Measuring Spin Precession in the Ringdown

Interim Report II, LIGO Summer Undergraduate Research Fellowship, California Institute of Technology

Serena Fink
University of Montana, Missoula, MT 59801

Mentors: Simona Miller and Eliot Finch
LIGO Laboratory, Caltech, Pasadena, CA 91125

(Dated: August 4, 2025)

The effect of progenitor masses and spins on the inspiral phase of a gravitational wave signal from binary black hole coalescence is relatively well understood for the case where spins are aligned with the orbital angular momentum. Less is known about how progenitor properties – especially misaligned spins – impact the ringdown. The accuracy and precision with which inspiral parameters (component masses and spins) can be inferred from the ringdown alone remains an open question. As several high-mass (most ringdown-dominated) systems observed by LIGO have misaligned spins, the ability to predict inspiral properties from the ringdown would improve existing measurements, and potentially give new insight into this portion of the binary black hole population. We introduce a least-squares fitting method for inverting ringdown surrogate models, allowing us to map from measured quasinormal modes back to inspiral parameters. This method varies in quality depending on the amplitude surrogate. The surrogate that has been explored the most in this project, NRSur3dq8_RD, reasonably estimates parameters for systems with near-equal masses and smaller spin magnitudes, but provides less accurate estimates in systems with highly unequal masses and larger spins. To better assess the uncertainty and correlational structure in our parameter space, we additionally perform parameter estimation using ringdown surrogate models.

I. INTRODUCTION

A binary black hole (BBH) coalescence can be broken into three parts: the inspiral, the merger, and the ringdown. The inspiral period occurs when the two black holes (BHs) are orbiting each other and emitting lower frequency gravitational waves (GWs) as they get closer together, and then eventually collide in the merger. Most BBHs thus far observed by LIGO-Virgo-KAGRA (the LVK) are inspiral-dominated. However, for high-mass systems, the inspiral is difficult to observe with ground-based interferometers because the majority of GWs emitted before the merger are at a lower frequency than their sensitivity band [1]. Because the inspiral period is difficult to probe directly for these systems, it is useful to be able to infer the properties of the progenitor BHs (i.e., masses and spins) based on the frequencies observed in the more readily visible ringdown. The ringdown is the period directly after the merger in which the two BHs have now coalesced into a single remnant BH, and the remnant rotates asymmetrically about its axis. This perturbed BH remnant still emits GWs as it equilibrates to its final stage, a Kerr BH [2].

The ringdown portion of a GW signal consists of a sum of damped oscillations known as quasinormal modes (QNMs) [3–7], which are emitted from the perturbed BH at fundamental and overtone frequencies. "Quasi" refers to the fact that the modes decay characteristically over time, whereas normal modes do not. The longest lasting of these are referred to as the fundamentals, while any shorter lived modes are the overtones. The observed ringdown signal from a BBH merger is a sum of all the QNMs with some relative amplitudes and phases. In this project, we measure QNMs from simulated GW signals to see if our predictions are consistent with the true parameters in the simulation.

We simulate full inspiral-merger-ringdown (IMR) signals using numerical relativity (NR) surrogate waveform models [8]. Then, using the QNM model, we dissect the ringdown portion of the simulated waveform into its individual QNM components. These QNMs are then inverted using a ringdown surrogate, which maps between modes and remnant properties to the progenitor masses and spins (which we call "inspiral parameters"). If the QNM model predicts inspiral parameters that are consistent with the injected values used to generated the IMR waveform, this method will then be applied to real GW data to estimate inspiral parameters that cannot be observed directly.

Presently, inferring inspiral parameters from ringdown analysis is best understood when the progenitor BHs have spins aligned with the binary's orbital angular momentum. For the aligned-spin case, the relationships between QNM amplitudes and phases and the masses/spins of the inspiraling black holes have been encoded in various fitting formulas or ringdown surrogate models [9–12]. However, in reality, most BHs possess at least some degree of spin-orbit misalignment [13], leading to spin-orbit precession. Since applying aligned-spin models to systems with precession may bias the recovery of system parameters, it is important to understand how precession affects the ringdown. Currently, the mapping from QNM am-

plitudes to inspiral properties is less clear for precessing binaries, although a relationship is expected [2]. The precessing case is particularly of interest because of its potential to give astrophysical insight into how a BBH system formed [e.g., 14, 15]. By studying if and how the properties of the individual progenitor BHs in precessing systems can be predicted and inferred using just their ringdown data, the population of BHs that undergo a merger can be better examined.

The need to understand the relationship between the inspiral and ringdown properties for high-mass (e.g. ringdown-dominated) precessing systems is further motivated by the fact that several of the BBHs observed by the LVK are indeed high-mass systems with intriguing spins. The new event GW231123 [16] is the most massive BBH detected yet by a factor of ~ 2, and is rapidly spinning–but the specifics of its inferred spin configuration are highly dependent on the specific waveform model used. GW190521, the second most-massive BBH system observed by the LVK, is highly precessing [17]; its precession measurement stems from the subtle interplay between the quiet final inspiral cycle and the loud merger [18]. GW191109 and GW200129 are other high-mass systems with interesting spin configurations, but are plagued by data-quality issues: the anti-aligned spin of GW191109 spins [19] and the precession in GW200129 [20] are both degenerate with glitches coincident with their inspiral data. For all of these events, the information about spins is coming primarily from their inspiral, which is short, quiet and/or affected by data quality issues or waveform systematics. Having ways to infer the inspiral properties *without* looking at this finicky inspiral itself serves as motivation for our study.

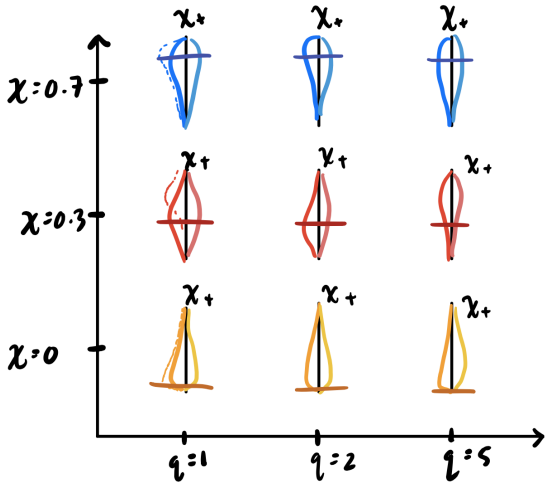

FIG. 1: Sketch of a violin plot showing posterior distributions on the inspiral spin parameter χ_+ (also known as $\chi_{\rm eff}$) [12]. In this sketch, the lighter color represents the *reconstructed* χ_+ parameter using the QNM fits; the darker color is χ_+ *directly inferred* with the NR surrogate when the full (dashed) or just ringdown (solid) data are analyzed. Each violin is from a simulated signal with a different primary spin $\chi_1 = \chi$ (vertical axis) and mass ratio q (horizontal axis). In practice, a version of this plot would be generated for each inspiral parameter we are measuring, for both the aligned and precessing cases.

II. OBJECTIVES

The purpose of this project is to test how well QNM models in the ringdown can be used to back-infer properties of the inspiral. We start with the aligned-spin case, where the mapping from QNM amplitudes to inspiral properties is comparatively well understood [9–12], and then move to the precessing case. We begin by analyzing simulated signals (for which we know the inspiral properties exactly); from these we can test the validity of our mappings between QNMs and inspiral properties before moving on to apply the method to real GW data. Our objectives are as follows:

1. Generate a mapping between QNMs and inspiral properties. Use this mapping in *least-squares fitting* to get point-estimate fits for inspiral properties of a simulated IMR signal given only that signal's ringdown.

2. Examine how the accuracy of the inferred inspiral parameters changes for different simulated signals, e.g., equal vs. unequal masses, large vs. small spins aligned vs. precessing spins. Determine how many and which QNMs are needed to make the most accurate fit for the different simulated signals.

3. Repeat the above analyses, but instead of using point estimates, perform *parameter estimation* (PE) to obtain posterior distributions on QNMs and inspiral properties inferred with a ringdown surrogate. Compare these posterior distributions to those obtained with traditional full IMR PE [21]. Figure 1 shows a sketch of what these posteriors might look like.

4. Apply our method to a real GW signal, like GW231123 or GW190521.

While we aspire to investigate the full spin degrees of freedom, mappings for the QNM amplitudes in the precessing case are not well understood, although see Refs. [2, 22] for work in this direction. Consequently, a simpler question we could ask is when do the aligned-spin mappings break down (say, as a function of SNR and how much precession there is).

III. APPROACH AND RESULTS

So far in this project, we have been working on developing a method for mapping between QNMs and inspiral

properties (see Point 1 in Sec.II). The method consists of utilizing NR surrogate waveforms and QNM models to compare ringdown signals and determine if the QNM model can accurately infer the inspiral parameters. Section III A outlines the process for testing various ringdown surrogates for the aligned spin case. After finding surrogates that agree well with the QNM model, we inverted these surrogates to return inspiral parameters. This process is discussed in Sec. III B. A summary of the surrogates and their associated code-bases is given in Table I.

Surrogate/Model	Code Base	Description	Ref.
–	`qnmfits`	Fit QNMs with least squares	[23]
–	`Jaxqualin`	QNM amplitude fits based on polynomials	[24]
NRSur3dq8_RD	`surfinBH`	QNM amplitude fits based on Gaussian processes	[25]
NRHybSur3dq8	`gwsurrogate`	IMR waveform surrogate	[26]
q8_3dAl	`qnmpredictor`	QNM amplitude fits based on Gaussian processes	[27]

TABLE I: Summary of the surrogates and models discussed in this paper.

The following is an overview of our method. The NR surrogate NRHybSur3dq8 [26] yields a complete IMR waveform. Then `qnmfits` is implemented to make a fit for just the ringdown of the NRHybSur3dq8 waveform. By inverting an additional *remnant* NR surrogate (also called a "ringdown surrogate") NRSur3dq8_RD with `surfinBH`, the amplitude and phase observed in the ringdown are mapped back to inspiral parameters.

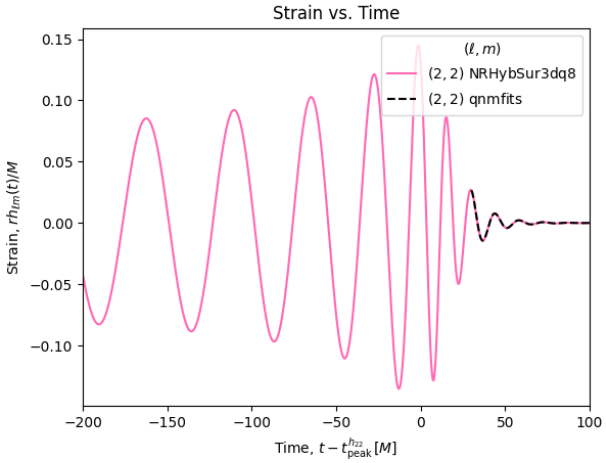

FIG. 2: Plot of full waveform from NRHybSur3dq8 and model from `qnmfits` for $t_0 = 30M$

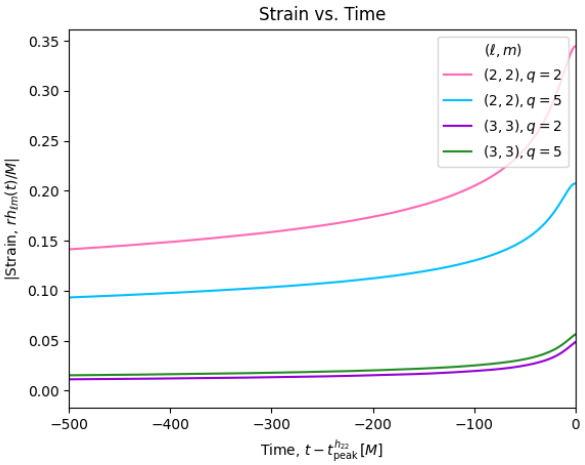

FIG. 3: Plot of strain vs. time for the (3,3) and (2,2) modes for $q = 2$ and $q = 5$.

Closer to the time of the merger, nonlinearities cause the QNM model to break down [28]. To make the QNM model valid, we must start at some time after the merger, which approximately corresponds to the time of peak strain of the quadrupole mode. The quadrupole mode is expected to be dominant in systems with equal masses and aligned spins, but higher harmonics are expected to be much more important in precessing systems [29]. In Figure 2, the start time for the QNM fit is $t_0 = 30M$. Other surrogates define t_0 at different times. We use a least-squares fit with the Python package `qnmfits` to determine which combination of QNMs yields the observed NRHybSur3dq8 waveform.

A. Testing Surrogates for Aligned Spin Systems

Figure 3 shows strain vs. time for example waveforms generated with NRHybSur3dq8 [26] with different mass ratios. This demonstrates how the amplitudes vary for different modes, and motivates modeling the relationship between amplitude and mass ratio directly for different modes. We plot the relationship between amplitude and mass ratio for different modes at $t_0 = 30M$ in Figure 4. This relationship is useful because it directly relates amplitude, a property directly measurable in the ringdown, and mass ratio, a property relating the progenitor BHs.

Dimensionless strains for QNMs are related to the amplitude by:

$$\begin{aligned} h_{\ell mn} &\sim C_{\ell mn} e^{-i\omega_{\ell mn} t} \\ &\sim A_{\ell mn} e^{i\phi_{\ell mn}} e^{-i(2\pi f_{\ell mn} - i/\tau_{\ell mn})t} \end{aligned}$$

where h is dimensionless strain, C is the complex amplitude, ω is the complex frequency, A is the real amplitude, f is the real frequency, and τ is the damping time. The complex amplitudes obtained from `qnmfits`, C, require a

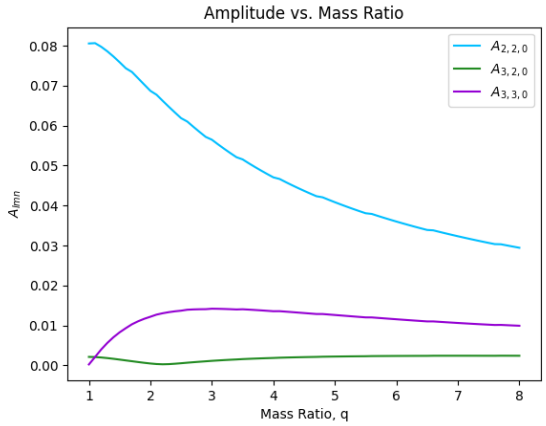

FIG. 4: Plot of corrected amplitude vs. mass ratio for selected modes from NRHybSur3dq8. This plot shows how for different mass ratios, amplitude varies differently for different modes. This could affect which combination of modes would be optimal to look at for different systems.

correction to account for different start times in different fits,

$$C_{\text{corrected}} = Ce^{-i\omega\Delta t}$$

where Δt is the difference in time between the start of a fit and a reference time t_{ref} we are propagating back to. In this project, we propagate all fits back to $t_{\text{ref}} = 0$ to compare them more easily.

Taking the absolute value of the complex amplitudes yields the real amplitudes, A. Similarly, finding the angle counterclockwise from the positive real axis in the complex plane gives the phase, ϕ. These quantities will be useful when a remnant surrogate is eventually inverted to take amplitude and phase as inputs and return mass ratio and spins.

Figure 5a shows corrected amplitudes vs. mass ratio from `qnmfits` compared to those implemented in `jaxqualin` code [24]. `Jaxqualin` code provides functions of mass ratio for both the amplitude and phase. Although the `qnmfits` and `jaxqualin` amplitude models agreed reasonably well, the phases were substantially discrepant, as can be seen in Figure 5b. We implement an additional surrogate, NRSur3dq8_RD (via the `surfinBH` code; see Table I), to see if it more closely aligns with NRHybSur3dq8 (via `qnmfits`). NRSur3dq8_RD matches NRHybSur3dq8 more closely in amplitude and relative phase than the `jaxqualin` fit.

Finally, we test an additional code-base for ringdown surrogates, `qnmpredictor`, which uses models trained on fits produced in Ref. [27]. This one appeared to align reasonably with `qnmfits`, but has other discrepancies that will be discussed further in Section III B. Moving forward, `surfinBH` and NRSur3dq8_RD has shown to be a more effective surrogate to compare with `qnmfits` for NRHyb-

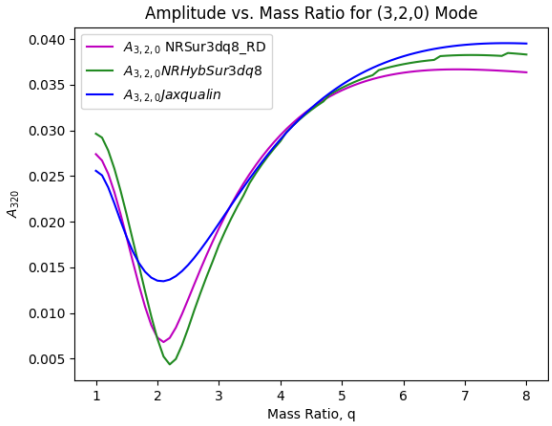

(a) Plot of amplitude vs. mass ratio for NRSur3dq8_RD, NRHybSur3dq8, and `jaxqualin`. This plot shows the $(3, 2, 0)$ mode from Figure 4 close up, as well as other surrogates in the same mode. In this plot, `jaxqualin` appears to have reasonable agreement with `qnmfits`, but still more of a difference than `surfinBH`. These plots were made using different spin parameters, so the vertical axes do not align perfectly. Regardless, the overall shape is the same, and clearly shows the variance of this mode for different mass ratios. For example, the $(3, 2, 0)$ appears to be less excited in systems with a mass ratio close to 2.

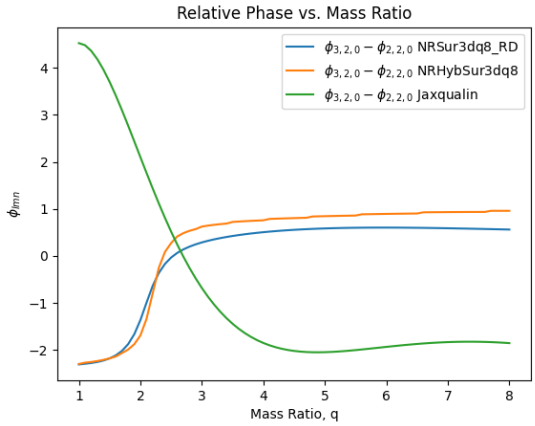

(b) Plot of relative phase for $(3, 2, 0)$ mode vs. mass ratio forNRSur3dq8_RD, NRHybSur3dq8, and `jaxqualin`. Since overall phase is determined by a sum of all the phases of modes, having different relative phases changes these sums, and the overall observed phases. It is clearly shown in this plot that `surfinBH` aligns much more with `qnmfits` than `jaxqualin`. This could have been a result of a misinterpretation of the outputs of the `jaxqualin` equations, but it is unclear.

FIG. 5: These plots compare amplitude and relative phase plotted against mass ratio for selected modes in NRSur3dq8_RD and `surfinBH`, NRHybSur3dq8 and `qnmfits`, and `jaxqualin`.

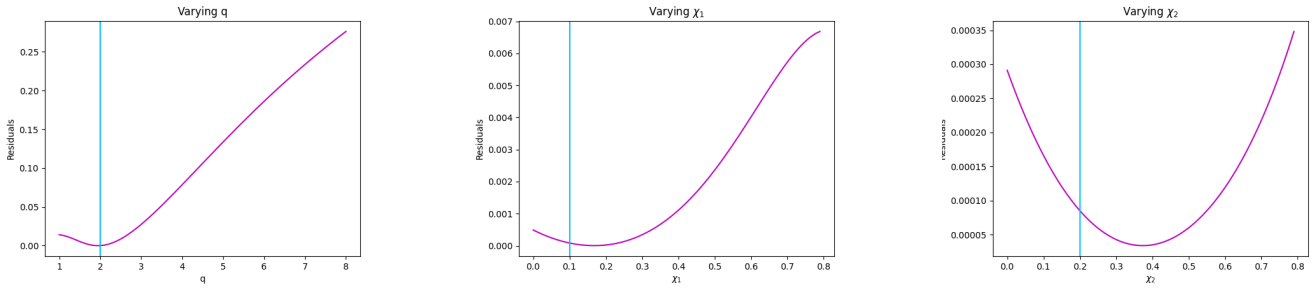

(a) Varying q while $\chi_1 = 0.1$ and $\chi_2 = 0.2$

(b) Varying χ_1 while $q = 2$ and $\chi_2 = 0.2$

(c) Varying χ_2 while $q = 2$ and $\chi_1 = 0.1$

FIG. 6: Sum of the squared residuals of the complex amplitudes (magenta) from the NR surrogate model and the QNM model plotted against q, χ_1, and χ_2, where one of these parameters varies on the horizontal-axis and the other two are held constant. The true values for q, χ_1, and χ_2 are shown by the vertical turquoise lines. The modes included in this plot are $(2, 2, 0)$, $(2, 2, 1)$, and their mirror modes.

Sur3dq8. `Jaxqualin` will not be used due to the large discrepancy in phases with `qnmfits`.

B. Inverting a Surrogate

The purpose of using a surrogate that requires inspiral parameters as inputs and generates a waveform is to invert the surrogate to return inspiral parameters as outputs when given amplitudes and phases from a QNM model.

We calculate a residual vector of differences between the real and imaginary parts of the complex amplitudes from `qnmfits` for NRHybSur3dq8 to those produced by the remnant surrogate `surfinBH` and NRSur3dq8_RD. We run `scipy.optimize.least_squares` and determine where the sum of the squared residual is minimized:

$$\begin{aligned} r_{\mathrm{Re},i} &= \mathrm{Re}\left[C_{\mathrm{model}} - C_{\mathrm{observed}}\right] \\ r_{\mathrm{Im},i} &= \mathrm{Im}\left[C_{\mathrm{model}} - C_{\mathrm{observed}}\right] \\ \mathrm{residual} &= \sum_{i \in \{(\ell,m,n)\}} \left(r^2_{\mathrm{Re},i} + r^2_{\mathrm{Im},i}\right) \end{aligned}$$

where i is indexed over the modes of interest. Inside this least squares fit, `scipy.optimize.least_squares` cycles through q, χ_1, and χ_2 values and finds the values at which this residual is minimized. This is where the remnant surrogate model and the QNM model have the smallest difference.

Shown in Figure 6 for the aligned-spin case, the true values of the full IMR waveform (turquoise vertical line) were not perfectly predicted by the remnant surrogate (minima of the magenta curve). This motivates investigating other remnant surrogate models to determine how much these discrepancies are surrogate dependent or if they are a result of varying true system parameters. Indeed, we followed the same process for `qnmpredictor` and successfully inverted this surrogate as well, but encountered problems with related to the the initial guess needed for running `scipy.optimize.least_squares`. This is discussed further in Section III D.

C. Further Tests with `surfinBH`

Because the `surfinBH` implementation of NRSur3dq8_RD is a working inverted surrogate which yields reasonable estimates for true parameters, we use it to explore the Objectives given in Section II. In Figure 7, we investigate how well this model can estimate mass ratio in systems with more vs. less extreme spins. The upper left panel shows a case where $\chi_1 = \chi_2 = 0.1$ (small spins), and the lower right is when $\chi_1 = \chi_2 = 0.7$ (large spins). The off-diagonals are the case where one spin is small and the other is large. To better quantify these findings, we plot the absolute errors of these fits in Figure 8:

$$\mathrm{Error} = |q_{\mathrm{inferred}} - q_{\mathrm{true}}| \ .$$

The error of the fit increases as spin magnitude increases, particularly in the higher mass ratio region of the parameter space. In the cases with one large spin and one small spin, the fit is worse when χ_1 is high rather than χ_2. This is because since χ_1 is the spin of the more massive BH, in systems with a greater mass ratio, it is expected that χ_1 will be the more prevalent spin and have a greater impact on the observed waveform.

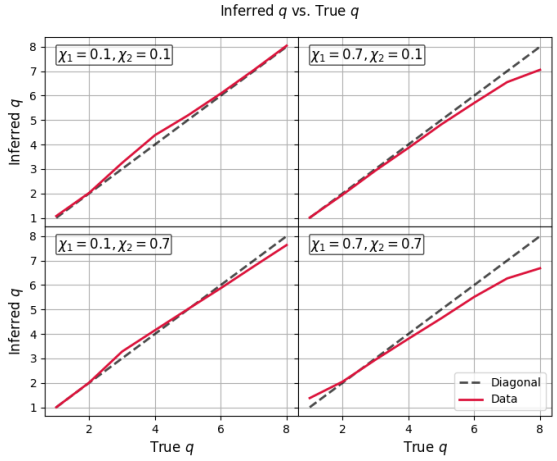

FIG. 7: This figure shows inferred q vs. true q for four different combinations of high and low spins for χ_1 and χ_2. The fit is better in systems with lower spins, and begins to lose accuracy in systems with more extreme spins, especially for large χ_1. The modes included in this fit are $(2, 2, 0)$ and $(3, 2, 0)$.

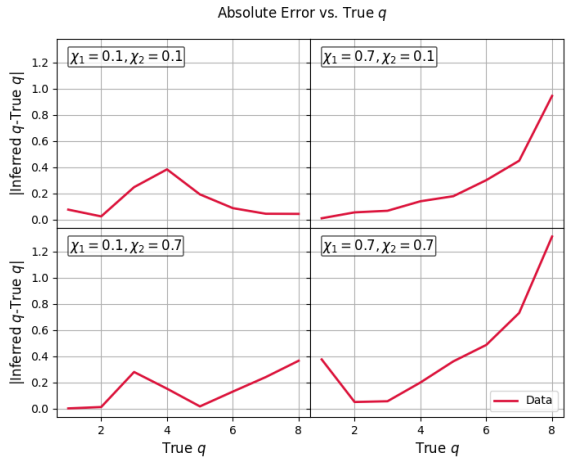

FIG. 8: This plot shows the absolute errors of the fits in Figure 7. The highest error occurs in systems with high mass ratio and high spins.

Likewise, we explore how well the inverted surrogate estimates χ_1 and χ_2 for varying q, as shown in Figure 9. With the exception of the $q = 1$ case for χ_1, the fit for spin magnitudes performs substantially worse than for q (Figure 7). In the case of χ_2 specifically, the fit is getting stuck on the upper boundary. This could be happening because it is not finding a good fit at all, and not necessarily because it estimates the value at the boundary for all q. This is something we hope to solve with parameter estimation; with parameter estimation, we will be able to determine if the fit is uninformative due to poorly constrained parameters or incorrect. As a test, we changed the bounds of the fit, and it still got stuck at the upper boundary regardless of what the value was.

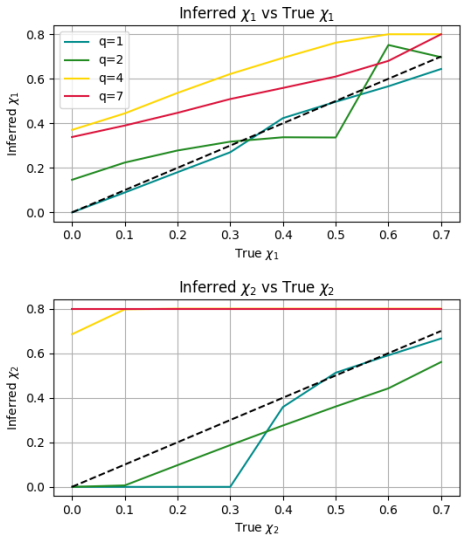

FIG. 9: This plot shows estimates for χ_1 and χ_2 for four different values of q, which are colored the same in each plot. The spin that is not being fit for in each plot is held constant at $\chi_1 = \chi_2 = 0.1$. In all cases, the fit for χ_1 performs better than for χ_2. Higher mass ratios have worse estimates. The modes included in this fit are $(2, 2, 0)$ and $(3, 2, 0)$.

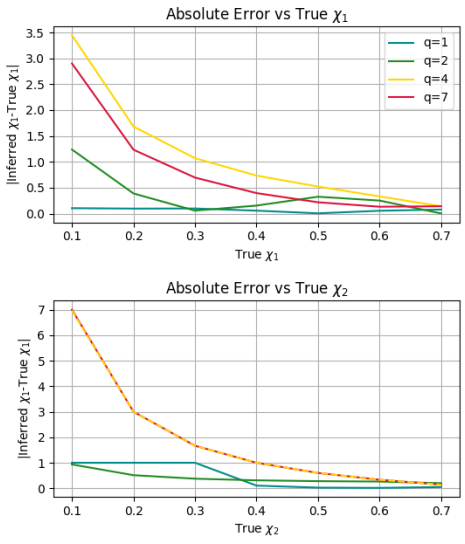

FIG. 10: Absolute error of the fits shown in Figure 9

In both cases of looking at systems with extreme mass ratios and extreme spins, fits generally worsen for higher

q, χ_1, and χ_2. We do not yet have an algorithmic way of determining where these fits can or cannot be used in the parameter space, making it hard to assess when a fit is "good enough." This is something we hope to better confine in the future, but as of now it is qualitative assessments.

D. Initial Guess Testing

The least squares fitting method requires some initial guess for the inspiral parameters. A problem we are encountering is that some of the ringdown surrogates only yield reasonable estimates if the initial guess is equal to the true value; this case is shown in Figure 11. Particularly, the `qnmpredictor` ringdown surrogate seems to also be getting stuck on a boundary unless a highly tailored initial guess is passed, leading us to believe that there is a bug in our implementation. Until we resolve this issue, further tests that can be done with `qnmpredictor` are limited.

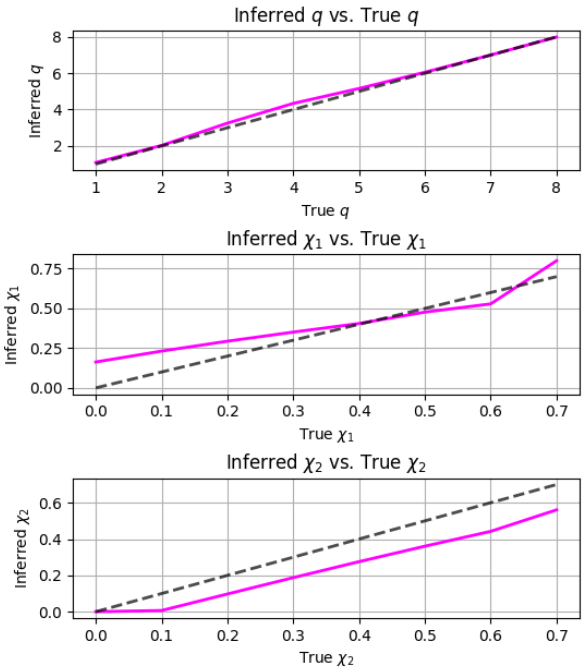

FIG. 11: Plot of inferred vs. true parameters for when the initial guess is equal to the true value using `surfinBH` and NRSur3dq8_RD. This performs reasonably well for mass ratio, and less so for spins. Each subplot varies one injected parameter at a time, with the other parameters fixed to their true values $q = 2$, $\chi_1 = 0.1$, and $\chi_2 = 0.2$.

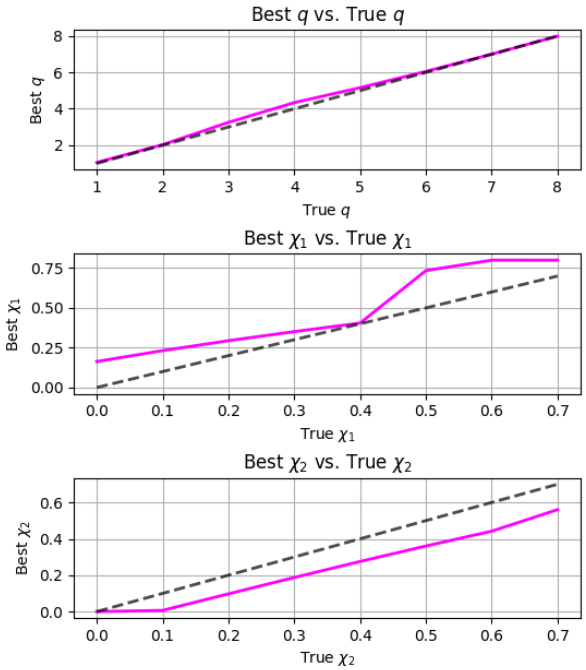

FIG. 12: This plot shows the same information as Figure 11, but with systematically chosen initial guesses instead. It appears to be a slightly worse fit, particularly in χ_1, but still reasonable. The fit for χ_2 appears to be influenced by initial guess less than the fit for χ_1. This plot uses `surfinBH` and NRSur3dq8_RD.

To mitigate the need for knowing *a priori* the true parameters of the system, we implement a method for systematically choosing a value for the initial guess. We cycle through an array of initial guesses for each parameter, and chose the guess where the returned best fit value has the lowest residual. Results for the inferred value using this method are shown in Figure 12.

IV. FUTURE WORK

The remnant surrogate NRSur3dq8_RD does not predict the true parameters as accurately in cases with more extreme mass ratios and spins, as we show in Figures 7 and 9. This motivates investigating if accuracy of the fit varies between as different QNMs are included/excluded, and if so, which QNMs are necessary to most accurately recover parameters. Other modes we want to try including in fits are $(3, 3, 0)$ and $(4, 4, 0)$. There are other combinations that could be considered as well, including mirror modes.

Similarly, it is worth further exploring the limits of q, χ_1, and χ_2 and what values cause inferred parameters to be less accurate. A useful approach to solving this will be to test other remnant surrogates and determine which

model can most accurately predict the true parameters. This addresses Point 2 in Sec. II. A problem we are encountering in multiple surrogates is least squares fitting getting stuck on boundaries and not giving accurate estimates. We are investigating looking at sampling methods to determine if we can see how these fits actually behave in particular cases where least squares is failing.

So far, we have only generated point estimates for different combinations of QNMs that could yield an observed waveform. We have just begun doing preliminary exploration using a novel branch of the `ringdown` code to do full parameter estimation. Moving forward, we will continue to work with this code and test increasing precession, addressing Point 3 in Sec. II.

The ultimate goal is to produce a similar model to map between QNMs and inspiral parameters for the precessing case. However, modeling the precessing case is much more difficult. A useful intermediate is determining at what level of in-plane spin the aligned spin analysis breaks down. If we can produce an effective model for precessing systems, the result could be applied to real GW data, such as GW231123 or GW190521 (Point 4 in Sec. II).

[1] A. Buikema *et al.*, Phys. Rev. D **102**, 062003 (2020), arXiv:2008.01301 [astro-ph.IM].
[2] H. Zhu *et al.*, Phys. Rev. D **111**, 064052 (2025), arXiv:2312.08588 [gr-qc].
[3] S. A. Teukolsky, Astrophys. J. **185**, 635 (1973).
[4] S. Chandrasekhar and S. Detweiler, Proceedings of the Royal Society of London Series A **344**, 441 (1975).
[5] K. D. Kokkotas and B. G. Schmidt, Living Rev. Rel. **2**, 2 (1999), arXiv:gr-qc/9909058.
[6] S. Detweiler, Astrophys. J. **239**, 292 (1980).
[7] O. Dreyer, B. J. Kelly, B. Krishnan, L. S. Finn, D. Garrison, and R. Lopez-Aleman, Class. Quant. Grav. **21**, 787 (2004), arXiv:gr-qc/0309007.
[8] 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].
[9] L. London, D. Shoemaker, and J. Healy, Phys. Rev. D **90**, 124032 (2014), [Erratum: Phys.Rev.D 94, 069902 (2016)], arXiv:1404.3197 [gr-qc].
[10] L. T. London, Phys. Rev. D **102**, 084052 (2020), arXiv:1801.08208 [gr-qc].
[11] S. Borhanian, K. G. Arun, H. P. Pfeiffer, and B. S. Sathyaprakash, Class. Quant. Grav. **37**, 065006 (2020), arXiv:1901.08516 [gr-qc].
[12] M. H.-Y. Cheung, E. Berti, V. Baibhav, and R. Cotesta, Phys. Rev. D **109**, 044069 (2024), [Erratum: Phys.Rev.D 110, 049902 (2024)], arXiv:2310.04489 [gr-qc].
[13] R. Abbott *et al.* (KAGRA, VIRGO, LIGO Scientific), Phys. Rev. X **13**, 041039 (2023), arXiv:2111.03606 [gr-qc].
[14] C. L. Rodriguez, M. Zevin, C. Pankow, V. Kalogera, and F. A. Rasio, Astrophys. J. Lett. **832**, L2 (2016), arXiv:1609.05916 [astro-ph.HE].
[15] I. Mandel and A. Farmer, Phys. Rept. **955**, 1 (2022), arXiv:1806.05820 [astro-ph.HE].
[16] "GW231123: a Binary Black Hole Merger with Total Mass 190-265 $M_\odot$," (2025), arXiv:2507.08219 [astro-ph.HE].
[17] R. Abbott *et al.*, Physical Review Letters **125** (2020), 10.1103/physrevlett.125.101102.
[18] S. J. Miller, M. Isi, K. Chatziioannou, V. Varma, and I. Mandel, Physical Review D **109** (2024), 10.1103/physrevd.109.024024.
[19] R. Udall, S. Hourihane, S. Miller, D. Davis, K. Chatziioannou, M. Isi, and H. Deshong, Physical Review D **111** (2025), 10.1103/physrevd.111.024046.
[20] E. Payne, S. Hourihane, J. Golomb, R. Udall, D. Davis, and K. Chatziioannou, Physical Review D **106** (2022), 10.1103/physrevd.106.104017.
[21] S. J. Miller, M. Isi, K. Chatziioannou, V. Varma, and S. Hourihane, "Measuring spin precession from massive black hole binaries with gravitational waves: insights from time-domain signal morphology," (2025), arXiv:2505.14573 [gr-qc].
[22] F. Nobili, S. Bhagwat, C. Pacilio, and D. Gerosa, (2025), arXiv:2504.17021 [gr-qc].
[23] L. Magaña Zertuche, L. Gao, E. Finch, and G. B. Cook, (2025), arXiv:2502.03155 [gr-qc].
[24] M. H.-Y. Cheung, E. Berti, V. Baibhav, and R. Cotesta, Physical Review D **109** (2024), 10.1103/physrevd.109.044069.
[25] V. Varma, D. Gerosa, L. C. Stein, F. Hébert, and H. Zhang, Physical Review Letters **122** (2019), 10.1103/physrevlett.122.011101.
[26] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, Physical Review D **99** (2019), 10.1103/physrevd.99.064045.
[27] K. Mitman, I. Pretto, H. Siegel, M. A. Scheel, S. A. Teukolsky, M. Boyle, N. Deppe, L. E. Kidder, J. Moxon, K. C. Nelli, W. Throwe, and N. L. Vu, "Probing the ringdown perturbation in binary black hole coalescences with an improved quasi-normal mode extraction algorithm," (2025), arXiv:2503.09678 [gr-qc].
[28] E. Finch, *Black-hole Ringdown*, Ph.D. thesis, University of Birmingham (2023).
[29] J. Roulet and T. Venumadhav, Annual Review of Nuclear and Particle Science **74**, 207–332 (2024).