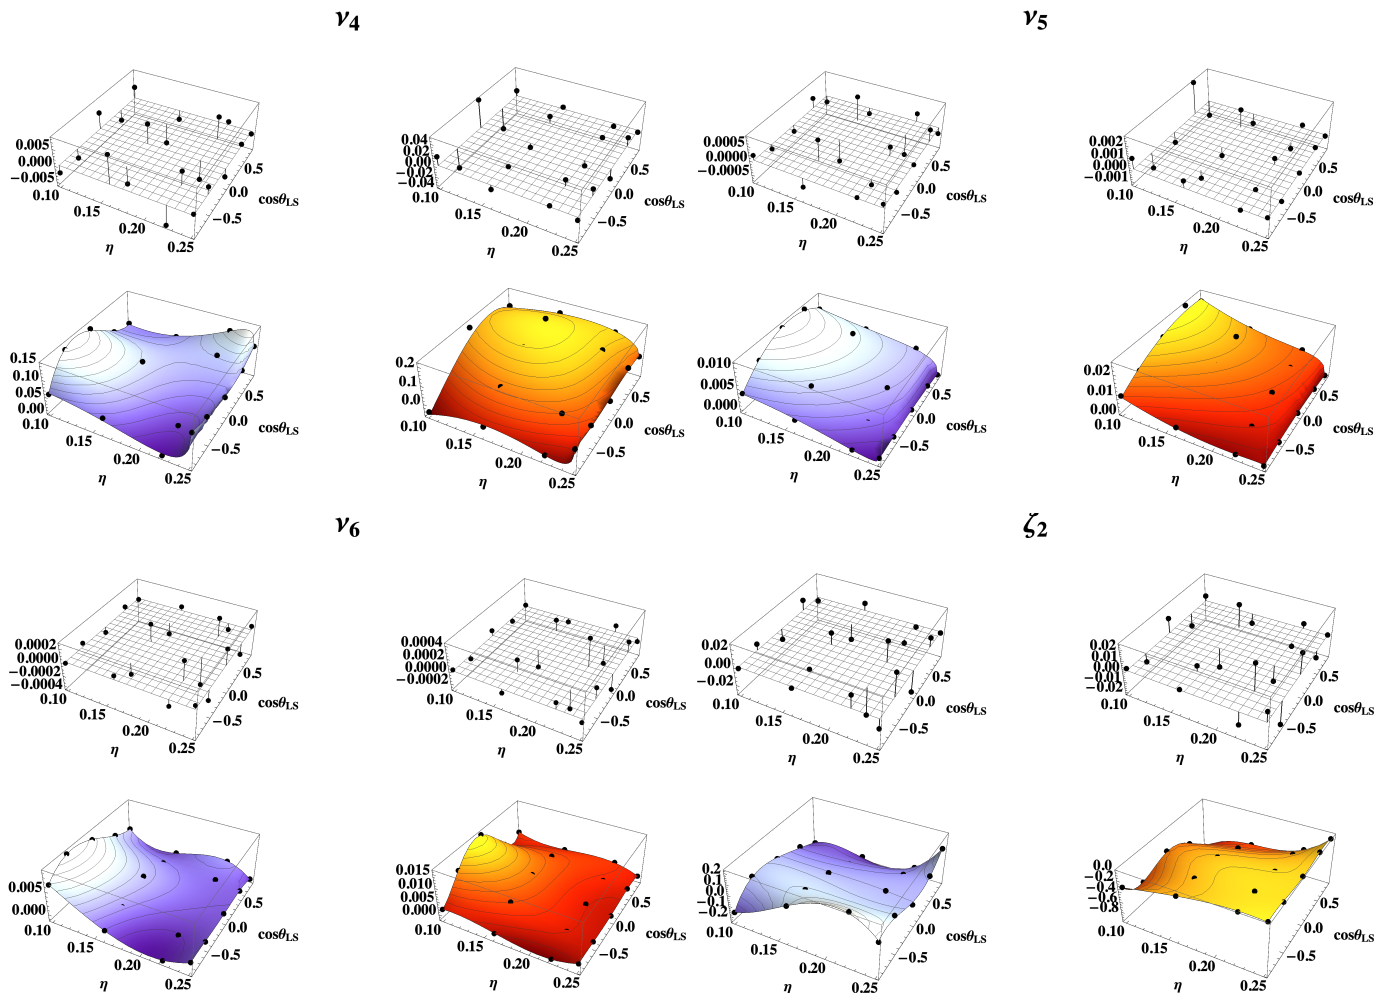


FIG. 26. Phase parameters for tuned co-precessing waveform model, PHENOMDCP. The fits are shown as two-dimensional surfaces covering the parameter space described by η and $\cos\theta_{LS}$. On the left in blue are the fits for the simulations with $\chi = 0.4$ and on the right in red are the fits for $\chi = 0.8$. Above each of these surfaces are shown the residuals.

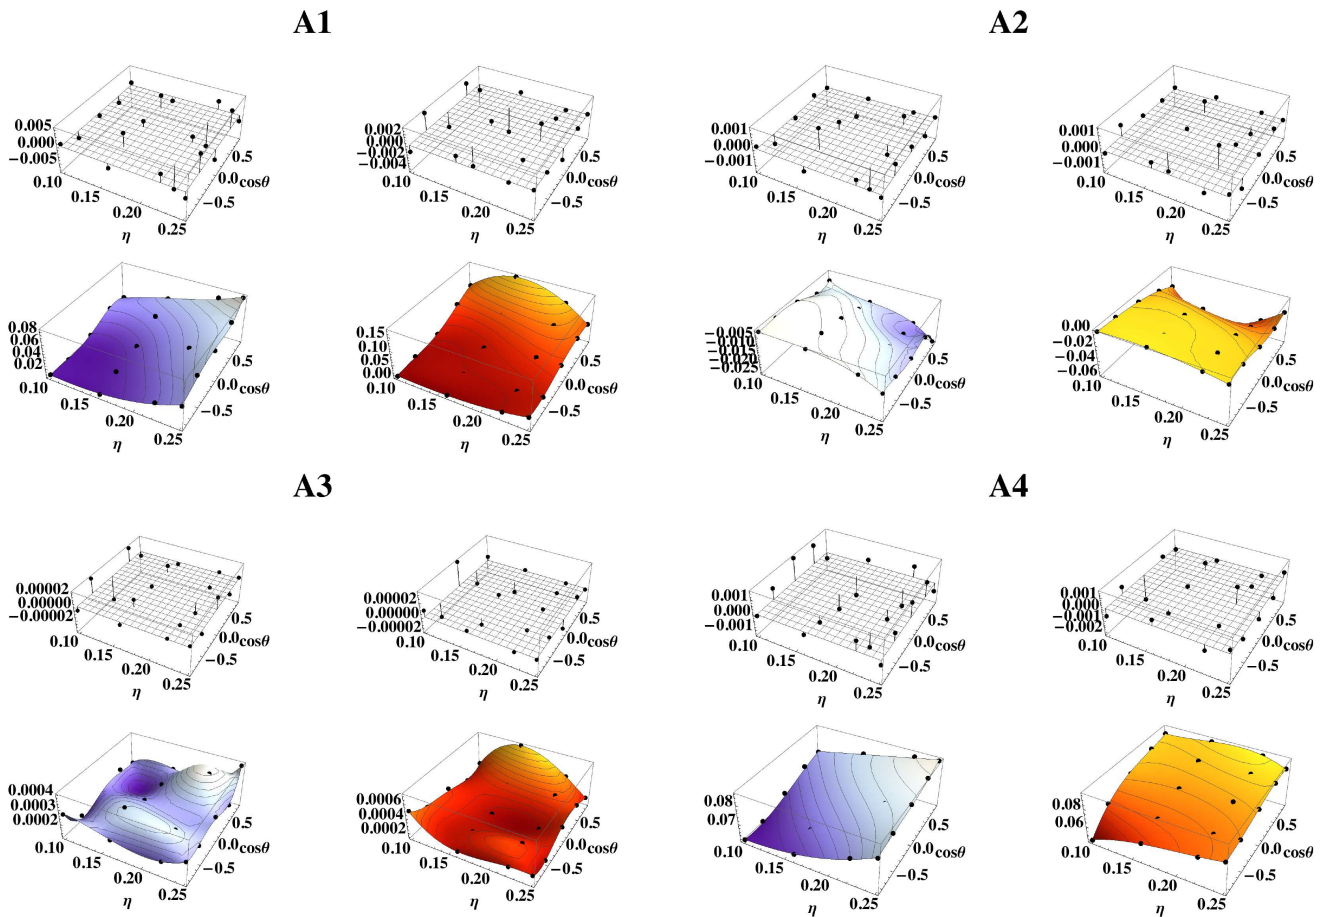


FIG. 27. Comparison of the fits for each of the co-efficients for the ansatz for α given in equation 50 with the co-efficients found from the data as described in Sec. VII. The fits are shown as two-dimensional surfaces covering the parameter space described by η and $\cos\theta_{LS}$. On the left in blue are the fits for the simulations with $\chi = 0.4$ and on the right in red are the fits for $\chi = 0.8$. Above each of these surfaces are shown the residuals.

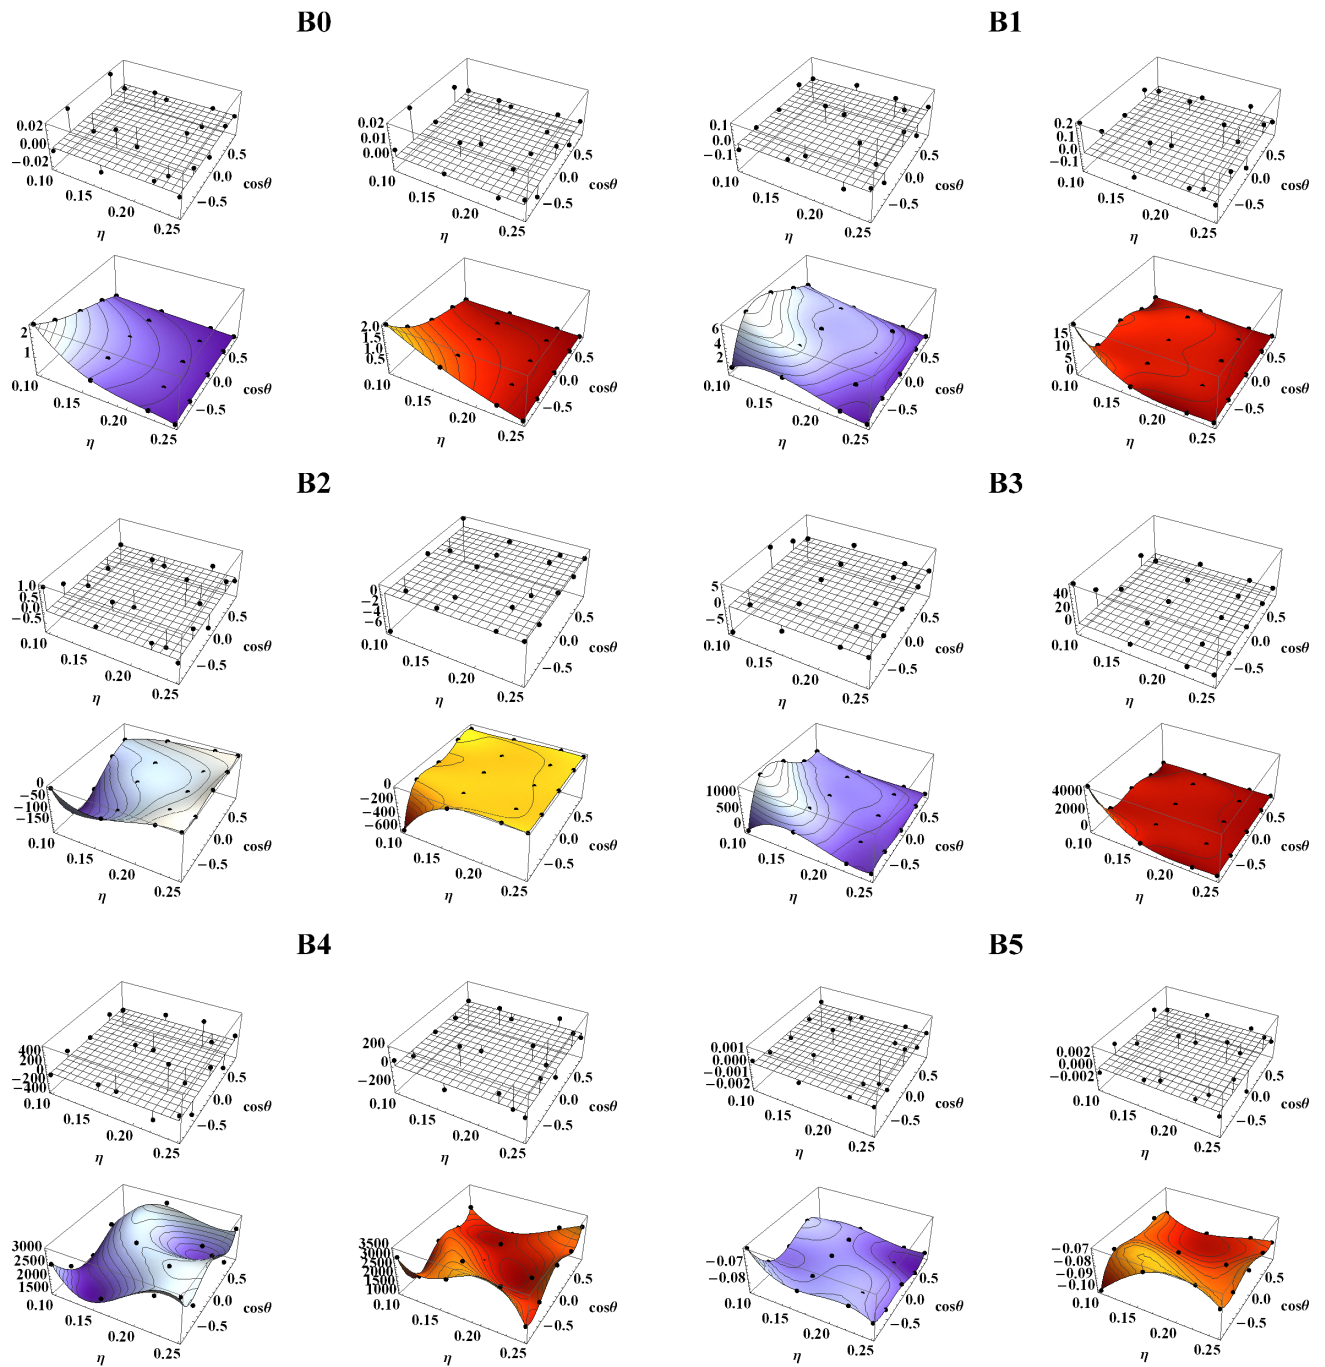


FIG. 28. Comparison of the fits for each of the co-efficients for the ansatz for β given in equation 51 with the co-efficients found from the data as described in Sec. VII. The fits are shown as two-dimensional surfaces covering the parameter space described by η and $\cos\theta_{\text{LS}}$. On the left in blue are the fits for the simulations with $\chi = 0.4$ and on the right in red are the fits for $\chi = 0.8$. Above each of these surfaces are shown the residuals.

- [1] J. Aasi *et al.* (LIGO Scientific), *Class. Quant. Grav.* **32**, 074001 (2015), [arXiv:1411.4547 \[gr-qc\]](#).
- [2] F. Acernese *et al.* (VIRGO), *Class. Quant. Grav.* **32**, 024001 (2015), [arXiv:1408.3978 \[gr-qc\]](#).
- [3] B. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. X* **9**, 031040 (2019), [arXiv:1811.12907 \[astro-ph.HE\]](#).
- [4] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. hARRY, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, *The Astrophysical Journal* **891**, 123 (2020).
- [5] B. Zackay, L. Dai, T. Venumadhav, J. Roulet, and M. Zaldarriaga, (2019), [arXiv:1910.09528 \[astro-ph.HE\]](#).
- [6] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **101**, 083030 (2020), [arXiv:1904.07214 \[astro-ph.HE\]](#).
- [7] R. Abbott *et al.* (LIGO Scientific, Virgo), (2020), [arXiv:2010.14527 \[gr-qc\]](#).
- [8] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. Jimnez Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016), [arXiv:1508.07250 \[gr-qc\]](#).
- [9] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jimnez Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044007 (2016), [arXiv:1508.07253 \[gr-qc\]](#).
- [10] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014), [arXiv:1308.3271 \[gr-qc\]](#).
- [11] L. London, S. Khan, E. Fauchon-Jones, C. Garca, M. Hannam, S. Husa, X. Jiménez-Forteza, C. Kalaghatgi, F. Ohme, and F. Pannarale, *Phys. Rev. Lett.* **120**, 161102 (2018), [arXiv:1708.00404 \[gr-qc\]](#).
- [12] S. Khan, K. Chatziioannou, M. Hannam, and F. Ohme, *Phys. Rev. D* **100**, 024059 (2019), [arXiv:1809.10113 \[gr-qc\]](#).
- [13] S. Khan, F. Ohme, K. Chatziioannou, and M. Hannam, *Phys. Rev. D* **101**, 024056 (2020), [arXiv:1911.06050 \[gr-qc\]](#).
- [14] G. Pratten, S. Husa, C. Garcia-Quiros, M. Colleoni, A. Ramos-Buades, H. Estelles, and R. Jaume, (2020), [arXiv:2001.11412 \[gr-qc\]](#).
- [15] C. Garcia-Quirós, M. Colleoni, S. Husa, H. Estellés, G. Pratten, A. Ramos-Buades, M. Mateu-Lucena, and R. Jaume, (2020), [arXiv:2001.10914 \[gr-qc\]](#).
- [16] G. Pratten *et al.*, (2020), [arXiv:2004.06503 \[gr-qc\]](#).
- [17] J. E. Thompson, E. Fauchon-Jones, S. Khan, E. Nitz, F. Pannarale, T. Dietrich, and M. Hannam, *Phys. Rev. D* **101**, 124059 (2020), [arXiv:2002.08383 \[gr-qc\]](#).
- [18] H. Estellés, A. Ramos-Buades, S. Husa, C. García-Quirós, M. Colleoni, L. Haegel, and R. Jaume, (2020), [arXiv:2004.08302 \[gr-qc\]](#).
- [19] H. Estellés, S. Husa, M. Colleoni, D. Keitel, M. Mateu-Lucena, C. García-Quirós, A. Ramos-Buades, and A. Borchers, (2020), [arXiv:2012.11923 \[gr-qc\]](#).
- [20] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, T. Chu, G. Lovelace, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **86**, 024011 (2012), [arXiv:1202.0790 \[gr-qc\]](#).
- [21] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **89**, 084006 (2014), [arXiv:1307.6232 \[gr-qc\]](#).
- [22] A. Taracchini *et al.*, *Phys. Rev. D* **89**, 061502 (2014), [arXiv:1311.2544 \[gr-qc\]](#).
- [23] A. Bohé *et al.*, *Phys. Rev. D* **95**, 044028 (2017), [arXiv:1611.03703 \[gr-qc\]](#).
- [24] R. Cotesta, A. Buonanno, A. Bohé, A. Taracchini, I. Hinder, and S. Ossokine, *Phys. Rev. D* **98**, 084028 (2018), [arXiv:1803.10701 \[gr-qc\]](#).
- [25] S. Ossokine *et al.*, *Phys. Rev. D* **102**, 044055 (2020), [arXiv:2004.09442 \[gr-qc\]](#).
- [26] A. Matas *et al.*, *Phys. Rev. D* **102**, 043023 (2020), [arXiv:2004.10001 \[gr-qc\]](#).
- [27] A. Ramos-Buades, P. Schmidt, G. Pratten, and S. Husa, *Phys. Rev. D* **101**, 103014 (2020), [arXiv:2001.10936 \[gr-qc\]](#).
- [28] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, D. A. Hemberger, P. Schmidt, and R. Smith, *Phys. Rev. D* **95**, 104023 (2017), [arXiv:1701.00550 \[gr-qc\]](#).
- [29] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, *Phys. Rev. D* **99**, 064045 (2019), [arXiv:1812.07865 \[gr-qc\]](#).
- [30] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, *Phys. Rev. D* **49**, 6274 (1994).
- [31] L. E. Kidder, *Phys. Rev. D* **52**, 821 (1995), [arXiv:gr-qc/9506022](#).
- [32] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, *Phys. Rev. D* **84**, 024046 (2011), [arXiv:1012.2879 \[gr-qc\]](#).
- [33] P. Schmidt, M. Hannam, and S. Husa, *Phys. Rev. D* **86**, 104063 (2012), [arXiv:1207.3088 \[gr-qc\]](#).
- [34] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, C. D. Ott, M. Boyle, L. E. Kidder, H. P. Pfeiffer, and B. Szilágyi, *Phys. Rev. D* **96**, 024058 (2017), [arXiv:1705.07089 \[gr-qc\]](#).
- [35] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, *Phys. Rev. Research* **1**, 033015 (2019), [arXiv:1905.09300 \[gr-qc\]](#).
- [36] P. Ajith *et al.*, *Class. Quant. Grav.* **24**, S689 (2007), [arXiv:0704.3764 \[gr-qc\]](#).
- [37] P. Ajith *et al.*, *Phys. Rev. D* **77**, 104017 (2008), [Erratum: *Phys. Rev. D* **79**, 129901 (2009)], [arXiv:0710.2335 \[gr-qc\]](#).
- [38] P. Ajith *et al.*, *Phys. Rev. Lett.* **106**, 241101 (2011), [arXiv:0909.2867 \[gr-qc\]](#).
- [39] L. Santamaria *et al.*, *Phys. Rev. D* **82**, 064016 (2010), [arXiv:1005.3306 \[gr-qc\]](#).
- [40] M. Pürrer, M. Hannam, and F. Ohme, *Phys. Rev. D* **93**, 084042 (2016), [arXiv:1512.04955 \[gr-qc\]](#).
- [41] B. Abbott *et al.* (KAGRA, LIGO Scientific, Virgo), *Living Rev. Rel.* **23**, 3 (2020).
- [42] S. Fairhurst, R. Green, M. Hannam, and C. Hoy, *Phys. Rev. D* **102**, 041302 (2020), [arXiv:1908.00555 \[gr-qc\]](#).
- [43] S. Vitale, R. Lynch, V. Raymond, R. Sturani, J. Veitch, and P. Graff, *Phys. Rev. D* **95**, 064053 (2017), [arXiv:1611.01122 \[gr-qc\]](#).
- [44] R. Green, C. Hoy, S. Fairhurst, M. Hannam, F. Pannarale, and C. Thomas, *Phys. Rev. D* **103**, 124023 (2021), [arXiv:2010.04131 \[gr-qc\]](#).
- [45] H. Estellés, M. Colleoni, C. García-Quirós, S. Husa, D. Keitel, M. Mateu-Lucena, M. d. L. Planas, and A. Ramos-Buades, (2021), [arXiv:2105.05872 \[gr-qc\]](#).

- [46] E. Fauchon-Jones, E. Hamilton, C. Kalaghatgi, S. Ghosh, M. Hannam, C. Hoy, S. Khan, P. Kolitsidou, L. London, J. Thompson, A. Vañó-Viñuales, and D. Yeeles, (2021), in preparation.
- [47] L. Pekowsky, R. O’Shaughnessy, J. Healy, and D. Shoemaker, *Phys. Rev. D* **88**, 024040 (2013), [arXiv:1304.3176 \[gr-qc\]](#).
- [48] F. Hofmann, E. Barausse, and L. Rezzolla, *Astrophys. J. Lett.* **825**, L19 (2016), [arXiv:1605.01938 \[gr-qc\]](#).
- [49] R. O’Shaughnessy, B. Vaishnav, J. Healy, Z. Meeks, and D. Shoemaker, *Phys. Rev.* **D84**, 124002 (2011), [arXiv:1109.5224 \[gr-qc\]](#).
- [50] M. Boyle, R. Owen, and H. P. Pfeiffer, *Phys. Rev.* **D84**, 124011 (2011), [arXiv:1110.2965 \[gr-qc\]](#).
- [51] E. Ochsner and R. O’Shaughnessy, *Phys. Rev.* **D86**, 104037 (2012), [arXiv:1205.2287 \[gr-qc\]](#).
- [52] M. Boyle, L. E. Kidder, S. Ossokine, and H. P. Pfeiffer, (2014), [arXiv:1409.4431 \[gr-qc\]](#).
- [53] E. Hamilton and M. Hannam, *Phys. Rev.* **D98**, 084018 (2018), [arXiv:1807.06331 \[gr-qc\]](#).
- [54] K. Chatziioannou, A. Klein, N. Cornish, and N. Yunes, *Phys. Rev. Lett.* **118**, 051101 (2017), [arXiv:1606.03117 \[gr-qc\]](#).
- [55] R. O’Shaughnessy, L. London, J. Healy, and D. Shoemaker, *Phys. Rev. D* **87**, 044038 (2013), [arXiv:1209.3712 \[gr-qc\]](#).
- [56] S. Marsat and J. G. Baker, (2018), [arXiv:1806.10734 \[gr-qc\]](#).
- [57] B. Bruegmann, J. A. Gonzalez, M. Hannam, S. Husa, and U. Sperhake, *Phys. Rev.* **D77**, 124047 (2008), [arXiv:0707.0135 \[gr-qc\]](#).
- [58] C. Kalaghatgi and M. Hannam, (2020), [arXiv:2008.09957 \[gr-qc\]](#).
- [59] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite,” free software (GPL) (2018).
- [60] A. Buonanno, Y.-b. Chen, Y. Pan, and M. Vallisneri, *Phys. Rev. D* **70**, 104003 (2004), [Erratum: *Phys. Rev. D* **74**, 029902 (2006)], [arXiv:gr-qc/0405090](#).
- [61] P. Schmidt, F. Ohme, and M. Hannam, *Phys. Rev. D* **91**, 024043 (2015), [arXiv:1408.1810 \[gr-qc\]](#).
- [62] R. Abbott *et al.* (LIGO Scientific, Virgo), *Astrophys. J. Lett.* **896**, L44 (2020), [arXiv:2006.12611 \[astro-ph.HE\]](#).
- [63] B. Bruegmann, J. A. Gonzalez, M. Hannam, S. Husa, U. Sperhake, and W. Tichy, *Phys. Rev. D* **77**, 024027 (2008), [arXiv:gr-qc/0610128](#).
- [64] M. Boyle *et al.*, *Class. Quant. Grav.* **36**, 195006 (2019), [arXiv:1904.04831 \[gr-qc\]](#).
- [65] <http://www.black-holes.org>.
- [66] K. Jani, J. Healy, J. A. Clark, L. London, P. Laguna, and D. Shoemaker, *Class. Quant. Grav.* **33**, 204001 (2016), [arXiv:1605.03204 \[gr-qc\]](#).
- [67] <http://www.einstein.gatech.edu/catalog/>.
- [68] M. Campanelli, C. Lousto, and Y. Zlochower, *Phys. Rev. D* **74**, 041501 (2006), [arXiv:gr-qc/0604012](#).
- [69] G. B. Cook and J. York, James W., *Phys. Rev. D* **41**, 1077 (1990).
- [70] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-Time Signal Processing (2nd Ed.)* (Prentice-Hall, Inc., USA, 1999).
- [71] E. P. Wigner, *Group theory and its application to the quantum mechanics of atomic spectra* (Academic Press, New York, 1959).
- [72] J. M. Bowen and J. W. York, Jr., *Phys. Rev. D* **21**, 2047 (1980).
- [73] M. Ruiz, M. Alcubierre, D. Núñez, and R. Takahashi, *General Relativity and Gravitation* **40**, 1705 (2008).
- [74] M. Campanelli, C. O. Lousto, Y. Zlochower, B. Krishnan, and D. Merritt, *Phys. Rev. D* **75**, 064030 (2007), [arXiv:gr-qc/0612076](#).
- [75] V. Varma, M. Isi, and S. Biscoveanu, *Phys. Rev. Lett.* **124**, 101104 (2020), [arXiv:2002.00296 \[gr-qc\]](#).
- [76] E. Poisson and C. M. Will, *Phys. Rev. D* **52**, 848 (1995), [arXiv:gr-qc/9502040](#).
- [77] C. Cutler and E. E. Flanagan, *Phys. Rev.* **D49**, 2658 (1994), [arXiv:gr-qc/9402014 \[gr-qc\]](#).
- [78] L. M. Thomas, P. Schmidt, and G. Pratten, (2020), [arXiv:2012.02209 \[gr-qc\]](#).
- [79] D. Gerosa, M. Mould, D. Gangardt, P. Schmidt, G. Pratten, and L. M. Thomas, (2020), [arXiv:2011.11948 \[gr-qc\]](#).
- [80] P. Ajith, *Phys. Rev. D* **84**, 084037 (2011), [arXiv:1107.1267 \[gr-qc\]](#).
- [81] P. Kumar, T. Chu, H. Fong, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, M. A. Scheel, and B. Szilagy, *Phys. Rev. D* **93**, 104050 (2016), [arXiv:1601.05396 \[gr-qc\]](#).
- [82] K. Chatziioannou, A. Klein, N. Yunes, and N. Cornish, *Phys. Rev.* **D95**, 104004 (2017), [arXiv:1703.03967 \[gr-qc\]](#).
- [83] E. Racine, *Phys. Rev. D* **78**, 044021 (2008), [arXiv:0803.1820 \[gr-qc\]](#).
- [84] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner, *Phys. Rev.* **D79**, 104023 (2009), [Erratum: *Phys. Rev. D* **84**, 049901 (2011)], [arXiv:0810.5336 \[gr-qc\]](#).
- [85] L. London and E. Fauchon-Jones, *Class. Quant. Grav.* **36**, 235015 (2019), [arXiv:1810.03550 \[gr-qc\]](#).
- [86] L. London, E. Fauchon, and EZHamilton, “llondon6/positive: map,” (2020).
- [87] M. Kesden, D. Gerosa, R. O’Shaughnessy, E. Berti, and U. Sperhake, *Phys. Rev. Lett.* **114**, 081103 (2015), [arXiv:1411.0674 \[gr-qc\]](#).
- [88] A. Klein, N. Cornish, and N. Yunes, *Phys. Rev. D* **88**, 124015 (2013), [arXiv:1305.1932 \[gr-qc\]](#).
- [89] K. Chatziioannou, A. Klein, N. Yunes, and N. Cornish, *Phys. Rev. D* **88**, 063011 (2013), [arXiv:1307.4418 \[gr-qc\]](#).
- [90] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, *Nature Methods* **17**, 261 (2020).
- [91] E. Leaver, *Proc. Roy. Soc. Lond. A* **A402**, 285 (1985).
- [92] S. A. Teukolsky, *Astrophys. J.* **185**, 635 (1973).
- [93] I. Kamaretsos, M. Hannam, and B. Sathyaprakash, *Phys. Rev. Lett.* **109**, 141102 (2012), [arXiv:1207.0399 \[gr-qc\]](#).
- [94] I. Kamaretsos, M. Hannam, S. Husa, and B. S. Sathyaprakash, *Phys. Rev. D* **85**, 024018 (2012), [arXiv:1107.0854 \[gr-qc\]](#).
- [95] M. Giesler, M. Isi, M. A. Scheel, and S. Teukolsky, *Phys. Rev. X* **9**, 041060 (2019), [arXiv:1903.08284 \[gr-qc\]](#).

- [96] LSC, “LIGO Document T1800044-v5,” <https://dcc.ligo.org/LIGO-T1800044/public>.
- [97] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, *Phys. Rev. D* **94**, 024012 (2016), [arXiv:1603.02444](https://arxiv.org/abs/1603.02444) [gr-qc].
- [98] A. H. Mroue *et al.*, *Phys. Rev. Lett.* **111**, 241104 (2013), [arXiv:1304.6077](https://arxiv.org/abs/1304.6077) [gr-qc].
- [99] P. Schmidt, I. W. Harry, and H. P. Pfeiffer, (2017), [arXiv:1703.01076](https://arxiv.org/abs/1703.01076) [gr-qc].
- [100] J. D. Hunter, *Computing in Science & Engineering* **9**, 90 (2007).
- [101] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, *Nature* **585**, 357 (2020).
- [102] A. Nitz, I. Harry, D. Brown, C. M. Biwer, J. Willis, T. D. Canton, C. Capano, T. Dent, L. Pekowsky, A. R. Williamson, G. S. C. Davies, S. De, M. Cabero, B. Machenschalk, P. Kumar, D. Macleod, S. Reyes, dfinstad, F. Pannarale, T. Massinger, S. Kumar, M. Tápai, L. Singer, S. Khan, S. Fairhurst, A. Nielsen, S. Singh, shasvath, B. U. V. Gadre, and I. Dorrington, “gwastro/pycbc: Pycbc release 1.18.1,” (2021).
- [103] E. T. Newman and R. Penrose, *Journal of Mathematical Physics* **7**, 863 (1966), <https://doi.org/10.1063/1.1931221>.
- [104] A. G. Shah and B. F. Whiting, *Gen. Rel. Grav.* **48**, 78 (2016), [arXiv:1503.02618](https://arxiv.org/abs/1503.02618) [gr-qc].