Inferring Properties of Binary Black Hole Mergers from the Ringdown

Final Report, 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: November 1, 2025)

The effect of mass and spin on the inspiral phase of a gravitational-wave signal emitted by a binary black hole coalescence is relatively well understood when the black holes' spins are aligned with the system's orbital angular momentum. Less is known about how progenitor properties – especially misaligned spins – impact the ringdown, the portion of the signal emitted after binary has merged into a single remnant black hole. 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 large misaligned spins, the ability to predict inspiral properties from the ringdown would improve existing measurements, and potentially give new insight into the high-mass 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 in the ringdown back to inspiral parameters. This method varies in quality depending on the specific surrogate model. That which we explore the most, 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 with various combinations of quasi-normal modes. We find that increasing the number of modes tends to generate a more precise fit for component masses and spins, but that different single or two-mode combinations are required to better-fit different mass-spin configurations.

I. INTRODUCTION

A binary black hole (BBH) coalescence can be broken into three segments: inspiral, merger, and ringdown. During the inspiral, the two progenitor black holes (BHs) orbit each other and emit gravitational-waves (GWs). The GW frequency increases as the BHs draw nearer together until they eventually collide in the merger, resulting in a single remnant BH. Immediately in the time after the merger, this highly perturbed remnant BH rotates asymmetrically about its axis, still losing energy in the form of GWs as it equilibrates into its final stage, a Kerr BH [1].

Most BBHs observed by LIGO-Virgo-KAGRA [2–4] (the LVK) are inspiral-dominated, due to the range of frequencies over which these ground-based interferometric detectors are most sensitive [5]. This is convenient for measuring inspiral parameters of these binary systems, such as the masses and spins of the progenitor BHs. Measuring these parameters becomes less straightforward in systems that are not inspiral-dominated. Systems with a sufficiently high total mass emit GWs at low enough frequencies during the inspiral that this portion of the signal falls outside of the LVK sensitivity band, making it difficult to observe a complete signal from these systems [5]. Figure 1 shows a comparison of a lower mass and higher mass system, demonstrating how the detectable duration of the inspiral changes with total mass.

Although masses and spins of progenitor BHs are properties of the inspiral, these parameters also determine the characteristics of the ringdown signal, such as the amplitude and phase of modes in the ringdown. If there is a relationship between the inspiral parameters of the system and the ringdown signal, it stands to reason that studying the ringdown would give insight into what parameters of a system produced a given ringdown waveform.

The ringdown signal consists of a sum of damped sinusoids called quasinormal modes (QNMs) [6–10], which are emitted from the perturbed BH at fundamental and overtone frequencies which are determined by the system's masses and spins. "Quasi" refers to the fact that the modes decay characteristically over time, whereas normal modes do not. The longest lasting of these modes 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.

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 BHs have been encoded in various fitting formulae or ringdown surrogate models [11]. However, in reality, most BHs possess at least some degree of spin-orbit misalignment [12], leading to spin-orbit precession.

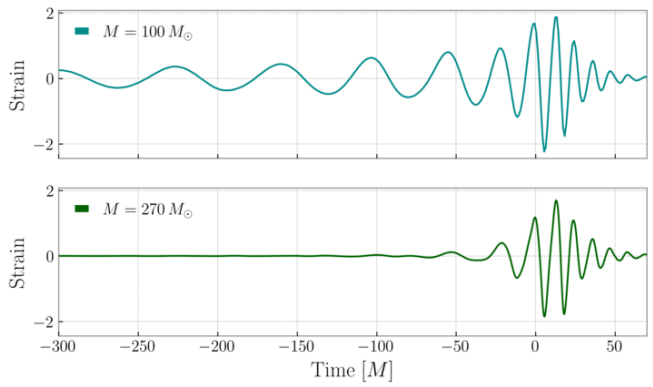

FIG. 1: The top panel shows an example of a full inspiral merger ringdown (IMR) signal of a BBH merger. The bottom panel shows the analogous data for a system with a higher total mass. These plots show detector strain vs. time, whitened using a representative noise power spectral density from LIGO's third observing run, such as to visualize the absence of inspiral data in high mass detections. We use natural units for time where mass is scaled out of both axes to more clearly show comparative waveform morphology between the two signals.

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 amplitudes to inspiral properties is less clear for precessing binaries, although a relationship is expected [1]. The precessing case is particularly of interest because of its potential to give astrophysical insight into how a BBH system formed [e.g., 13, 14]. 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. However, this analysis is difficult to do for precessing systems; many necessary tools remain in development. As such, this paper focuses on the analysis of aligned spin systems, which will produce a useful framework for repeating a similar analysis for precessing systems.

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 [15] 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 [16]; its precession measurement stems from the subtle interplay between the quiet final inspiral cycle and the loud merger [17]. 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 [18] and the precession in GW200129 [19] are both degenerate with transient non-Gaussian detector noise (i.e.,"glitches") the overlap 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 data itself serves as motivation for this study.

Section II discusses the function and use of ringdown surrogates. Section III details the process of inverting a surrogate using least-squares in subsections III A and III B and discusses tests done using this method in subsection III C. Section IV describes inverting a ringdown surrogate using sampling. Section V discusses the findings of this project and future work to be done.

II. RINGDOWN AMPLITUDE SURROGATES

There is no known way to analytically map between the ringdown signal and inspiral parameters. To accurately infer inspiral parameters from the ringdown, other methods must be explored. This paper focuses on using a ringdown amplitude surrogate to predict parameters. We begin by analyzing simulated signals (for which we *a priori* 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.

The function of a ringdown amplitude surrogate is to take the mass ratio (q) and spins of the progenitor BHs (χ_1 and χ_2) as inputs and generate ringdown amplitudes and phases as outputs. A corresponding ringdown waveform can then be constructed by summing QNMs with the corresponding amplitudes and phases. These amplitudes and phases are indexed with (ℓ, m, n), which relate to the spherical harmonics of the QNMs. The ℓ index is the multipole number and describes where the mode is located spatially around a BH. The m index describes the projection of the angular momentum onto the rotational axis of the BH and determines frequency. The n index is the overtone number, which indicates decay time. The fundamentals where $n = 0$ take the longest to decay, but higher overtones can be more significant at earlier times in the ringdown [20].

This paper discusses a variety of surrogates and models which are summarized in Table I. The package `qnmfits` implements least squares to fit for amplitudes and phases from a reference time series, such as an inspiral-merger-ringdown (IMR) waveform [21]. We here use the IMR waveform NRHybSur3dq8 [22], which models aligned-spin systems and is based off of interpolating full numerical relativity waveforms. The `Jaxqualin` package predicts amplitudes and phases based on polynomial approximations [11]. NRHybSur3dq8_RD and q8_3dAl are

both ringdown amplitude/phase surrogates constructed on the premise of modeling and summing QNMs [23, 24]. These two ringdown surrogates are trained using machine learning techniques based on Gaussian processes.

Surrogate/Model	Code Base	Description	Ref.
–	`qnmfits`	Fit QNMs with least squares	[21]
–	`Jaxqualin`	QNM amplitude fits based on polynomials	[11]
NRSur3dq8_RD	`surfinBH`	QNM amplitude fits based on Gaussian processes	[23]
NRHybSur3dq8	`gwsurrogate`	IMR waveform surrogate	[22]
q8_3dAl	`qnmpredictor`	QNM amplitude fits based on Gaussian processes	[24]

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

Ringdown amplitudes and phases are determined by q, χ_1, and χ_2. The relationships between these inspiral parameters and the resultant QNMs can be mapped with surrogates to better understand this effect in different modes and across parameter space. Figure 2 shows an example of this relationship between mass ratio and amplitude for three different QNMs: the $(\ell, m, n) = (2,2,0)$, (3,2,0), and (3,3,0) modes. We can see, for example, that the (2,2,0) mode becomes less excited as mass ratio becomes more extreme, while the opposite is true for the (3,3,0) mode. Relationships like this are useful because they directly relate amplitude, a property directly measurable in the ringdown, and mass ratio, a property relating the progenitor BHs. The vertical axis of this plot is the real part of the amplitude, which is related to the GW strain h by:

$$\begin{aligned} h_{\ell mn} &\sim C_{\ell mn} e^{-i\omega_{\ell mn} t} \\ &\sim A_{\ell mn} e^{i\phi_{\ell mn}} e^{-i(2\pi f_{\ell mn} - i/\tau_{\ell mn})t} \end{aligned} \tag{1}$$

where C is the complex amplitude, ω is the complex frequency, f is the real frequency, and τ is the damping time.

III. INVERTING A RINGDOWN SURROGATE USING LEAST SQUARES

A. Method

If a ringdown surrogate generates some unique amplitudes and phases for different QNMs, it follows that those amplitudes and phases could be studied to back-infer what progenitor parameters produced them. This idea motivates *inverting a ringdown surrogate*. The following is a brief overview of our method for implementing

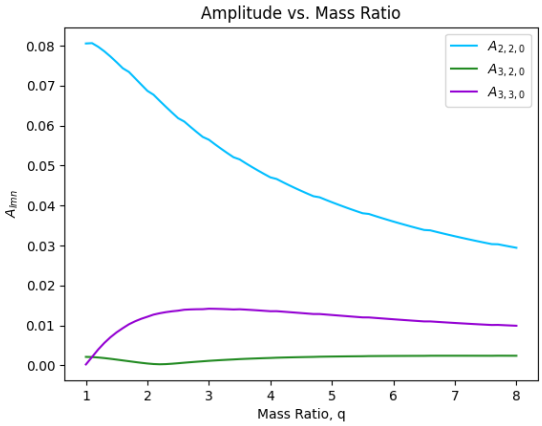

FIG. 2: Plot of the real amplitude vs. mass ratio for selected modes, generated with NRHybSur3dq8. For different modes, amplitude varies differently with increasing mass ratio. Understanding these individual components gives better insight into the construction of an observed waveform.

an inverted surrogate, a process which is non-trivial and non-analytic.:

1. Create simulated IMR data using a numerical relativity (NR) surrogate, e.g., NRHybSur3dq8. This waveform (shown in pink in Figure 4) acts as the "observed" data to which the inverted ringdown surrogate is applied. Using a simulated IMR waveform with known inspiral parameters allows us to investigate how accurately the inverted surrogate predicts q, χ_1, and χ_2.

2. Fit for amplitudes and phases in the ringdown of the IMR waveform from the NR surrogate. This fit, shown in black in Figure 4, is done with the package `qnmfits`. We treat this single best-fit point estimate as our "true observed" QNM spectrum.

3. Use inverted ringdown surrogate to infer values for q, χ_1, and χ_2 based on the "true observed" QNMs. A variety of ringdown surrogates and fitting formulas were tested with this method of inversion, which is further discussed below.

4. Compare inferred parameters from inverted ringdown surrogate to true parameters used to generate IMR data.

We note that the complex amplitudes obtained from `qnmfits`, which we denote as C, require a correction to account for different start times intrinsic to the different ringdown surrogates explored, i.e., we must compute a $C_{\text{corrected}}$ such that

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

where Δt is the difference in time between the start of a fit and a reference time t_{ref} we are propagating back

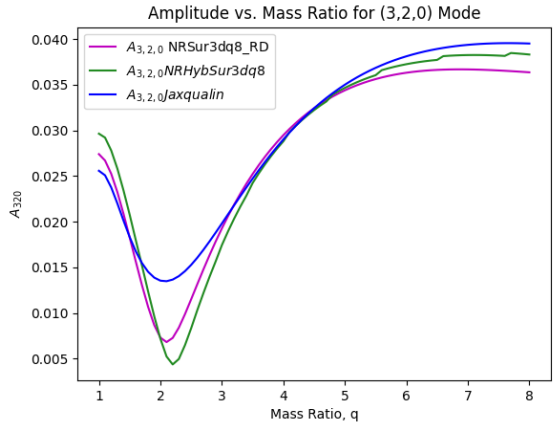

(a) Plot of amplitude vs. mass ratio for NRSur3dq8_RD, NRHybSur3dq8, and `jaxqualin`. This plot shows the $(3, 2, 0)$ mode from Figure 2 close up, as well as other surrogates for the same mode. In this plot, `jaxqualin` agrees reasonably with `qnmfits`, but still more of a difference than `surfinBH`. Figure 3a and figure 2 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 least excited in systems with a mass ratio close to 2.

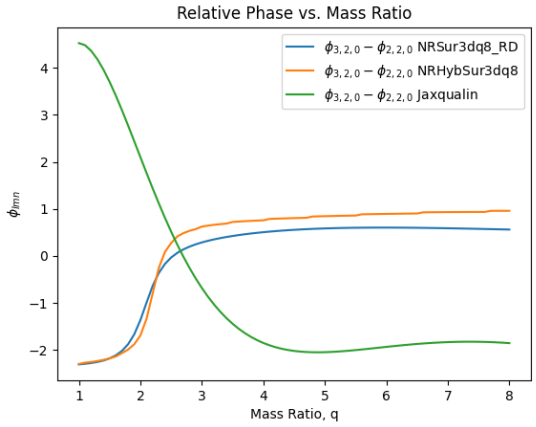

(b) Plot of relative phase for $(3, 2, 0)$ mode vs. mass ratio for NRSur3dq8_RD, NRHybSur3dq8, and `jaxqualin`. Since overall phase is determined by a sum of all the phases of modes, having different relative phases changes these sums, and the overall observed phases. It is clearly shown in this plot 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 and warrants future investigation.

FIG. 3: These plots compare amplitude and relative phase plotted against mass ratio for selected modes in NRSur3dq8_RD and `surfinBH`, NRHybSur3dq8 and `qnmfits`, and `jaxqualin`.

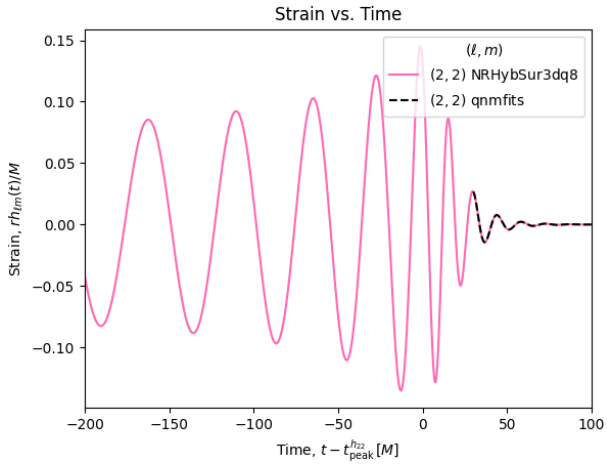

FIG. 4: The pink line shows the simulated IMR waveform from NRHybSur3dq8. The dashed black line shows the fit to the data with the package `qnmfits`; this line begins at some time after the merger $t_0 = 30M$ when nonlinear effects from the merger have subsided sufficiently so that a ringdown model based on QNMs is valid [25].

to, and ω is the complex frequency. In this project, we propagate all fits back to $t_{\rm ref} = 0$ (i.e., the time of merger) to compare them more easily.

To invert a ringdown surrogate, we calculate a residual vector of differences between the real and imaginary parts of the "true observed" complex amplitudes ($C_{\rm observed}$, i.e., from `qnmfits` on the IMR NRHybSur3dq8 waveform) to those produced by the remnant surrogate ($C_{\rm model}$). We here focus on the ringdown surrogate NRSur3dq8_RD, implemented with the codebase `surfinBH`. We run `scipy.optimize.least_squares` and determine where the sum of the squared residual is minimized:

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

where i is indexed over the modes of interest, $\{(\ell, m, n)\}$. Inside this least squares fit, `scipy` 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 5 for the aligned-spin case, the true values of q, χ_1, and χ_2 for the full IMR waveform (turquoise vertical line) were not perfectly predicted by the remnant surrogate (minima of the magenta curve). In this test case, we simulate a system with a mass ratio of $q = 2$, $\chi_1 = 0.1$, and $\chi_2 = 0.2$, i.e, comparable masses and small spins. The inverted surrogate inferred q most accurately. The curve for χ_1 is much wider and has a

less clear minimum. In other systems, this minimum can be even less clear, rendering the inverted surrogate incapable of returning a reasonable estimate at all. The same is true for χ_2 but it is typically less constrained, even in ideal systems. This lack of perfect recovery motivates investigating other remnant surrogate models to determine how much these discrepancies are surrogate-dependent, or if they are a result of poor fits in certain areas of parameter space across all different surrogates.

Figure 3 compares NRSur3dq8_RD, NRHybSur3dq8, and `jaxqualin` in amplitude (3a) and in phase (3b). These plots show the output of passing all the different ringdown surrogates the same initial parameters. In this case, spins are held constant at $\chi_1 = 0.1$ and $\chi_2 = 0.2$, and q is varied. Although the same input was used, there are clear discrepancies in the outputs, particularly in phase. This suggests that errors in inferred parameters from an inverted surrogate are not independent of the model used.

B. Initial Guess Testing

The least squares fitting method requires some initial guess for the inspiral parameters. A problem we encountered is that some of the ringdown surrogates only yield reasonable estimates for mass ratio and spins if the initial guess is equal to the true value; this case is shown in Figure 6 for the NRSur3dq8_RD surrogate over a grid of q, χ_1, χ_2 using the $(2, 2, 0)$ and $(3, 2, 0)$ modes.

Particularly, the q8_3dAl ringdown surrogate implemented with the `qnmpredictor` code-base seems to also be getting stuck on a parameter-space boundary unless a highly tailored initial guess is passed, leading us to believe that there is likely a bug in our implementation. Until we resolve this issue, further tests that can be done with `qnmpredictor` are limited; such test and debugging fall under future work (Section V).

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 7. Although we are using a systematically chosen initial guess, these fits still are not perfect; least-squares locates the value for q, χ_1, or χ_2 where the residual calculated in section III A is minimized. In cases where many different values have similar differences in residuals, there is no clear best estimate. On a plot similar to Figure 5, this would appear as a wide curve with no clear local minimum.

C. Further Tests with `surfinBH`

Because our inverted `surfinBH` implementation of NRSur3dq8_RD yields reasonable estimates for true parameters, we use it to explore other interesting cases in the mass ratio and spin parameter space. In Figure 8, 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 9, where

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

We find that 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. We believe that this is because–since it is the spin of the more massive BH–in systems with a greater mass ratio, χ_1 is expected that to have a greater impact on the observed waveform. Typically in GW signals, the mass-weighted spin is used to better compare to inspiral measurements [11], which is reflected in our findings.

Likewise, we explore how well the inverted surrogate estimates χ_1 and χ_2 for varying q, as shown in Figure 10. With the exception of the $q = 1$ case for χ_1, the fit for spin magnitudes performs substantially worse than for q (Figure 8). In the case of χ_2 specifically, the fit often gets stuck on the upper boundary. This could be happening because our method 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 (Section IV); 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.

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.

Another interesting question to test with this inverted surrogate is what combination of modes produce the best fit. In a real ringdown signal, an extremely limited number of modes can be observed, usually one or two. Exploring quality of fit from different modes provides a baseline for determining how well this parameter estimation could be done on real systems based on what modes are observed.

As shown in Figure 13, fitting with three modes does not improve estimates, compared to Figure 12 where combinations of two modes are shown. This is not expected, as including more modes should be a closer representation of an actual observed signal. The reason for

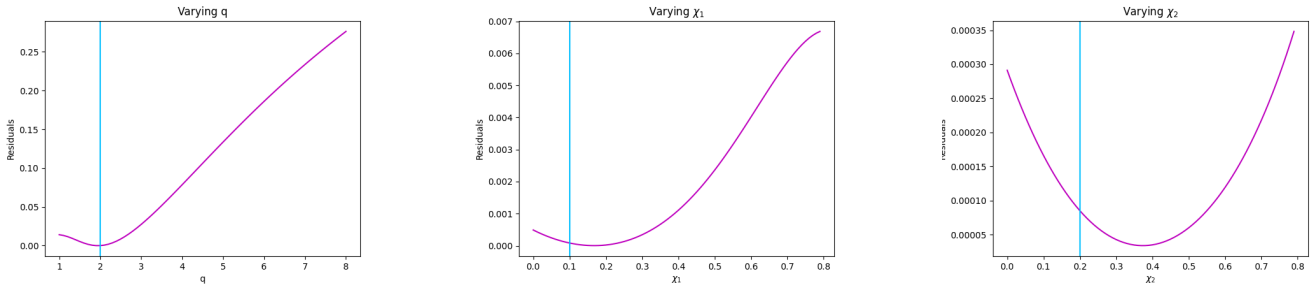

(a) Varying q while $\chi_1 = 0.1$ and $\chi_2 = 0.2$

(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. 5: Sum of the squared residuals of the complex amplitudes (magenta) from the fit to the NR surrogate 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.

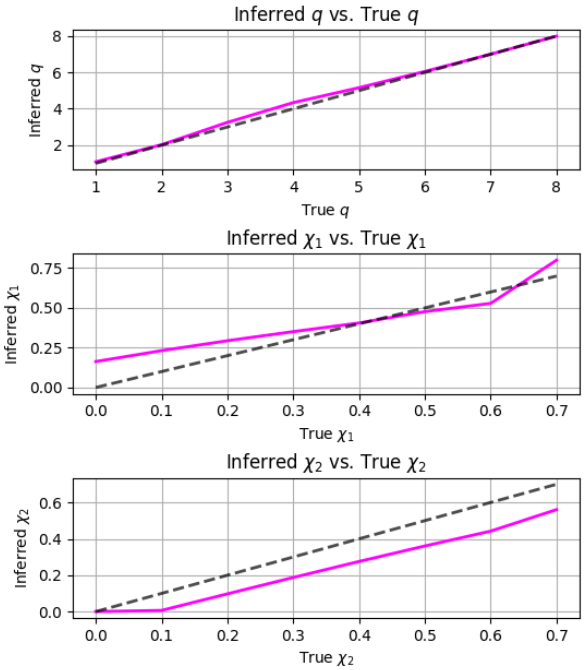

FIG. 6: Plot of inferred vs. true parameters for when the initial guess is equal to the true value using NRSur3dq8_RD with `surfinBH`. The fit performs reasonably well for mass ratio, and less so for spins. If the inverted surrogate was predicting these parameters perfectly, the magenta line would follow the diagonal where true parameters equal inferred parameters, indicated by the black dashed line. 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 $\chi_2 = 0.2$.

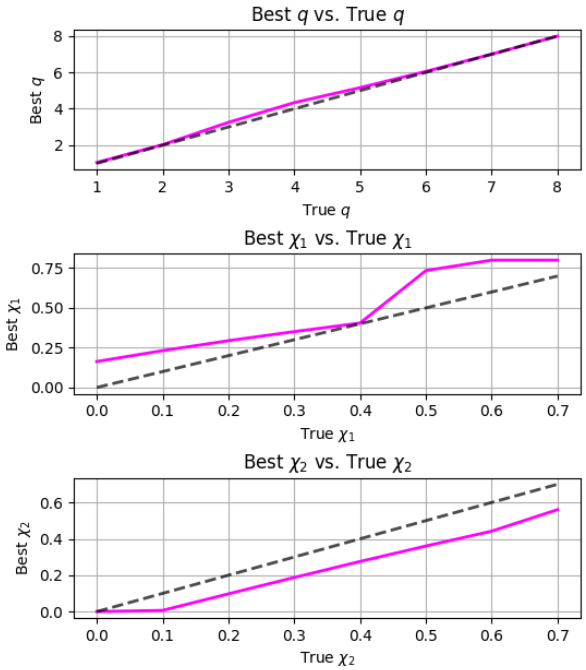

FIG. 7: This plot shows the same information as Figure 6, but with systematically chosen initial guesses instead. The fit appears slightly worse, particularly in χ_1, but still reasonable. The fit for χ_2 is impacted by initial guess less than that for χ_1.

this behavior could be that using a least squares fit depreciates as more degrees of freedom are introduced. This further motivates checking this behavior using sampling, which we discuss in the following section.

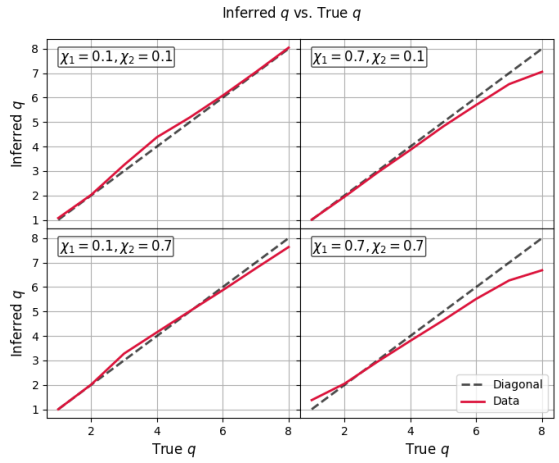

FIG. 8: Inferred q vs. true q for four different combinations of high and low spins for χ_1 and χ_2. The fit is better in systems with lower spins, and begins to lose accuracy in systems with more extreme spins, especially for large χ_1. The modes included in this fit are $(2, 2, 0)$ and $(3, 2, 0)$.

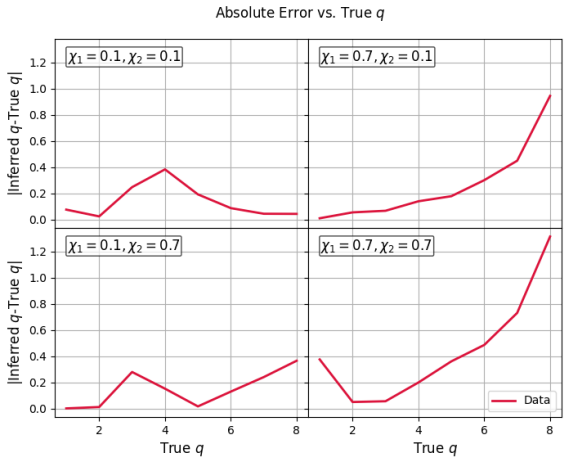

FIG. 9: Absolute errors of the fits in Figure 8. The highest error occurs in systems with high mass ratio and high spins.

IV. INVERTING A RINGDOWN SURROGATE USING SAMPLING

Using the least squares fit to obtain a point estimate fails in cases where there are many different combinations of parameters which yield local minima of our residuals function. By switching to using sampling to invert a surrogate and defining a likelihood, we were able to see where these degeneracies occur and experiment with using different combinations of QNMs to determine which subset are necessary to include to break these degeneracies.

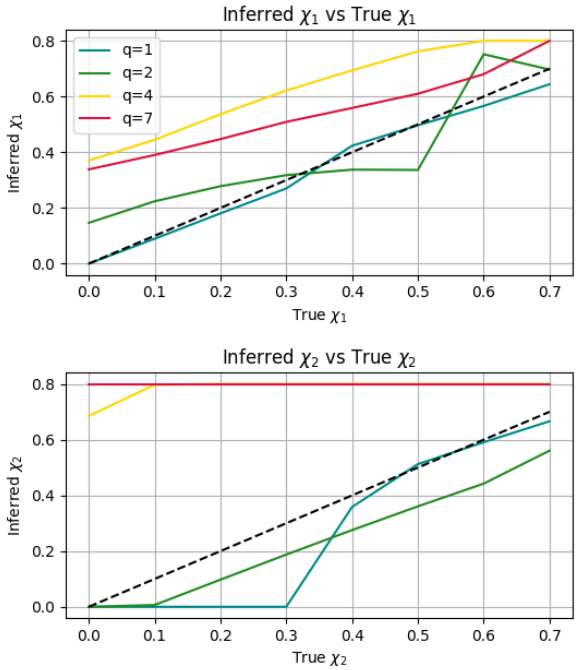

FIG. 10: Estimates for χ_1 and χ_2 for four different values of q, which are colored the same in each plot. The spin that is not being fit for in each plot is held constant at $\chi_1 = \chi_2 = 0.1$. In all cases, the fit for χ_1 performs better than for χ_2. Higher mass ratios have worse estimates. The modes included in this fit are $(2, 2, 0)$ and $(3, 2, 0)$.

The likelihood function is as follows:

$$\log \mathcal{L}_{Re,i}(Re[C_i]|q, \chi_1, \chi_2) \propto -\frac{1}{2}\left(\frac{Re[C_i(q, \chi_1, \chi_2) - C_{i,\text{observed}}]}{\sigma}\right. \quad (2)$$

where C is the complex amplitude indexed over different modes and σ is the uncertainty in the amplitude measurement, which comes from noise of a system. In this project, we have chosen $\sigma = 0.001$ for all results shown. This corresponds to a high signal to noise ratio (SNR). We sample the likelihood in Eq. (2) using a Markov Chain Monte Carlo implemented with the code-base `emcee` [].

Figure 14 is an example corner plot showing the probability distributions for q, χ_1, and χ_2 obtained by sampling Eq. (2) with two different mode combinations. The three plots on the diagonal show histograms of the probability distribution for individual parameters, and the remaining three are two dimensional probability densities. The true system parameters are represented by black lines; here, $q = 2$, $\chi_1 = 0.1$, and $\chi_2 = 0.1$. In the plot for q and χ_1, when we fit with only the fundamental $(2, 2, 0)$ mode, there is a visible track where there are many different combinations of q and χ_1 with similar probabili-

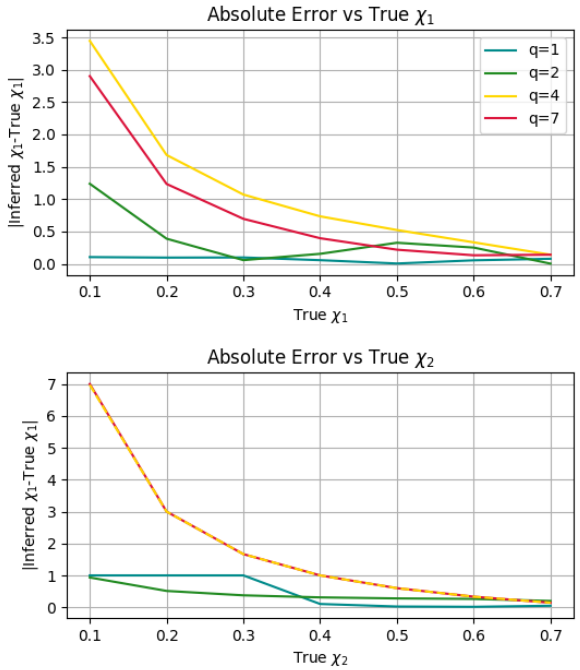

FIG. 11: Absolute error of the fits shown in Figure 10. These trends suggest that error decreases for higher χ_1 and χ_2, but this is likely due to the least squares method failing and getting stuck on a boundary. As spin values increase and approach the upper boundary, error decreases although the fit is not performing properly.

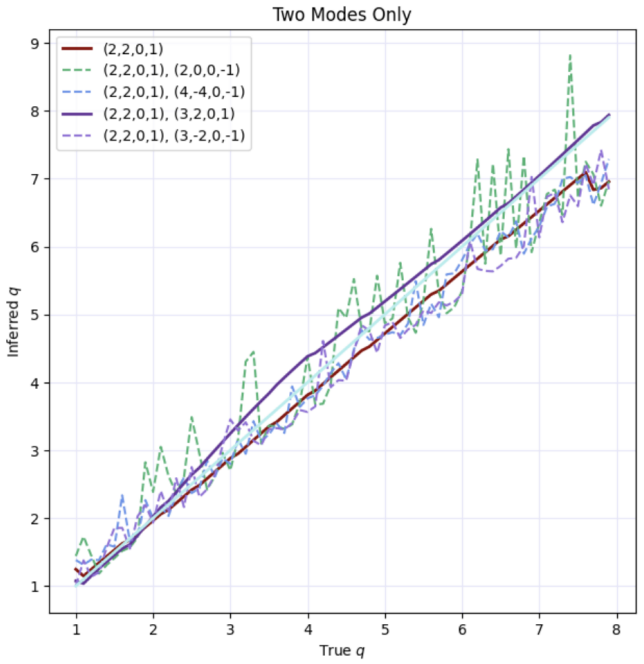

FIG. 12: The turquoise line is a perfect diagonal, which is what the data would show if the inverted surrogate was working perfectly and predicting the true parameters exactly. Dashed lines represent a combination that contains a mirror mode. In general, fitting with two modes does not make the fit noticeably better than using just the fundamental.

ties. These degeneracies cannot be mapped with point estimates, which is a benefit of using sampling and a likelihood function. This particular degeneracy is mitigated by fitting with the $(3, 3, 0)$ mode in addition to the $(2, 2, 0)$ (turquoise).

In general, the quality of fit increases substantially as more modes are used. In practice, when observing a real GW ringdown signal, only a finite number of modes that are actually observable. In most cases, only one or two modes can be resolved. This motivates testing the quality of fits using single modes or combinations of two modes (including the fundamental). In reality, what modes are observed cannot be chosen, but it is useful to know how good these fits are for certain combinations so when those combinations are observed, we have a reference for knowing how much fits using a certain combination can be trusted for parameter estimation.

Using this method, we tested four different configurations of extreme mass ratio and spin: low q and χ, high q and low χ, low q and high χ, and high q and χ. In these tests, we set $\chi_1 = \chi_2$. For the sake of observing the functionality of an inverted surrogate in unobstructed conditions, we are not using `qnmfits` in the sampling method, but rather the amplitudes and phases from the ringdown

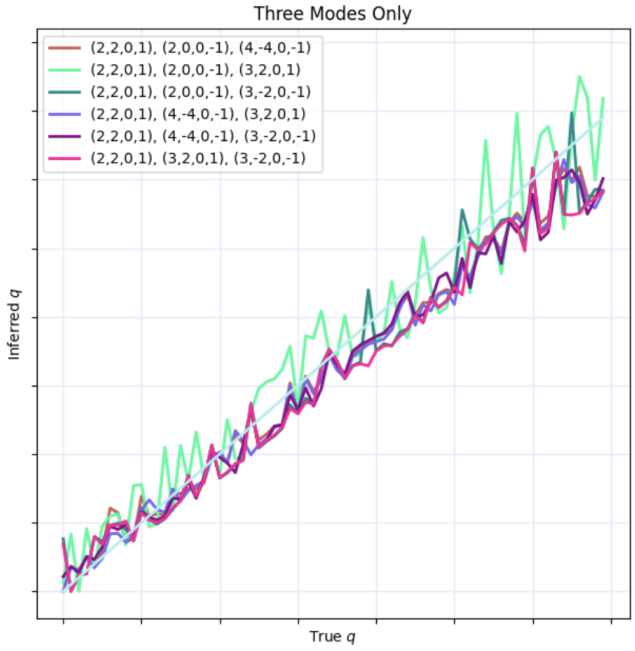

FIG. 13: Compared to the data plotted in Figure 12, fitting with three modes does not improve results.

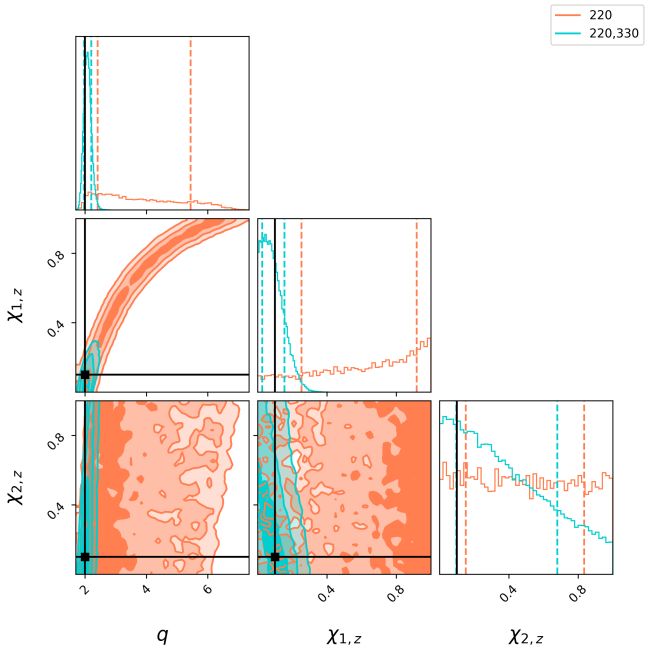

FIG. 14: This corner plot shows some odd confidence interval for the estimated parameters q, χ_1, and χ_2. The coral corresponds to fitting for these parameters using just the fundamental mode, the $(2,2,0)$, and the turquoise is the fit using the $(3,3,0)$ mode in addition to the fundamental. The corner plot with q and χ_1 shows a degeneracy using just the fundamental for fitting, but this degeneracy is greatly mitigated by using additional modes.

surrogate q8_3dAl.

Figure 15 shows the results for the parameter estimation in the low mass ratio and low spin magnitude case. Parameters are poorly constrained in the one mode fit, as seen in Figure 15a. The same degeneracy shown in Figure 14 can be seen in the $(2,2,0)$ mode, but none others. In Figure 15b, fitting with the fundamental as well as an additional mode improves the constraints greatly, but all combinations tend toward the degeneracy seen in $(2,2,0)$. Aside from the fit with all modes (pink), the combinations $(2,2,0)+(3,3,0)$ and $(2,2,0)+(4,4,0)$ mitigate this degeneracy most effectively in this mass-spin configuration. We show results for the remaining three mass-spin configurations in Appendix A.

Increasing the number of modes tends to generate a more precise fit for component masses and spins. However, different single or two-mode combinations are required to better-fit different mass-spin configurations. Notably, the combination of the (2,2,0)+(3,2,0) tends to be the most informative for constraining mass ratio and the spin of the primary (more massive) black hole, while the (2,1,0) mode significantly helps constrain the spin of the secondary (less massive) for the more-equal-mass cases. As expected, the fit using all modes was the best in every case. The $(4,4,0)$ mode was also notable for inferring mass ratio in lower mass ratio cases.

V. CONCLUSIONS AND FUTURE WORK

In the case of least squares fitting, it became clear that using least squares to invert a surrogate is not an effective method when fitting for parameters with many degrees of freedom. It is also not optimal for exposing degeneracies. Although this method was not the most effective, it laid a solid framework for following the same process of ringdown surrogate inversion using sampling. It also directed attention to specific problems to investigate with sampling, such as getting stuck on boundaries and worse fits for including more modes. These things could all easily be investigated by sampling.

As expected, including more modes in the fits when using sampling greatly improved constraints on parameters. In particular, the $(3,2,0)$ and $(2,1,0)$ modes in combination with the fundamental provided reasonable estimates for various systems. In general, switching to the sampling method made clear any degeneracies present and allowed us to see how likely an estimate was, as opposed to just a point estimate returned by least squares. In this project, the σ we chose corresponded to a high SNR and was the same for all systems. This motivates investigating what SNR is required to constrain different parameters. Even in the high SNR case, χ_2 was poorly constrained.

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.

VI. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation Research Experience for Undergraduates (NSF REU) program, the LIGO Laboratory Summer Undergraduate Research Fellowship program (NSF LIGO), and the California Institute of Technology Student-Faculty Programs.

Appendix A: Corner Plots for Various Cases

Figure 17 shows the results for the system with a high mass ratio and low spins. In a system with such discrepant masses, we expect χ_1 to dominate and therefore poorer constraints on χ_2. This is apparent in both the one and two mode fits. The degeneracy in q and χ_1 is present in the $(2,2,0)$ mode. Since the fundamental is the main mode observed, combinations of two modes also follow the degeneracy.

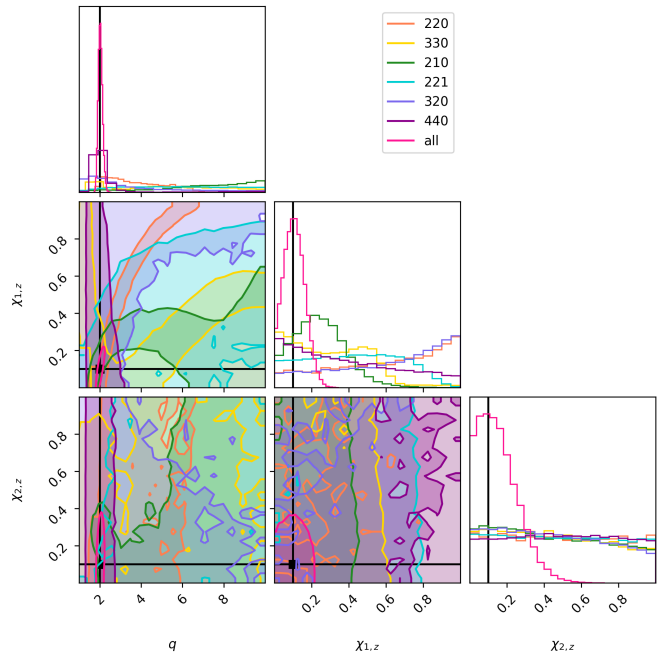

(a) Spins are not very well constrained for any single mode fit. Mass ratio is relatively constrained in few modes, particularly the $(4, 4, 0)$ and $(3, 3, 0)$ modes.

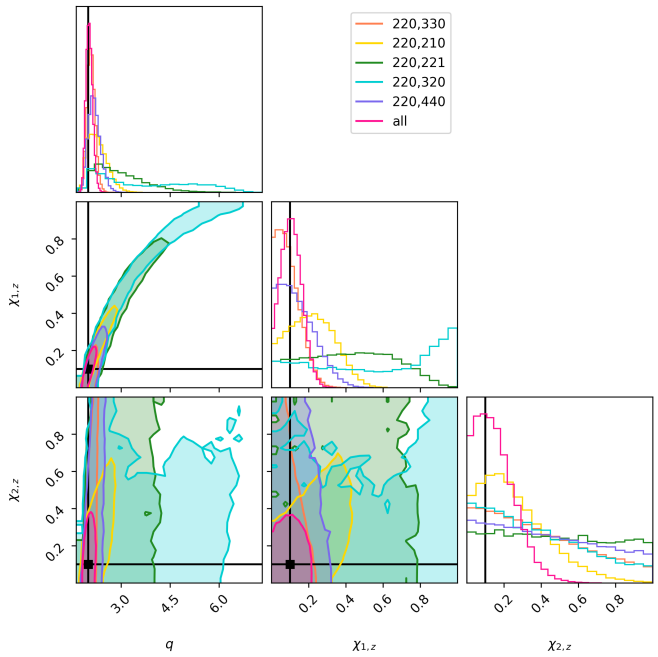

(b) For the low mass ratio and low spin case, the degeneracy between q and χ_1 is apparent in every combination of modes, but substantially less so in $(3, 3, 0)$ and $(4, 4, 0)$.

FIG. 15: Corner plots showing 90% confidence interval for configuration of $q = 2$, $\chi_1 = \chi_2 = 0.1$ for one and two mode fits.

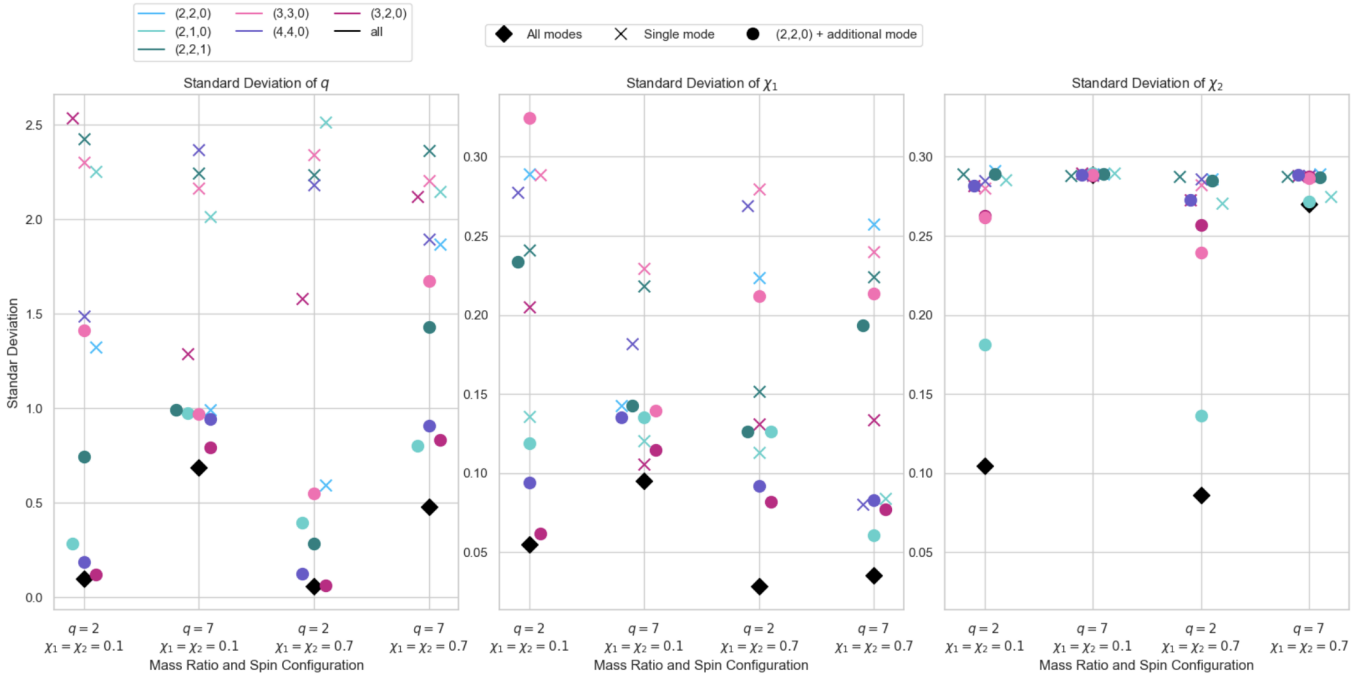

FIG. 16: Standard deviation of probability distributions for q, χ_1, and χ_2. Modes with $\ell = 2$ are represented in a shade of blue, $\ell = 3$ is in pink, and $\ell = 4$ is in purple. Fit using all modes is plotted as a black diamond. Especially in the cases of q and χ_1, fits with two modes have a smaller standard deviation than fits with one mode. χ_2 is poorly constrained in all different mass ratio and spin configurations that were tested.

The system with low mass ratio and high spins is shown in Figure 18. Unlike high mass ratio systems, χ_2 is better constrained in this system, particularly in the fit including the $(2, 2, 0)$ and $(2, 1, 0)$ modes. This is expected when the BHs are closer to equal mass.

In the system with a high mass ratio and high spins, shown in Figure 19, χ_2 is also poorly constrained. This is for the same reason as the system shown in Figure 17. The degeneracy is present in this case also.

[1] H. Zhu *et al.*, Phys. Rev. D **111**, 064052 (2025), arXiv:2312.08588 [gr-qc].
[2] J. Aasi *et al.* (LIGO Scientific), Class. Quant. Grav. **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].
[3] F. Acernese *et al.* (VIRGO), Class. Quant. Grav. **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
[4] T. Akutsu *et al.* (KAGRA), PTEP **2021**, 05A101 (2021), arXiv:2005.05574 [physics.ins-det].
[5] A. Buikema *et al.*, Phys. Rev. D **102**, 062003 (2020), arXiv:2008.01301 [astro-ph.IM].
[6] S. A. Teukolsky, Astrophys. J. **185**, 635 (1973).
[7] S. Chandrasekhar and S. Detweiler, Proceedings of the Royal Society of London Series A **344**, 441 (1975).
[8] K. D. Kokkotas and B. G. Schmidt, Living Rev. Rel. **2**, 2 (1999), arXiv:gr-qc/9909058.
[9] S. Detweiler, Astrophys. J. **239**, 292 (1980).
[10] 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.
[11] M. H.-Y. Cheung, E. Berti, V. Baibhav, and R. Cotesta, Physical Review D **109** (2024), 10.1103/physrevd.109.044069.
[12] R. Abbott *et al.* (KAGRA, VIRGO, LIGO Scientific), Phys. Rev. X **13**, 041039 (2023), arXiv:2111.03606 [gr-qc].
[13] 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].
[14] I. Mandel and A. Farmer, Phys. Rept. **955**, 1 (2022), arXiv:1806.05820 [astro-ph.HE].
[15] "GW231123: a Binary Black Hole Merger with Total Mass 190-265 $M_\odot$," (2025), arXiv:2507.08219 [astro-ph.HE].
[16] R. Abbott *et al.*, Physical Review Letters **125** (2020), 10.1103/physrevlett.125.101102.
[17] S. J. Miller, M. Isi, K. Chatziioannou, V. Varma, and I. Mandel, Physical Review D **109** (2024), 10.1103/physrevd.109.024024.
[18] 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.
[19] E. Payne, S. Hourihane, J. Golomb, R. Udall, D. Davis, and K. Chatziioannou, Physical Review D **106** (2022), 10.1103/physrevd.106.104017.
[20] M. Giesler, S. Ma, K. Mitman, N. Oshita, S. A. Teukolsky, M. Boyle, N. Deppe, L. E. Kidder, J. Moxon, K. C. Nelli, H. P. Pfeiffer, M. A. Scheel, W. Throwe, and N. L. Vu, Physical Review D **111** (2025), 10.1103/physrevd.111.084041.
[21] L. Magaña Zertuche, L. Gao, E. Finch, and G. B. Cook, (2025), arXiv:2502.03155 [gr-qc].
[22] 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.
[23] V. Varma, D. Gerosa, L. C. Stein, F. Hébert, and H. Zhang, Physical Review Letters **122** (2019), 10.1103/physrevlett.122.011101.
[24] 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].
[25] E. Finch, *Black-hole Ringdown*, Ph.D. thesis, University of Birmingham (2023).

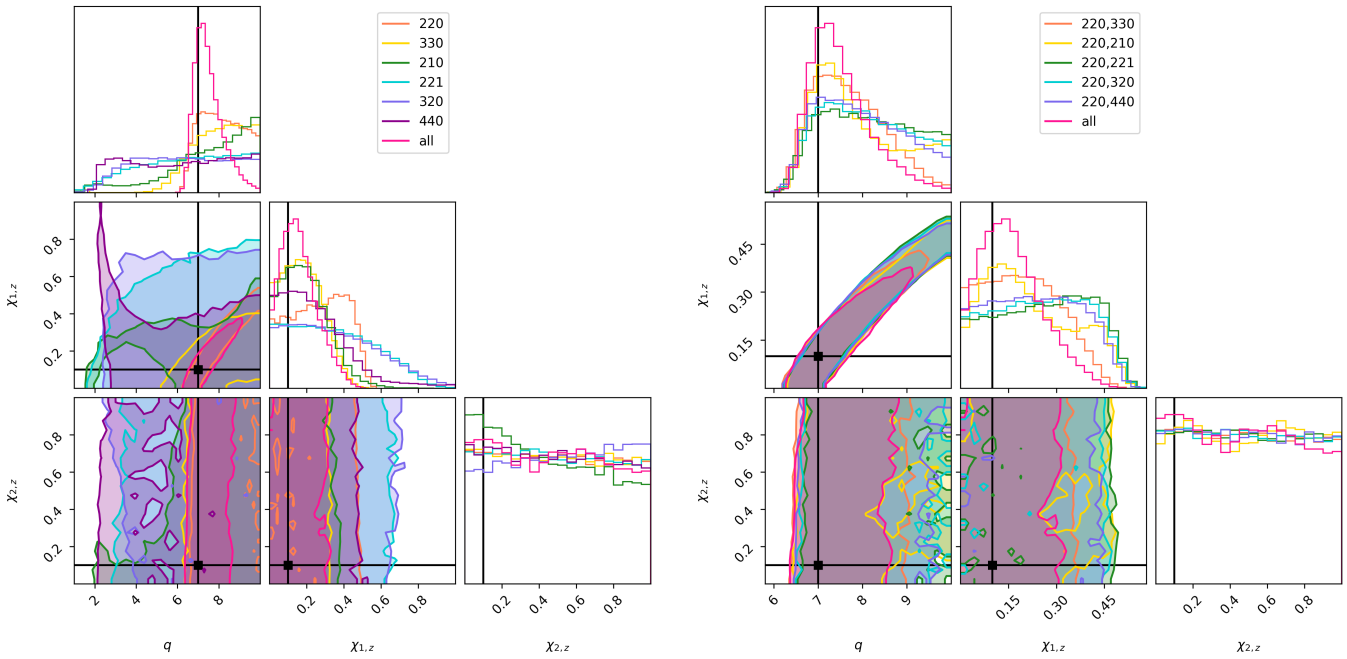

(a) (b)

FIG. 17: Corner plots showing 90% confidence interval for configuration of $q = 7$, $\chi_1 = \chi_2 = 0.1$ for one and two mode fits.

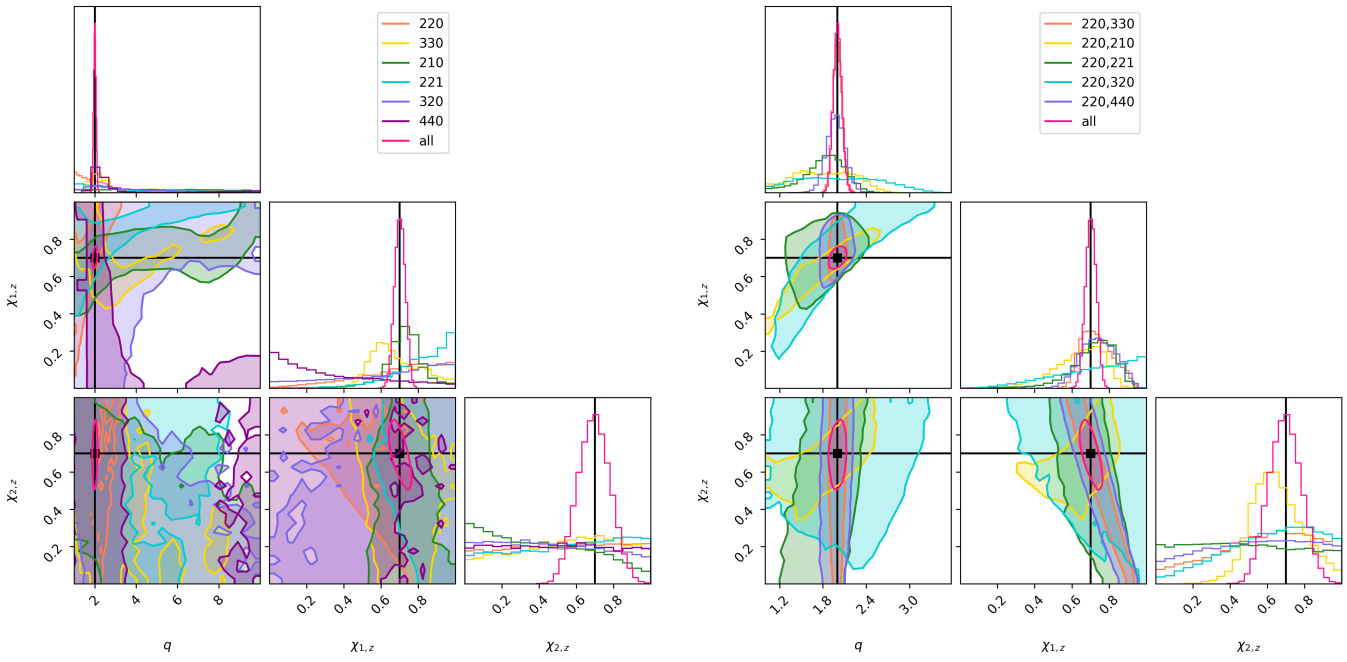

(a) (b)

FIG. 18: Corner plots showing 90% confidence interval for configuration of $q = 2$, $\chi_1 = \chi_2 = 0.7$ for one and two mode fits.

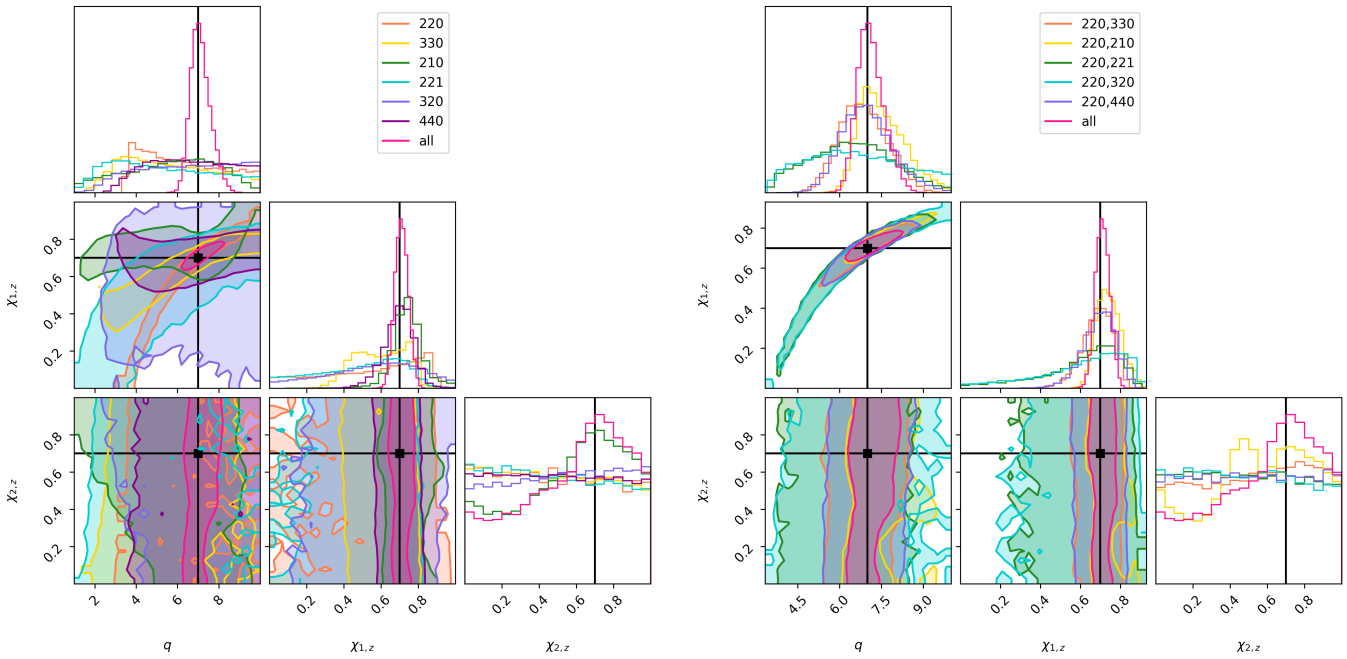

(a) (b)

FIG. 19: Corner plots showing 90% confidence interval for configuration of $q = 7$, $\chi_1 = \chi_2 = 0.7$ for one and two mode fits.