LIGO-T1600152

LIGO SURF Final Report: Studying the effects of higher-order modes on the parameter estimation of precessing low-mass binary black holes

Rupini V. Kamat,[1] Patricia Schmidt,[2] Rory Smith,[2] and Yanbei Chen[2]

[1]*LIGO SURF Student, California Institute of Technology*
[2]*LIGO SURF Mentors, California Institute of Technology*

(Dated: October 20, 2016)

Precessing low-mass binary black holes are promising detection candidates in Advanced LIGO's upcoming observation runs. The gravitational radiation emitted by these systems is distributed across various gravitational wave modes. While the waveforms of symmetric, non-precessing systems can be well described by lower-order modes alone, higher-order modes can carry a significant amount of energy, and thus information, which may be particularly relevant for precessing systems. During parameter estimation, model waveforms are compared against detected signals to recover the parameters of the observed systems. Often, these models only include a few dominant modes, potentially neglecting important information about the source that may be encoded in other modes of its gravitational waveform. This project aims to investigate whether the inclusion of higher-order modes in model waveforms has a significant effect on the accuracy of parameter estimation for precessing, low-mass binary black holes systems. The methods we employ include a Fisher matrix analysis as well as a maximum likelihood estimator. While our analysis is limited to inspiral waveforms due to a lack of availability of complete inspiral-merger-ringdown waveforms with higher-order modes for precessing systems, future research can extend this analysis to merger and ringdown waveforms, where higher-order modes are expected to be even more significant.

I. INTRODUCTION

Advanced LIGO, the Laser Interferometer Gravitational-Wave Observatory [1], has made huge breakthroughs in the field of gravitational wave (GW) astronomy with the first ever direct detections of gravitational waves [2, 3]. For both observations, the GWs were generated by merging binary black holes (BBH) with a total mass of $\sim 70M_{\odot}$ respectively of $\sim 22M_{\odot}$.

As theoretically predicted, both GW signals detected by Advanced LIGO evolved in three stages: the inspiral, the merger, and the ringdown. The inspiral stage of a black hole binary's evolution is characterized by the large separation of the binary as well as orbital speeds much smaller than the speed of light. Because the motion of the binary is non-relativistic during the inspiral phase, the binary dynamics and the inspiral waveform can be calculated by expanding the Einstein field equations in terms of v/c using the post-Newtonian formalism (see e.g., [4] for a comprehensive overview). However, as gravitational waves carry angular momentum and orbital energy away from the system, the orbit of the binary decays, the separation of the binary shrinks, and the speeds of the black holes approach the speed of light. At these speeds, the non-linearity of General Relativity becomes significant, and the waveform must be calculated using the numerical solution to the relativistic two-body problem, or numerical relativity. Finally, after the binary black holes merge into a single black hole, the system enters its final ringdown stage, during which quasinormal modes are emitted by the perturbed remnant black hole until it settles down to the stationary Kerr solution for spinning black holes. The waveform emitted during the ringdown stage of the binary's evolution must either be calculated using numerical relativity or via perturbation theory (see e.g., [5] for details).

Compact binary coalescences, such as the merger of two black holes, are among the most promising candidates for present and future observations by Advanced LIGO. The gravitational waveforms of these objects encode a tremendous amount of information about each system that would otherwise be inaccessible, as systems like binary black holes are not thought to be observable through electromagnetic radiation.

Thus, the gravitational waveforms of these systems can help us answer astrophysically interesting questions, for example how compact binaries are formed, what their masses and spins are, how they evolve in time, and how these types of systems are distributed in space. However, while LIGO has now proven its capability of detecting gravitational waves from these astronomical sources, our ability to use observed gravitational waveforms to extract the properties of the systems they originated from is still limited. This may be due to a number of factors, including the fact that many of the currently used waveform models use various approximations to the full solutions of General Relativity, amongst them for example the quadrupole approximation [23].

This approximation may lead to inaccuracies in parameter estimation and parameter degeneracies that prevent us from more accurately determining the physical parameters of an observed system. Such degeneracies include the spin-mass ratio degeneracy, the orientation-distance degeneracy, and many more. For example, Figs. 3 and 4 in [6] illustrate the spin-mass degeneracy for binary black holes.

Thus, in order for us to use gravitational-wave observations to confidently answer astronomically interesting questions, we must also improve our ability to probe the properties of the systems we observe. One potential way to improve our parameter estimation abilities is through the inclusion of higher modes in our theoretical waveforms [7, 8]. While the quadrupole approximation may be sufficient for highly symmetric binaries, binary systems that contain asymmetries, such as unequal mass ratios and unequal spins, emit a significant amount of gravitational-wave energy in higher modes, as, for example, shown in Figures 1 and 2 in [9]. Thus, higher modes carry a significant amount of information about the source system, and the inclusion of higher modes in template waveforms has the potential to allow us to recover large amounts of information from an observed waveform.

II. PROJECT DESCRIPTION

We propose that the inclusion of higher-order modes in our template waveforms has the potential to break parameter degeneracies and decrease parameter uncertainties, thus allowing us to more accurately recover the properties of an observed system through its detected gravitational waveform. Similar work was done by O'Shaughnessy et al. [7, 10], who analyzed the inspiral waveforms of select black hole–neutron star binary systems to investigate the possibility of breaking some of these parameter degeneracies. O'Shaughnessy et al. found that for the fiducial system they analyzed, the inclusion of higher order modes broke degeneracies regarding the direction of the system's angular momentum, but otherwise had little impact on the estimation of the systems' parameters.

The goal of this project is to extend this analysis to some precessing binary black hole systems to systematically assess what impact higher order modes have on parameter estimation for the systems we considered. Specifically, we are focusing on precessing, low-mass binary black holes of masses m_1 and m_2 with various spin configurations and a total mass of less than $M = m_1 + m_2 \leq 12M_\odot$. The reason for our restriction on the total mass of the system is because we only considered inspiral waveforms, and for black hole binaries with a total mass $\geq 12M_\odot$, the merger phase appears in the LIGO band. Using pure inspiral waveforms for such system can result in biased results during parameter estimation [11], hence the restricted analysis. While it is expected that the effects of higher modes are more significant during the merger phase than during the inspiral phase, we do not currently have NR models for the merger waveforms of precessing binary black holes that would enable us to study these effects, which is why our analysis was limited to inspiral waveforms.

We evolve the binary in a source frame as shown in Fig. 1, and compute the gravitational-wave modes following the Eqn. given in [12]. We define the inclination of the binary w.r.t. the line-of-sight by θ and its azimuthal orientation by ϕ. These two angles define the orientation of the system and are measured with respect to the total angular momentum of the system $\vec{J}$, as its direction, $\hat{J}$, remains almost constant throughout the inspiral. Contrary to the binary orientation, the dimensionless spin vectors $\vec{\chi}_1$ and $\vec{\chi}_2$ are

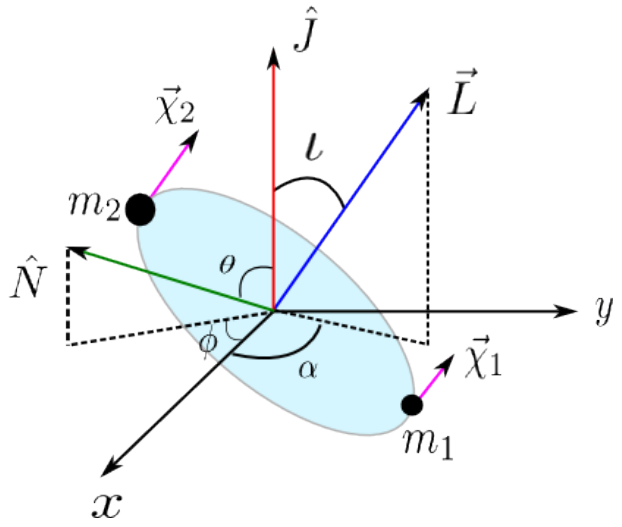

FIG. 1: The source frame for a precessing compact binary system (adapted from [14]). $\hat{J}$ indicates the direction of the total angular momentum of the system, $\vec{L}$ indicates the orbital angular momentum, $\vec{\chi}_1$ and $\vec{\chi}_2$ are the dimensionless spin vectors of each black hole of masses m_1 and m_2, and $\hat{N}$ is the direction of the observer (line-of-sight).

defined relative to the instantaneous orbital angular momentum, $\hat{L}$, which precesses around $\hat{J}$ throughout the evolution. The x- and y-components of the spin vectors define the spin components that lie in the instantaneous orbital plane (indicated by the blue ellipse in Fig. 1), which are the source of the precession of the binary. The z-components of the spins are the components parallel to $\hat{L}$, i.e., perpendicular to the instantaneous orbital plane, which mainly influence the inspiral rate of the binary (see e.g. [13] for further details). The precession of the orbital plane is reflected in the time-dependence of the orientation of $\hat{L}$, which is determined by the two polar angles ι and α, where ι describes the opening angle of the precession cone and α the rotation of $\hat{L}$ around $\hat{J}$.

The analysis we perform is an error analysis, as described in Sections III A and III B, which uses the Fisher information matrix as a tool to estimate both the statistical errors, or the error due to noise, and the systematic errors, or the error due to an incorrect or incomplete waveform model, in the various recovered parameters of a given system. Our aim in this analysis is to see how the inclusion of higher order modes effects both of these errors, in order to understand if any significant parameter bias occurs from using higher-mode exclusive waveform templates for certain systems.

We have completed a statistical error analysis for one aligned spin case and one precessing case. From this, we have seen that the effect of higher order modes on the statistical errors of each parameter varies, but that in general the inclusion of higher modes decreases parameter uncertainties. In the aligned spin case, this decrease in parameter uncertainty is most significant for the inclination angle θ, while in the precessing case it is most noticeable for the spin parameters, in particular for the out-of-plane spin components of both black holes. The effect of higher order modes on parameter correlations however varies. While increased parameter correlations are found in some cases, we also find decreased parameter correlations in others. The significance of our statistical error analysis can only be assessed once a systematic error analysis has also been completed. We have made some initial steps towards performing a systematic error analysis through the maximum likelihood estimator, defined in Section III B, as shown in Section IV D, but further work is required to fully assess which error source is the dominant one. Our work on the maximum likelihood estimator has not yet returned satisfactory results, but it has enabled us to better understand some of the general features of the likelihood function, such as its multi-modality with variation of the individual spin magnitudes, its relatively strong sensitivity to variations in the mass ratio, and the dependence of its sensitivity to changes in inclination on polarization angle.

Once a systematic error analysis has been completed for the cases analyzed in Sections IV C 1 and IV C 2 and the entire error analysis has been expanded to a larger range of precessing systems, we can begin to understand whether or not the inclusion of higher order modes in template waveforms results in any significant improvement in LIGO's ability to recover the parameters of the systems it observes, thus improving the ability of LIGO's current and future detections to answer various questions in astrophysics.

III. METHODS

First, we chose a set of precessing low-mass black hole binary systems to analyze. In addition to a total mass constraint of $\leq 12M_{\odot}$, we constructed our cases by choosing systems we expect to emit a significant amount of radiation in higher modes. From a series of match calculations (Section IV B), we found that the systems for which higher order modes carry a significant amount of information have parameters in the neighborhood of mass ratio $q = 6$, dimensionless spin vectors $\vec{\chi}_1 = \vec{\chi}_2 = (0.6,\ 0,\ 0)$, and an inclination angle [24] of $\theta \approx \pi/2$. We then generated the gravitational waveforms produced by our chosen systems, both with and without higher modes. For our higher modes analysis, we used the post-Newtonian inspiral waveforms for precessing binary black holes given in Appendix A in Ref. [14].

We then analyzed the resulting simulated waveforms using a Fisher information matrix analysis for select cases. We used Fisher analysis to get a best estimate of the improvement in parameter uncertainty, if any, that results from the inclusion of higher waveforms.

A. Fisher Matrix Analysis

Due to the high computational cost of a Bayesian parameter estimation analysis with LALInference, we decided to focus our analysis of the effects of higher modes on parameter estimation with a Fisher matrix analysis. The Fisher information matrix is a tool that allows one to measure the amount of information that a given set of data provides about some unknown parameter that can be used to model the data set. It is defined as follows:

$$\Gamma = -[\nabla\nabla \ln(L(s\,|\,\vec{\lambda}))], \tag{1}$$

where $L(s\,|\,\vec{\lambda})$ is, in Bayesian statistical analysis, the likelihood of the observed data, or signal s, given the parameters $\vec{\lambda}$. The inverse of the Fisher matrix gives the covariance matrix of the posterior probability distribution, where the *statistical uncertainty* of a measured parameter λ_i due to noise is given by

$$(\Delta\lambda_i)_{\text{stat}} = \sqrt{\Gamma^{-1}_{ii}}, \tag{2}$$

and the remaining entries in the matrix $\Gamma_{ij} = \Gamma_{ji}$ give the correlation between two parameters λ_i and λ_j. In the case of a gravitational wave signal, we can assume that an observed signal $s \in \mathbb{R}$ contains some noise n and a gravitational wave signal $h_0(\vec{\lambda}_{\text{true}})$, where $\vec{\lambda}_{\text{true}}$ are the true parameters of the observed system

$$s = n + h_0(\vec{\lambda}_{\text{true}}). \tag{3}$$

Thus, if we model the signal as $h(\vec{\lambda})$, we can represent the likelihood function as the following:

$$L(s|\vec{\lambda}) \propto e^{-\langle s-h(\vec{\lambda}),s-h(\vec{\lambda})\rangle/2}, \tag{4}$$

where $\langle \bar{h}, \bar{g} \rangle$ is the overlap, or the noise-weighted inner product, of two real functions $\bar{h}(t)$ and $\bar{g}(t)$, as defined below:

$$\langle \bar{h}, \bar{g} \rangle = \mathrm{Re} \int_{-\infty}^{\infty} \frac{\tilde{h}(f)\tilde{g}^*(f)}{S_n(|f|)} df e^{2i(\psi-\sigma)} + \mathrm{Re} \int_{-\infty}^{\infty} \frac{\tilde{h}^*(-f)\tilde{g}^*(f)}{S_n(|f|)} df e^{-2i(\psi+\sigma)}, \tag{5}$$

where $h, g \in \mathbb{C}$ are the *complex* strains, $\bar{h} = \mathrm{Re}[h e^{2i\psi}]$ and $\bar{g} = \mathrm{Re}[g e^{2i\sigma}]$, where the polarization angles are ψ and σ respectively. The Fourier transform of h is denoted by $\tilde{h}$ and the complex conjugate by h^*. We use the one-sided noise power spectral density (PSD) $S_n(|f|)$ for the advanced, zero-detuned LIGO detector. The algorithm used to calculate the match is discussed in Appendix A.
From Eq.(5) we define the norm of a waveform as

$$||h|| = \sqrt{\langle h, h \rangle}, \tag{6}$$

and the normalized waveform $\hat{h}$ as

$$\hat{h}(\vec{\lambda}) = \frac{h(\vec{\lambda})}{||h||}. \tag{7}$$

If we assume that the noise in a gravitational wave signal is Gaussian with a mean of zero and that our prior probability distribution for the data is a flat distribution, the elements of the Fisher information matrix are given by:

$$\Gamma^{ij}(\vec{\lambda}) = \left\langle \frac{\partial h}{\partial \lambda_i}(\vec{\lambda}), \frac{\partial h}{\partial \lambda_j}(\vec{\lambda}) \right\rangle. \tag{8}$$

The Fisher matrix scales quadratically with the signal-to-noise ratio (SNR), which is a detection statistic that quantifies signal strength and is defined as follows:

$$\rho = \max_{\lambda} \langle s, \hat{h}(\lambda) \rangle, \tag{9}$$

where s is the observed signal and $\hat{h}$ is the normalized template waveform used to model the data.
The optimal SNR for a given gravitational waveform h is given by

$$\rho_{opt} = \langle h, h \rangle. \tag{10}$$

Thus, we can define an SNR-independent, normalized Fisher matrix as shown in Eq. (23) in Ref. [15]:

$$\hat{\Gamma}^{ij}(\vec{\lambda}) = \frac{1}{\rho_{opt}^2} \Gamma^{ij}(\lambda). \tag{11}$$

We developed code to calculate the Fisher matrix for a given gravitational waveform in `Mathematica`. Our implementation uses finite differencing methods to calculate the derivative of the waveform with respect to a given parameter, as shown in Eq. (12).

$$\frac{\partial h}{\partial \lambda_i}(\vec{\lambda}) \approx \frac{h(\lambda_1, \lambda_2, ..., \lambda_i + \delta\lambda_i, ...\lambda_n) - h(\lambda_1, \lambda_2, ..., \lambda_i - \delta\lambda_i, ..., \lambda_n)}{2\,\delta\lambda_i} \tag{12}$$

We then calculated the overlap, as defined in Eq. (5), between derivatives of the waveform using the algorithm outlined in Appendix A to produce the Fisher information matrix of the waveform h. We then normalized the Fisher matrix, using Eq. (7), and assumed an SNR of $\rho = 15$ (i.e. multiplied the normalized

Fisher matrix by ρ^2), as done by Cho et al. in Ref. [15]. Inverting this Fisher information matrix gives us the covariance matrix of the system, which has the following form:

$$C_{ij} = \begin{cases} \sigma_{ij} & i \neq j \\ \sigma_i^2 & i = j \end{cases}, \tag{13}$$

where σ_{ij} is the covariance between two parameters λ_i and λ_j, and σ_i is the statistical error of the parameter λ_i, as stated in Eq. (2). We can use the information in the covariance matrix of a given system's waveform to construct two dimensional confidence ellipses that show visually both how well each individual parameter is constrained, as well as how various parameters are correlated with one another. We do so by first constructing a set of 2x2 matrices of the form:

$$c(\lambda_i, \lambda_j) = \begin{pmatrix} \sigma_i^2 & \sigma_{ij} \\ \sigma_{ji} & \sigma_j^2 \end{pmatrix}, \tag{14}$$

Then, we can construct the 1σ and 2σ confidence ellipses, which correspond to confidence levels of 68.3% and 95.4% respectively. The ellipse for each confidence level represents the region in which we can state that for repeated measurements, the probability that the true values for λ_i and λ_j lie within the ellipse is given by the confidence level. The parameters of a confidence ellipse is given by:

$$a^2 = \alpha^2 \left(\frac{\sigma_m^2 + \sigma_b^2}{2} + \sqrt{\frac{(\sigma_m^2 - \sigma_b^2)^2}{4} + \sigma_{mb}^2} \right), \tag{15}$$

$$b^2 = \alpha^2 \left(\frac{\sigma_m^2 + \sigma_b^2}{2} - \sqrt{\frac{(\sigma_m^2 - \sigma_b^2)^2}{4} + \sigma_{mb}^2} \right), \tag{16}$$

and

$$\tan(2\theta) = \frac{2\sigma_{mb}}{\sigma_m^2 - \sigma_b^2}, \tag{17}$$

where a and b are the axes of the ellipse, θ is the counterclockwise rotation of the ellipse, and α changes based on the confidence interval the ellipse to represents. For a 1σ, or 68.3%, confidence level, $\alpha = 1.52$. For a 2σ, or 95.4%, confidence level, $\alpha = 2.48$ [16].

B. Maximum Likelihood Estimator

The *maximum likelihood estimate* $\vec{\lambda}'$ is given by the set of parameters that maximize the likelihood function $L(s|\vec{\lambda})$ given in Eq. (4). The maximum likelihood estimator can be found either by maximizing the likelihood $L(s|\vec{\lambda})$ over $\vec{\lambda}$ or by maximizing the log-likelihood function, $\log[L(s|\vec{\lambda})]$ over $\vec{\lambda}$. As stated in Ref. [17], the maximization of the log-likelihood function is equivalent to maximizing the *fitting factor* (FF) between the true waveform $h_0(\vec{\lambda})$ of the observed system and some template waveform $h(\vec{\lambda}')$ used to model the system at some fixed polarization angle. The fitting factor is the match between two real valued waveforms maximized over *all* parameters:

$$\mathrm{FF} = \max_{\vec{\lambda}'} \frac{\langle \bar{h}_0(t,\vec{\lambda}),\, \bar{h}(t,\vec{\lambda}')\rangle}{||h_0||||h||} \tag{18}$$

If h is an accurate model for the system's true waveform h_0, the posterior probability distribution, or the probability distribution of parameters that would result from a thorough parameter estimation analysis, can then be approximately described by the *maximum likelihood estimate* $\vec{\lambda}'$, and the parameter uncertainties given by Eq. (2). If, however, the model waveform h deviates from the true waveform h_0, such that

$$\delta h(t;\vec{\lambda}) = h_0(t;\vec{\lambda}) - h(t;\vec{\lambda}) \neq 0, \tag{19}$$

then the maximum likelihood estimate $\vec{\lambda}'$ for the true parameters has some *systematic* error that is not described by Eq. (2), which quantifies only the *statistical* error due to the noise. Using the derivation shown in Section 2.3.3 of Ref. [18], we find that the systematic error that arises from an incomplete waveform model is given by:

$$(\Delta\lambda_i)_{\mathrm{sys}} = [\Gamma^{-1}(\vec{\lambda}')]_{ij} \left\langle \delta h_0(\vec{\lambda}'), \frac{\partial h}{\partial \lambda_j}(\vec{\lambda}') \right\rangle. \tag{20}$$

The total error in the estimate for a single parameter is given by the sum of the statistical error and the systematic error. Because, as previously mentioned, Γ_{ij} scales quadratically with SNR, we have that $(\Delta\lambda_i)_{\mathrm{stat}}$ increases linearly with decreasing SNR, while $(\Delta\lambda_i)_{\mathrm{sys}}$ is independent of the SNR. Thus, systematic error is the dominant source of parameter estimation error at high SNR, but becomes negligible compared to statistical error at low SNR.

Thus, we can define some SNR-dependent "indistinguishability criterion" to determine whether or not two different waveforms h and h_0 can be considered *indistinguishable* from one another. Following Ref. [19], we call two waveforms indistinguishable if $(\Delta\lambda)_{\mathrm{sys}} < (\Delta\lambda)_{\mathrm{stat}}$. Thus, we can use Fisher matrices to analyze whether or not the quadrupole and $\ell = 2$ waveforms of a system are "indistinguishable" from the higher-mode-inclusive waveform of the same system. This enables us to determine if there exist areas in the parameter space for low mass, precessing binary black holes where significant systematic error arises from using only quadrupole or $\ell = 2$ waveforms during parameter estimation.

C. Bayesian parameter estimation

In addition to the Fisher matrix analysis, in the future it may also be beneficial employ Bayesian parameter estimation algorithms using the LALInference software library.

In general terms, the process of Bayesian analysis is as follows. First, a probability distribution, known as the *prior distribution* and represented as $p(\vec{\lambda})$, is constructed for each set of parameters $\vec{\lambda}$ in the binary parameter space. Each set of parameters in the binary parameter space includes values like component masses, spins, position on the sky, distance to the source, orientation of the binary, and many more.

After a prior distribution has been constructed, the likelihood $L(s\,|\,\vec{\lambda})$ of the observed signal s given the parameters $\vec{\lambda}$ is calculated using Eq. 4. Then, we multiply the likelihood function by the prior distribution $p(\vec{\lambda})$ and normalize by $p(s)$, the probability of the signal independent of the distribution of the parameters, also known as the evidence. This gives us the posterior probability distribution, $p(\vec{\lambda}\,|\,s)$, as described by Bayes' Theorem:

$$p(\vec{\lambda}\,|\,s) = \frac{L(s\,|\,\vec{\lambda})\; p(\vec{\lambda})}{p(s)}. \tag{21}$$

The parameter estimate is then given by the median value of the posterior distribution, and the error is given by the values that correspond to the 90% credible interval, or the parameter values that enclose 90% of the probability of the distribution.

Using the algorithms in the LALInference software library, we would first generate a gravitational-wave signal using the post-Newtonian (PN) waveform approximation *SpinTaylorT4*. The signal would then be then injected into simulated, recolored Gaussian noise at a given SNR, and Bayesian statistics, as described above, would be used to recover the parameters of the source.

Since we are not able to directly access the $h_{\ell m}$-modes, described below in Section IV A, of the waveform through LALInference, instead we would perform a study where we truncate the post-Newtonian expansion of the waveform $h(t)$ at different PN orders v^n, where v is the PN expansion parameter:

$$h(t) = A(t, v^m) e^{i\phi(t, v^n)}, \tag{22}$$

where $A(t, v^m)$ is the GW amplitude at a specific PN order m and $\phi(t, v^n)$ denotes the GW phase at the PN order n.

The injected waveform would have PN order of 3.5, i.e. $O(v^7)$, in both amplitude and phase, but the template waveforms used for parameter recovery would have PN order of 0 in amplitude, and PN order of 3.5 in the phase. While it is not possible for us to directly map the PN order of the waveform's amplitude and phase to a spherical harmonic mode distribution, through these runs we can do an additional study on whether using a lower PN order in amplitude results in parameter biases. Since this is how parameter recovery is typically done for low-mass systems with LIGO, this can give us useful information about how well lower order PN approximations constrain parameters for precessing systems, and whether or not the parameter accuracy can be improved when higher order PN amplitudes are used.

IV. RESULTS

A. Visualization of the gravitational radiation field

The gravitational radiation field of a binary black hole can be decomposed into a set of gravitational-wave modes, $h_{\ell m}$, in a basis of spin-weighted spherical harmonics. We can write the full waveform of a system of binary black holes as:

$$h(t, \vec{\lambda}; \theta, \phi) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h_{\ell m}(t, \vec{\lambda})\ {}^{-2}Y_{\ell m}(\theta, \phi), \tag{23}$$

where (θ, ϕ) denote the angles on the unit sphere. In particular, θ is the inclination angle between the total angular momentum vector $\vec{J}_0$ of the binary and the observer, and ϕ is the corresponding azimuthal angle in the binary source frame; $h_{\ell m}(t, \vec{\lambda})$ is the amount of radiation emitted in a given mode (defined by the values of ℓ and m) as a function of time and the binary parameters $\vec{\lambda}$, and ${}^{-2}Y_{\ell m}(\theta, \phi)$ are the spin-weighted spherical harmonic modes.

Thus, using the $h_{\ell m}(t, \vec{\lambda})$ equations given in Appendix A of [14], we can use Eq.23 to construct the projection of the gravitational radiation field of a given system onto the unit sphere at various points in time. In doing so, we generated a set of animations showing the time evolution of the radiation emitted by a precessing black hole binary during its inspiral phase.

The first animation included only the quadrupole modes, $(l = 2, m = \pm 2)$, the second included all $(\ell = 2)$-modes, and the last included all $(\ell = 2)$- and all $(\ell = 3)$-modes. The system we used had the following parameters $\vec{\lambda}$: a mass ratio of $q = m_1/m_2 = 4$, an initial separation of $a_0 = 50M$, and dimensionless spin vectors of $\vec{\chi}_1 = (0.9, 0, 0)$ and $\vec{\chi}_2 = (0, 0.9, 0)$.

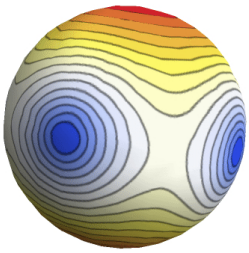

(a) $\ell = 2$ modes at $t = 1000$

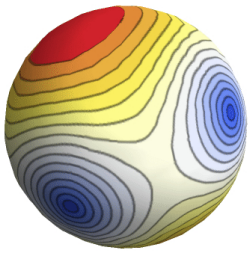

(b) $\ell = 2$ modes at $t = 7000$

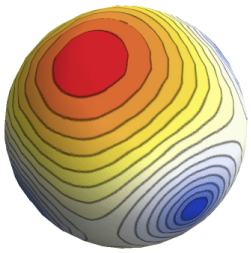

(c) $\ell = 2$ modes at $t = 12000$

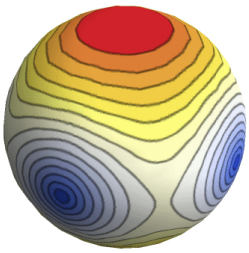

(d) quadrupole modes at $t = 1000$

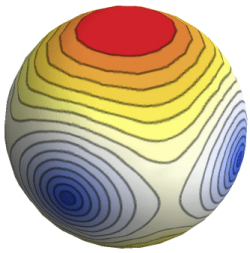

(e) quadrupole modes at $t = 7000$

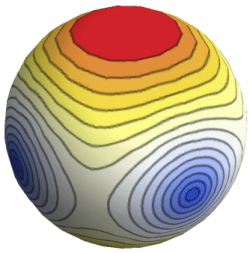

(f) quadrupole modes at $t = 12000$

FIG. 2: The top panel ((a)-(c)) shows the radiation field projected onto the unit sphere for three different times during the time evolution of the entire $\ell = 2$ waveform emitted by a black hole binary system, while the bottom panel ((d)-(f)) displays still frames from an animation of the quadrupole-only waveform of the same system for the same times.

The purpose of this was to visualize how the radiation field of a black hole binary system changes with the inclusion of higher modes. This enables our understanding of just how much information is carried by the higher modes of a system's gravitational radiation field.

Frames from two time evolutions are shown in Fig. 2. The top panel corresponds with the animation of the $(\ell = 2)$-modes, and the bottom panel corresponds with the animation of the quadrupole modes.

We know that a binary emits gravitational radiation most strongly in the direction of the orbital angular momentum of the system. In Fig. 2, the red hotspot indicates the maximum emission at any given time. Similarly, we know that a binary black hole system emits the least amount of radiation in the orbital plane. Thus, the location of the blue coldspots in each image indicates the position of the orbital plane of the system at various points in time.

With this information, we can see that the images corresponding with the full $\ell = 2$ waveform of the system show a tilt of the orbital angular momentum vector with respect to the total angular momentum vector of the system, which is directed along the z-axis. They clearly show that as time goes on the orientation of the binary changes due to precession. The frames in the bottom panel however, show very little tilt of the orbital plane and no visible precession of the binary. These distinctions were displayed even more clearly in the animated versions of the frames above. Ultimately, we found that the time evolution of the quadrupole waveform of the system looks almost like the waveform of an aligned spin system, with very little visible

precession, while the entire $\ell = 2$ waveform of the system shows clear and significant precession. Hence, the exclusion of the non-quadrupolar ($\ell = 2$)-modes leads to a significant loss of information about the binary which is directly reflected into the waveforms as illustrated in Fig. 4.

The differences between the $\ell = 2$ inclusive waveform and the $\ell = 3$ inclusive waveform (i.e. a slight increase in the magnitude of the oscillations in size of the "hotspot" on the contour map) were less visually dramatic, but nonetheless indicated that important information about the dynamics of the binary is carried by the $\ell = 3$ mode of the binary's waveform.

These animations allowed us to get a better visual understanding of the effect that including higher modes has on the waveform of a system, and helped us visualize just how much information is really carried by the higher modes of a system's waveform.

B. Match Calculation

In order to find areas in the parameter space where we can expect to find large differences between the quadrupole waveform and the higher-mode-inclusive waveform of a given system, we examined the *match* between the two waveforms. The match is defined as the normalized inner product of two waveforms, maximized over phase and time shift, and is commonly used as a measure of similarity between two waveforms:

$$\mathscr{M}(h_1,h_2) = \max_{\Delta t,\Delta\Phi} \langle \frac{h_1}{||h_1||}, \frac{h_2}{||h_2||} \rangle = \max_{\Delta t,\Delta\Phi} \frac{2}{||h_1||||h_2||} \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_1(f)\tilde{h}_2^*(f)}{S_n(|f|)} df. \tag{24}$$

In our match calculation, we used a $f_{\min}$ of 10 Hz and a $f_{\max}$ of fISCO, or the frequency of the innermost stable circular orbit of the binary, which is given, in units of total mass M, by Eq. (A2). Additionally, the noise power spectral density, $S_n(f)$, used in our match calculations was the zero-detuned high power noise curve for LIGO. Two waveforms that are exactly the same have a match of 1, while waveforms that are orthogonal have a match of 0. Thus, we can calculate the match between waveforms at various spins, mass ratios, values of (θ,ϕ), etc. to identify areas in the parameter space where the quadrupole waveform does not match well with the full waveform of the system. For example, Fig. 3 shows the match between the quadrupole waveform and the full $\ell = 2$ waveform of a black hole binary system with a total mass of $M = 12M_\odot$, an initial separation of $a_0 = 40M$, dimensionless spin vectors of $\vec{\chi}_1 = (0.9,\ 0,\ 0)$ and $\vec{\chi}_2 = (0,\ 0.9,\ 0)$, and an azimuthal angle of $\phi = 0$, at various inclination angles θ and at four different mass ratios, $q \in [1,2,3,4]$.

We see that for each mass ratio, there exist values for the inclination angle θ where the match between the quadrupole waveform and the full $\ell = 2$ waveform dips well below 1. These dips are most pronounced for higher mass ratios. For example, the match between the $q = 4$ quadrupole and $\ell = 2$ waveforms drops by more than 35% from $\mathscr{M} = 1$ at $\theta = 0$ to $\mathscr{M} = 0.626$ at $\theta = \pi/2$. This indicates that, as anticipated, for more asymmetric systems, such as systems with higher mass ratios, there are significant differences between the system's quadrupole waveform and its higher mode-inclusive waveform.

To illustrate the differences between the quadrupolar and the higher modes waveform, we show the amplitudes of the waveforms in Fig. 4 for $q = 4$, $\vec{\chi}_1 = (0.9,\ 0,\ 0)$, $\vec{\chi}_2 = (0,\ 0.9,\ 0)$, and $a_0 = 40M$ system at the viewing angle $\theta = \pi/2$ and $\phi = 0$. We see that the low match between the two waveforms is reflected in the significant amplitude and phase difference between the actual signals.

We then considered the match between the $\ell = 2$ waveform and the $\ell = 4$ waveform (which includes all $\ell = 2$, $\ell = 3$, and $\ell = 4$ modes) for the same system, and found that the match remains relatively close to 1 for most orientations, as shown in Fig. 5.

Although there is some mismatch between the waveforms, particularly at $\theta = \pi/2$, it is nowhere near as dramatic as the mismatch between the quadrupole waveform and the full $\ell = 2$ waveform. This tells us

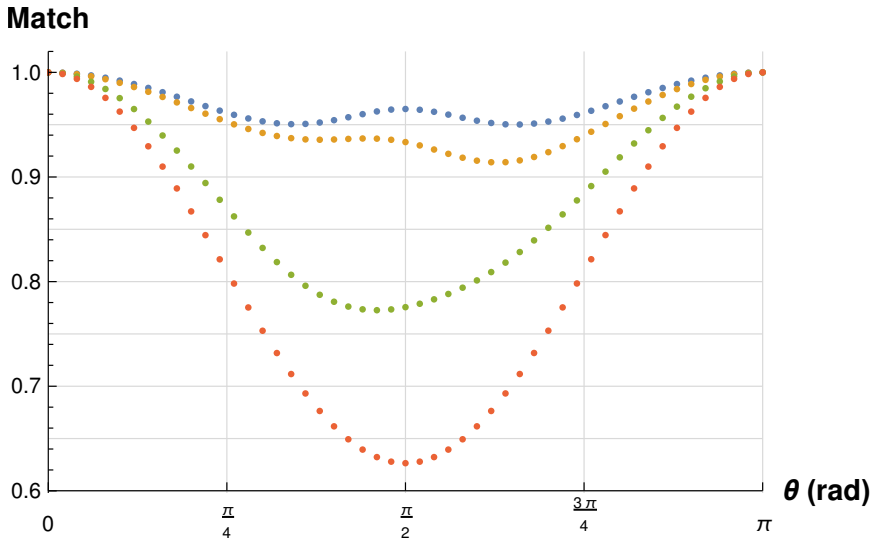

FIG. 3: The match between the quadrupole and $l = 2$ waveforms as a function of θ. The blue dots correspond to $q = 1$, the orange dots correspond to $q = 2$, the green dots correspond to $q = 3$, and the red dots correspond to $q = 4$.

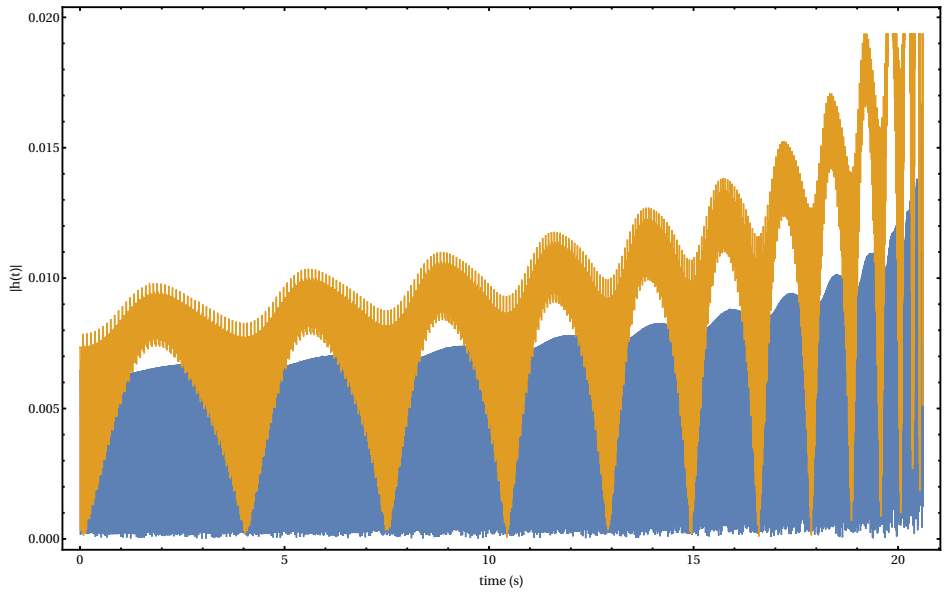

FIG. 4: The quadrupole (shown in blue) and $l = 2$ (shown in orange) waveform signals of a single black hole binary system at the inclination angle θ at which the match between the two waveforms is minimized.

is that, as expected, the $\ell = 3$ and $\ell = 4$ modes of the precessing binary black hole system are subdominant to the $\ell = 2$ modes. Again, while the matches we calculated for the different waveforms do not necessarily translate directly to a bias in parameter estimation, they do allow us to gain a better understanding of what areas in the parameter space we can expect to find a large difference in the information that can be extracted from a quadrupole or $\ell = 2$ waveform as opposed to from an $\ell = 3$ or $\ell = 4$ waveform.

By holding $a_0 = 40M$ and $M = 12M_\odot$ constant and stepping through the remaining parameters (mass ratio, spin magnitudes, spin orientations, and inclination angle) in a systematic way, we were able see how the match between the $\ell = 2$ waveform and $\ell = 4$ waveform changes as a function of each parameter. For example, we saw that the match between the two waveforms is smallest for systems with parallel spins, i.e., $\vec{\chi}_i = (\chi_i, 0, 0)$. This is shown in Fig. 6, which displays the match between the $\ell = 2$ and $\ell = 4$ waveforms

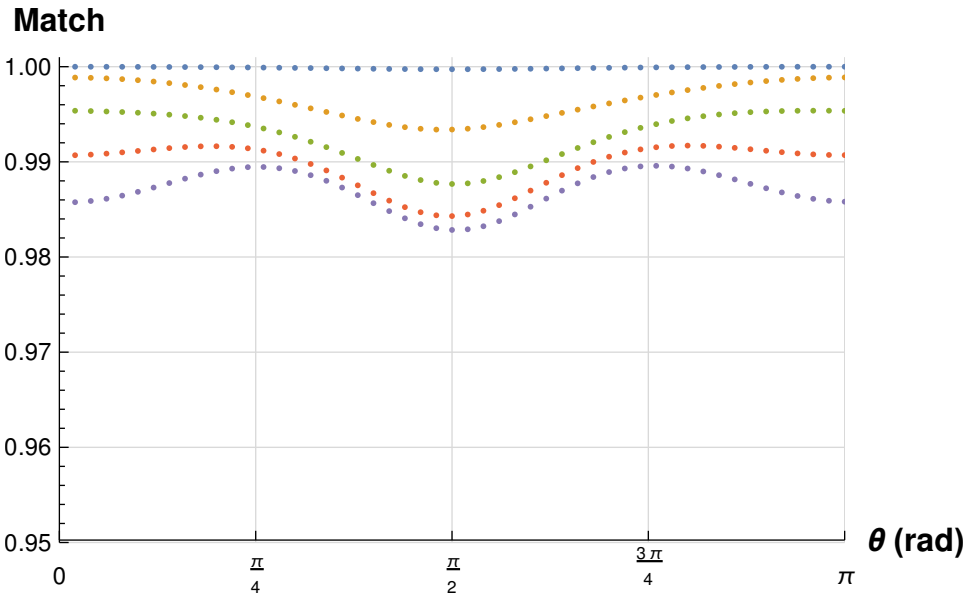

FIG. 5: The match between the $l = 2$ and $l = 4$ waveforms as a function of θ. The blue dots correspond to $q = 1$, the orange dots correspond to $q = 2$, the green dots correspond to $q = 3$, the red dots correspond to $q = 4$, and the purple dots correspond to $q = 5$.

for $M = 12M_{\odot}$, $\vec{\chi}_1 = (0.6,\ 0,\ 0)$, $\vec{\chi}_2 = (0.6,\ 0,\ 0)$, and an initial separation of $a_0 = 40M$ as a function of θ. It is apparent that the minimum matches reached by this system for each mass ratio q are smaller than the minimum matches reached by the systems shown in Fig. 5.

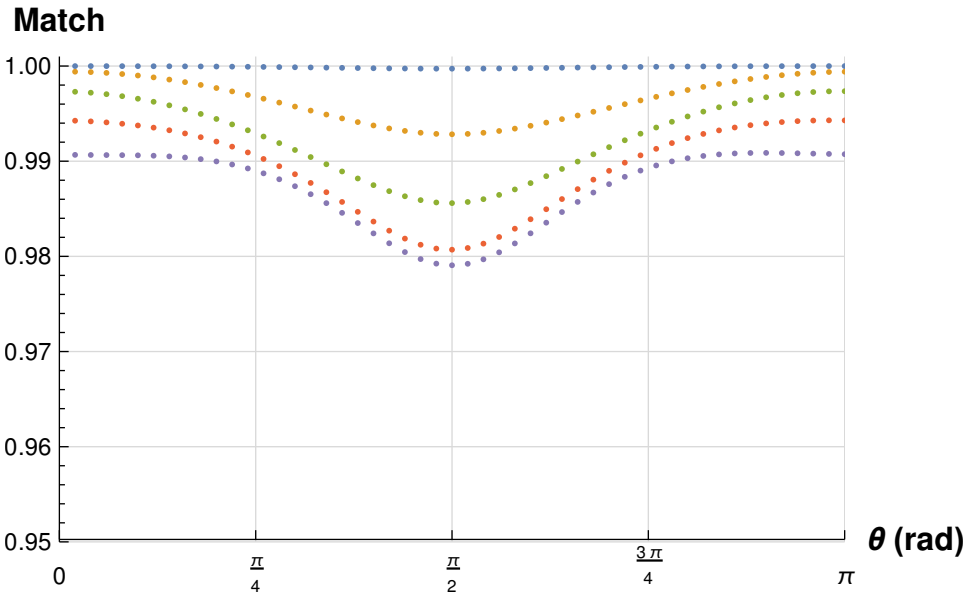

FIG. 6: The match between the $l = 2$ and $l = 4$ waveforms as a function of θ for a system with parallel spins. The blue dots correspond to $q = 1$, the orange dots correspond to $q = 2$, the green dots correspond to $q = 3$, the red dots correspond to $q = 4$, and the purple dots correspond to $q = 5$.

Using our calculated matches, we were able to determine that the systems with the highest mismatch between the $\ell = 2$ and the $\ell = 4$ waveforms had approximately the following parameters: $q = 6$, $\vec{\chi}_1 = \vec{\chi}_2 = (0.6,\ 0,\ 0)$, and $\theta = \pi/2$.

C. Statistical Error Analysis

1. *Aligned-spin case*

We first applied the Fisher analysis to an "aligned spin case", i.e., a binary configuration for which the spins of each black hole are aligned with the orbital angular momentum of the black hole binary. In this case, no precession of the orbital plane is present, so the inspiral waveform of the system depends only on the following parameters: the mass ratio q, the initial separation a_{init}, the spin magnitude of the less massive black hole $||\vec{\chi}_1||$, and the spin magnitude of the more massive black hole $||\vec{\chi}_2||$. The total mass of the system M_{total} and the distance to the system D_L are scaling factors that are added in later to rescale the waveform and insert physical units. This is true for the precessing case as well.

Using the method described in Section III A, we then generated Fisher information matrices and covariance matrices for the ($\ell = 2$), ($\ell = 2,3$), and ($\ell = 2,3,4$)-waveforms of a specific aligned spin system, where we only varied the mass ratio q, the dimensionless spin magnitudes $||\vec{\chi}_1||$ and $||\vec{\chi}_2||$, and inclination angle θ to calculating the Fisher information matrix. This resulted in a 4x4 covariance matrix of the following form:

$$C \equiv [\Gamma^{-1}] = \begin{pmatrix} \sigma_q^2 & \sigma_{q\chi_1} & \sigma_{q\chi_2} & \sigma_{q\theta} \\ \sigma_{\chi_1 q} & \sigma_{\chi_1}^2 & \sigma_{\chi_1\chi_2} & \sigma_{\chi_1\theta} \\ \sigma_{\chi_2 q} & \sigma_{\chi_2\chi_1} & \sigma_{\chi_2}^2 & \sigma_{\chi_2\theta} \\ \sigma_{\theta q} & \sigma_{\theta\chi_1} & \sigma_{\theta\chi_2} & \sigma_{\theta}^2 \end{pmatrix} \tag{25}$$

For certain parameters, we found the Fisher matrix to be ill-conditioned, i.e., resulting in negative values in the diagonals of the covariance matrix, which yields imaginary parameter uncertainties, and covariance matrices that are effectively zero. While this is expected for systems with a higher-dimensional parameter dependence, such as precessing systems, this behavior is not anticipated in the aligned-spin case. We were able to identify the error source (inconsistent units), we were then able to produce physically reasonable covariance matrices.

The binary configuration we analyze here has the following parameters: mass ratio $q = 4$, spin magnitudes $||\vec{\chi}_1|| = 0.1$ and $||\vec{\chi}_2|| = 0.6$, inclination angle $\theta = \pi/3$, azimuthal angle $\phi = 0$, total mass $M_{total} = 12\mathrm{M}_\odot$, and initial separation $a_{init} = 40\mathrm{M}$. We find the following covariance matrix for the ($\ell = 2$)-waveforms, where in the following all (ℓ, m)-modes are computed at v^4 post-Newtonian order:

$$C = [\Gamma^{-1}] = \begin{pmatrix} 0.000102585 & -2.15201\times10^{-6} & -2.29191\times10^{-6} & 0.0000135053 \\ -2.15201\times10^{-6} & 2.20217\times10^{-6} & -4.05466\times10^{-7} & -0.0000653031 \\ -2.29191\times10^{-6} & -4.05466\times10^{-7} & 1.47451\times10^{-6} & -0.0000154815 \\ 0.0000135053 & -0.0000653031 & -0.0000154815 & 0.00401989 \end{pmatrix} \tag{26}$$

In addition, the analysis was repeated twice: 1) using all ($\ell = 2,3$)-modes as v^4 PN-order, and 2) using all ($\ell = 2,3,4$)-modes at v^4 PN-order. The resulting 2σ, or 95.4% confidence level, error ellipses for each pair of parameters are shown in Fig 7. For all the parameter planes shown in Fig. 7, the orange curve, which depicts the results when the ($\ell = 3$)-modes are included, is completely obscured by the green curve, which corresponds to the results with the inclusion of ($\ell = 4$)-modes as well. This indicates that including the ($\ell = 4$)-modes has little effect on the statistical error in this particular aligned spin case.

From the covariance matrix itself, it is clear to see that the statistical uncertainty in the inclination angle θ is much larger than the statistical uncertainty in any of the other parameters. This is clearly reflected in the $q-\theta$, $\chi_1-\theta$, and $\chi_2-\theta$ error ellipses, which show a much larger spread in θ than in the other parameters. This is especially true when higher modes are included. In each case, we see that the inclusion

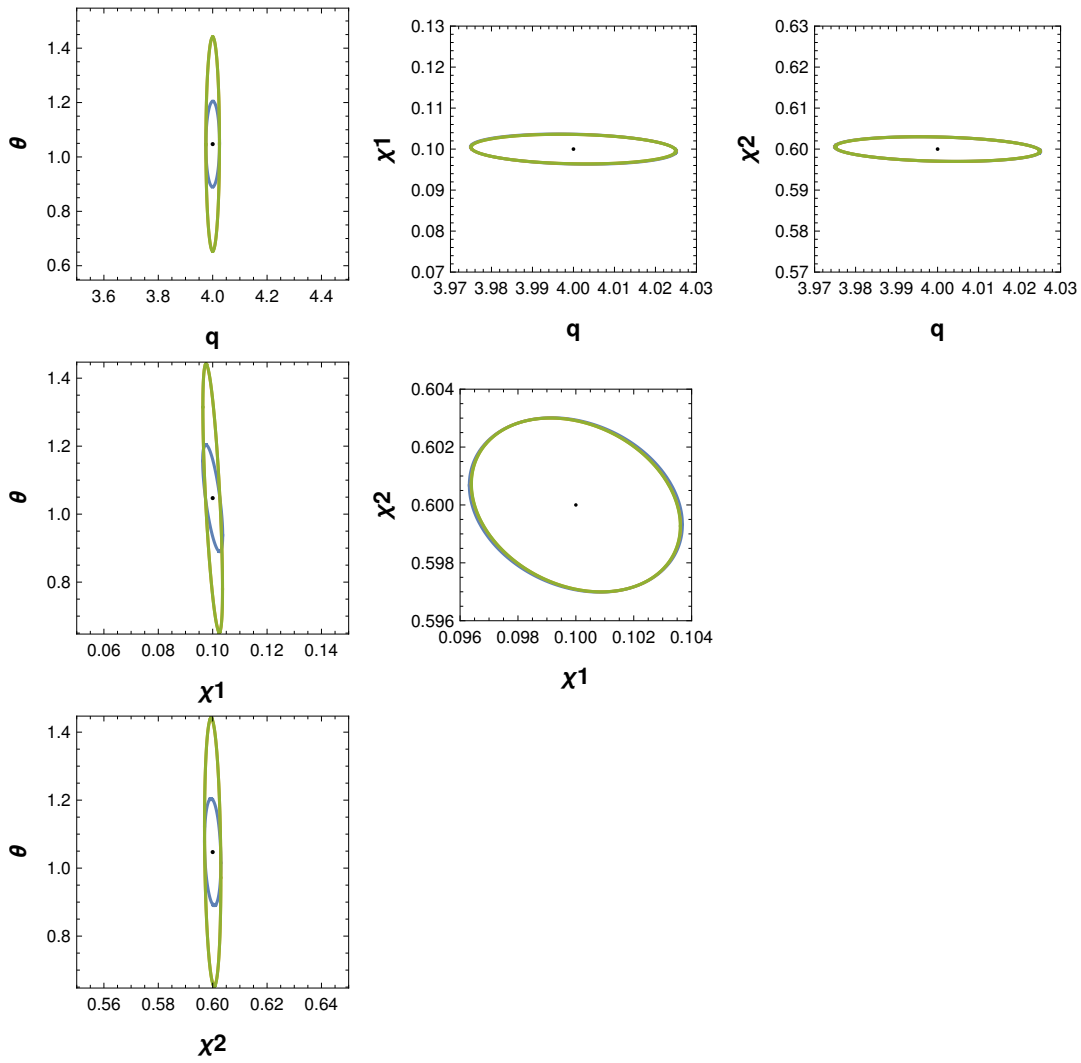

FIG. 7: 2σ confidence ellipses for each pair of parameters in the aligned spin case. The blue ellipse depict the results for the $(\ell = 2)$ waveform, the orange ellipse the $(\ell = 2,3)$ waveform, and the green ellipse the $(\ell = 2,3,4)$ waveform.

of the $(\ell = 3)$-modes (orange curves, hidden behind the green curves) significantly increases the statistical error in θ. Furthermore, in the $\theta - \chi_1$ error ellipse, we observe a small counter-clockwise rotation of the error ellipse with the inclusion of higher modes, which in indicates a very small correlation between the spin on the smaller black hole and the orientation of the binary. No such correlation behavior is observed for the spin on the larger black hole and θ. At this stage, we are unsure whether this is of physical origin or not and suggest that an ensemble of aligned configurations be analyzed. Further, we note that the scale of the χ_1 axis is much smaller than the scale of the θ axis, which artificially widens the ellipses along the χ_1 axis, thus somewhat exaggerating the inclination of the ellipses in that plane. However, this has no effect on the inclination of the ellipses with respect to each other, and as such does not change the fact that there exists a clear difference between the inclination angle of the error ellipse for the $(\ell = 2)$-waveform and inclination angle of the error ellipses for the higher-mode-inclusive waveforms.

The parameter with the second largest statistical uncertainty is the mass ratio q, while χ_1 and χ_2 have uncertainties of comparable size, as indicated both in the relatively low eccentricity of the $\chi_1 - \chi_2$ error ellipse, and in Eq. (26). Other notable features of the error ellipses in Fig. 7 include the lack of correlation, indicated by the inclination of the error ellipse, between most parameters, with the exception of $\theta - \chi_1$,

as mentioned previously, and $\chi_1 - \chi_2$. The ellipses imply that the χ_1 and χ_2 parameters are most strongly correlated, while the χ_1 and θ parameters are somewhat correlated in the $(\ell = 2)$-waveform and the remaining parameters are negligibly correlated with one another, both with and without the inclusion of higher multipoles. The relatively strong correlation between χ_1 and χ_2 is rather unsurprising, as in the aligned spin case, a small, positive offset in the magnitude of χ_1 can relatively easily be compensated for with a small, negative offset in the magnitude of χ_2 to produce a similar waveform, and vice versa. This is reflected by the fact that in the aligned spin case, the two spins can be combined into one "effective" dimensionless spin parameter $\chi_{\rm eff}$ [20] given by

$$\chi_{\rm eff} = \frac{m_1\chi_1 + m_2\chi_2}{m_1 + m_2}. \tag{27}$$

The waveform of a system is unchanged if χ_1 and χ_2 change, so long as $\chi_{\rm eff}$ remains constant. There is however a very slight clockwise rotation of the error ellipse in the $\chi_1 - \chi_2$ plane with the inclusion of higher modes, which may indicate that the inclusion of higher order modes has an effect on this correlation, but the rotation is so small that we cannot as of yet make any conclusive statements about the effect of higher order modes on the $\chi_1 - \chi_2$ degeneracy in the aligned spin case.

2. *Precessing case*

Next, we analyzed a precessing case. At this point we decided to re-parameterize our waveforms to be a function of m_1, the mass of the smaller black hole, m_2, the mass of the larger black hole, and η, the symmetric mass ratio of the system, defined below, as opposed to just the mass ratio q.

$$\eta = \frac{m_1 m_2}{M_{total}^2}. \tag{28}$$

We again encountered problems related to ill-conditioned Fisher matrices, in particular we found negative values in the diagonal entries of the covariance matrices, which has been noted previously for precessing systems in Refs. [15, 17]. While we were able to positive definite covariance matrices for several systems, for other cases, such as cases with mass ratio $q = 6$, we found that the covariance matrix was still unphysical, which may indicate that a Fisher matrix analysis is not well-suited for some parts of the precessing binary parameter space.

Nevertheless, we successfully generated covariance matrices for a few cases, including the following system, which our analysis is focused on: $m_1 = 2.4\ \mathrm{M}_\odot$, $m_2 = 9.6\ \mathrm{M}_\odot$, $\eta = 0.16$, $\chi_1 = \chi_2 = (0.6,\ 0,\ 0)$, $\theta = \pi/3$, $\phi = 0$, $M_{total} = 12\ \mathrm{M}_\odot$ and $a_{init} = 40$ M. The parameters we varied in the Fisher analysis were the following: $\vec{\lambda} = (m_1,\ m_2,\ \eta,\ \chi_{1x},\ \chi_{1y},\ \chi_{1z},\ \chi_{2x},\ \chi_{2y},\ \chi_{2z},\ \theta)$, resulting in an 8x8 covariance matrix. We once again generated three covariance matrices for this system, one using only $(\ell = 2)$-modes to generate $h(t, \vec{\lambda})$, one using all $(\ell = 2, 3)$-modes, and one using all $(\ell = 2, 3, 4)$-modes.

First, we consider the error ellipses in the $m_1 - m_2$ plane and in the $\chi_{1z} - \chi_{2z}$ plane, shown in Fig. 8. For these ellipses, the curve which includes the $(\ell = 2, 3)$-modes is again obscured by the curve showing the results for the $(\ell = 2, 3, 4)$-waveform, indicating that the inclusion of $(\ell = 4)$-modes has little effect on statistical uncertainty and correlation for the parameter pairs shown in Fig. 8.

In the $m_1 - m_2$ plane, we see that the error ellipse is roughly circular, indicating that the statistical uncertainties of m_1 and m_2 are of comparable size. Also, the error ellipses for the $(\ell = 2, 3)$ and $(\ell = 2, 3, 4)$-waveforms appear to be very slightly inside the error ellipse for the $(\ell = 2)$-waveform, potentially indicating that the inclusion of higher modes slightly improves the statistical uncertainty in m_1 and m_2.

In the $\chi_{1z} - \chi_{2z}$ plane, we see that the error ellipse is visibly inclined, indicating a strong correlation between the two parameters. This is analogous to the correlation seen between χ_1 and χ_2 in the aligned spin

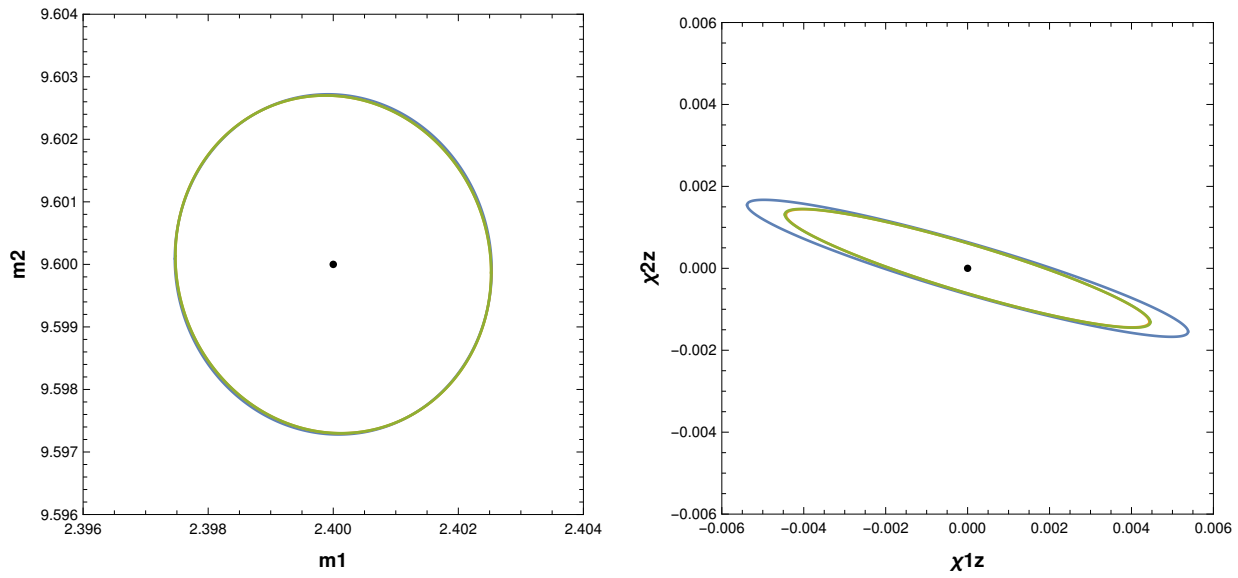

FIG. 8: The 2σ confidence ellipses in the $m_1 - m_2$ plane and the $\chi_{1z} - \chi_{2z}$ plane. Once again the blue ellipses show the results for $(\ell = 2)$ waveform, the orange ellipses for the $(\ell = 2,3)$ waveform, and the green ellipses for the $(\ell = 2,3,4)$-waveform.

case, and indicates that there exists some $\chi_{\rm eff}$ value that can be used to parametrize the waveform in place of χ_{1z} and χ_{2z} in the precessing case as well. There is also a clear decrease in the statistical uncertainty in both parameters with the inclusion of higher modes, as indicated by the decreased size of the error ellipses corresponding with higher mode inclusion.

Next we consider the error ellipses between various pairs of spin components: $\chi_{1x} - \chi_{1z}$, $\chi_{1y} - \chi_{1z}$, $\chi_{2x} - \chi_{2z}$, and $\chi_{2y} - \chi_{2z}$. In each of the error ellipses shown in Fig. 9, we see that there is little to no correlation between the spin components in the plane of the orbit (i.e. the x- and y-components) and the spin components out of the orbital plane (i.e. the z-components) for the more massive black hole. However, there does exist a slight correlation between χ_{1x} and χ_{1z} and χ_{1y} and χ_{1z}. This correlation does however decrease with the inclusion of higher order modes, as indicated by the decreased inclination of the $(\ell = 2,3)$ and $(\ell = 2,3,4)$ error ellipses with respect to the $(\ell = 2)$ error ellipse. The observed independence between the in-plane and out-of-plane components of a single black hole's spin vector has previously been noted in Refs. [13, 21]. We also see in the $\chi_{1y} - \chi_{1z}$ plane that the inclusion of higher order modes increases the spread in χ_{1y}, but decreases the statistical uncertainty in χ_{1z}.

Next, we consider the error ellipses of the in-plane spin components (i.e. in the $\chi_{1x} - \chi_{1y}$ plane and the $\chi_{2x} - \chi_{2y}$ plane). The results are shown in Fig. 10. For the in-plane components of the less massive black hole (i.e. χ_{1x} and χ_{1y}), we see that including more higher modes results in an increased correlation between the two parameters. For the more massive black hole, however, there is little visible change in the inclination of the error ellipse with the inclusion of higher modes, but instead we see an improvement in the statistical uncertainty in both χ_{2x} and χ_{2y}.

The remaining error ellipses for this case are shown in Appendix B. In general, for the inclination angle θ we see that the inclusion of higher modes results in a slight improvement in statistical uncertainty as can be seen in the panels of Fig. 18. Furthermore, while θ is only negligibly correlated with the component masses, and not correlated with either z-components of the spins, and the y-component of the spin of the less massive black hole. However, θ is strongly correlated with all other spin component, and is slightly correlated with the symmetric mass ratio η.

We see negligible correlations between the component masses and the z-components of the spins. Fur-

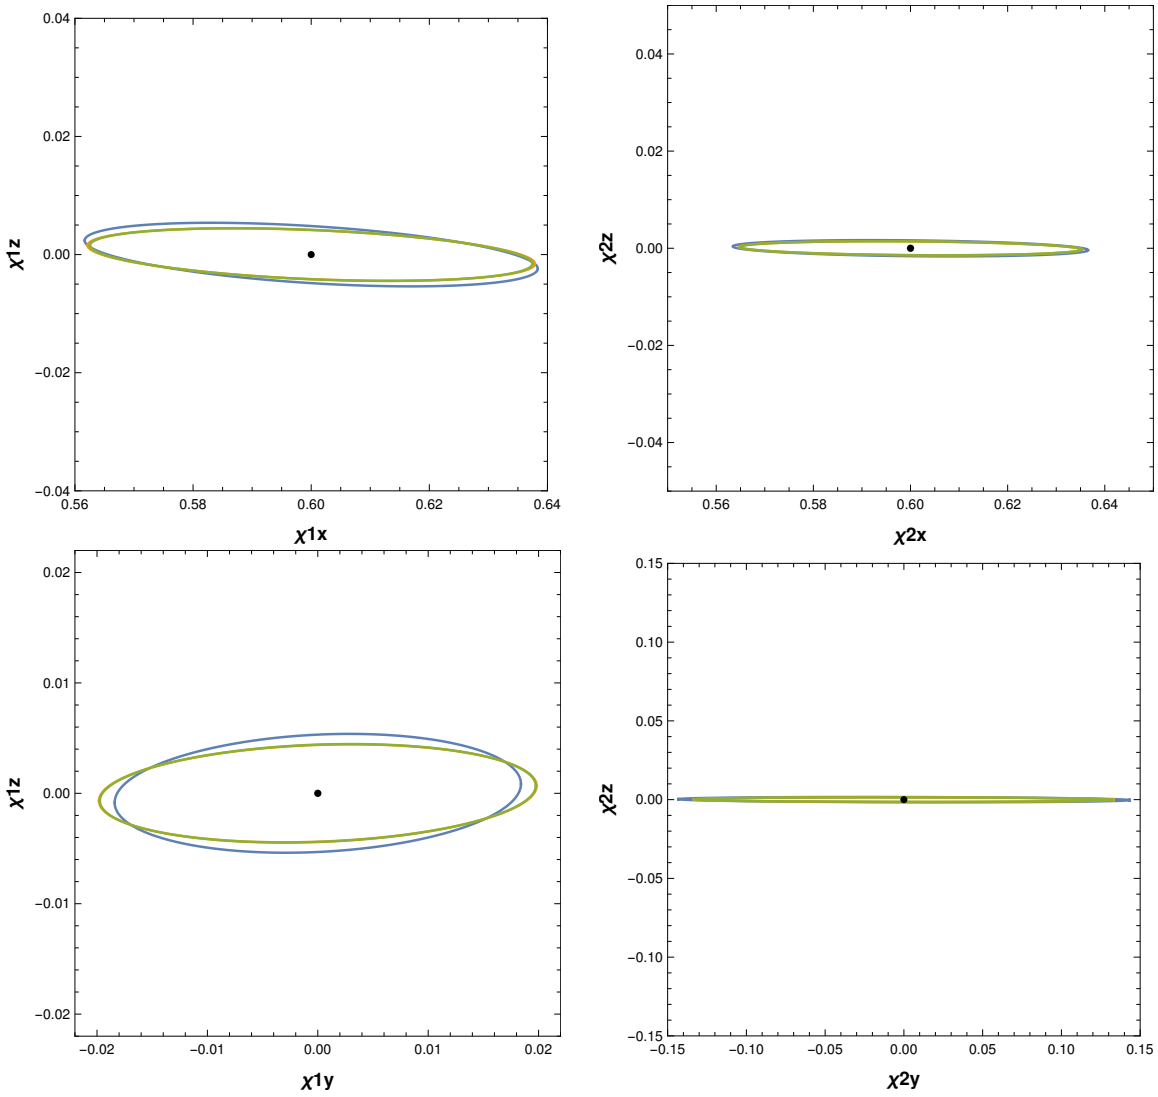

FIG. 9: The 2σ confidence ellipses in the $\chi_{1x}-\chi_{1z}$, $\chi_{1y}-\chi_{1z}$, $\chi_{2x}-\chi_{2z}$, and $\chi_{2y}-\chi_{2z}$ planes. As before, the blue ellipses shows the results for the $(\ell=2)$ waveform, the orange ellipses for the $(\ell=2,3)$ waveform, and the green ellipses for the $(\ell=2,3,4)$-waveform.

thermore, while there exists very little correlation between the component masses and the in-plane spin components, there do exist visible correlations between the symmetric mass ratio η and the spin components, particularly the x-components of both spins as can be seen in the panels in Fig. 20.

So far our analysis has only studied the effect of higher order modes on the *statistical* uncertainty of parameters for one binary case. In the future, it is necessary to both extend this analysis to a wider range of binary cases, and investigate the effect of higher order modes on *systematic* error. It is entirely possible that the uncertainties of the parameters we are studying are in fact dominated by systematic error, in which case the effect of higher order modes on systematic error is more relevant to parameter estimation accuracy than the effect of higher order modes on the statistical error.

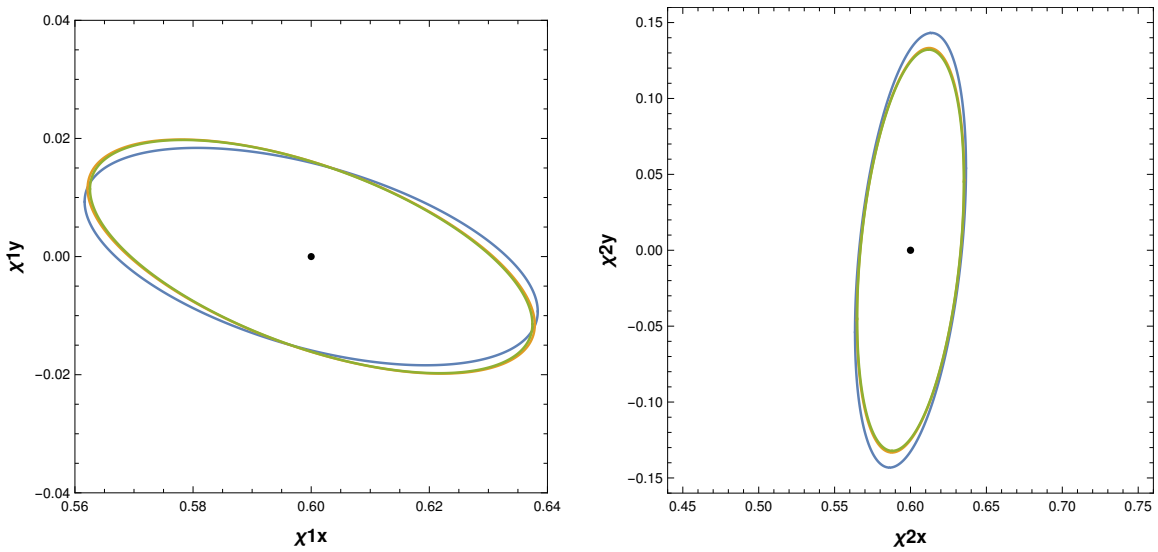

FIG. 10: The 2σ confidence ellipses in the $\chi_{1x} - \chi_{1y}$ and $\chi_{2x} - \chi_{2y}$ planes. The blue ellipses corresponds with the $(\ell = 2)$ waveform, the orange ellipses correspond with the $(\ell = 2,3)$ waveform, and the green ellipses correspond with the $(\ell = 2,3,4)$-waveform.

3. *Numerical Stability of Statistical Error*

Because our implementation of a Fisher matrix analysis uses numerical methods such as finite differencing, we must analyze the numerical stability of our analysis in order to determine whether or not we can trust our results. Particularly, we have to investigate the convergence of the obtained statistical uncertainty with the chosen parameter variation $\delta\lambda$.

In order to do so, we computed the statistical uncertainty $\Delta\lambda_{i,stat}$ for each binary parameter λ_i multiple times, varying the step size in each parameter $\delta\lambda_i$ used to generate each $\partial h/\partial\lambda_i$ in our Fisher matrix generation. In the cases presented in Sections IV C 1 and IV C 2, the $\delta\vec{\lambda}$ used to generate the Fisher matrices was, for the aligned spin case, $(\delta q,\ \delta\chi_1,\ \delta\chi_2,\ \delta\theta) = (0.1,\ 0.01,\ 0.01,\ \pi/50)$, and, for the precessing case, $(\delta m_1,\ \delta m_2,\ \delta\eta,\ \delta\chi,\ \delta\theta) = (0.01,\ 0.01,\ 0.01,\ 0.001,\ \pi/50)$, where $\delta\chi$ represents the variance used to calculate the derivative of $h(t,\ \vec{\lambda})$ with respect to the x, y and z spin components of both black holes.

Figs. 11 and 12 show how the value of the statistical uncertainty changes with each $\delta\lambda_i'$. In the aligned spin case, we see that for each parameter, for sufficiently small $\delta\lambda_i$, the resulting $\Delta\lambda_{i,stat}$ values are very close to one another, indicating minimal variations in the statistical uncertainty results for each aligned-spin binary parameter with variations in the $\delta\lambda_i$ value used to numerically generate the Fisher matrix. This indicates the numerical stability of the Fisher matrix analysis for the aligned spin case, as it indicates the convergence of each $\Delta\lambda_{i,stat}$ value to a single value as $\delta\lambda_i$ approaches smaller and smaller values.

In the precessing case, however, while the statistical uncertainty of some parameters shows the same numerical stability as for the parameters in the aligned spin case, such as for m_1, m_2, η, and θ, the statistical uncertainty of the spin parameters do not show the same convergence as the $\delta\chi$ values approach zero. The statistical uncertainties for the spin parameters of the larger black hole appear to be less convergent and show outliers. While the statistical uncertainties for the various $\vec{\chi}_1$-components are relatively similar to one another for small $\delta\vec{\chi}_1$, the statistical uncertainties for the various $\vec{\chi}_2$-components vary significantly, even at small $\delta\vec{\chi}_2$ values. As a result, we have to carefully chose a small enough $\delta\vec{\chi}$ when analyzing the statistical uncertainty results involving the spin parameters, particularly for the larger black hole, as our statistical uncertainty analysis appears to be less convergent for those parameters in the processing case.

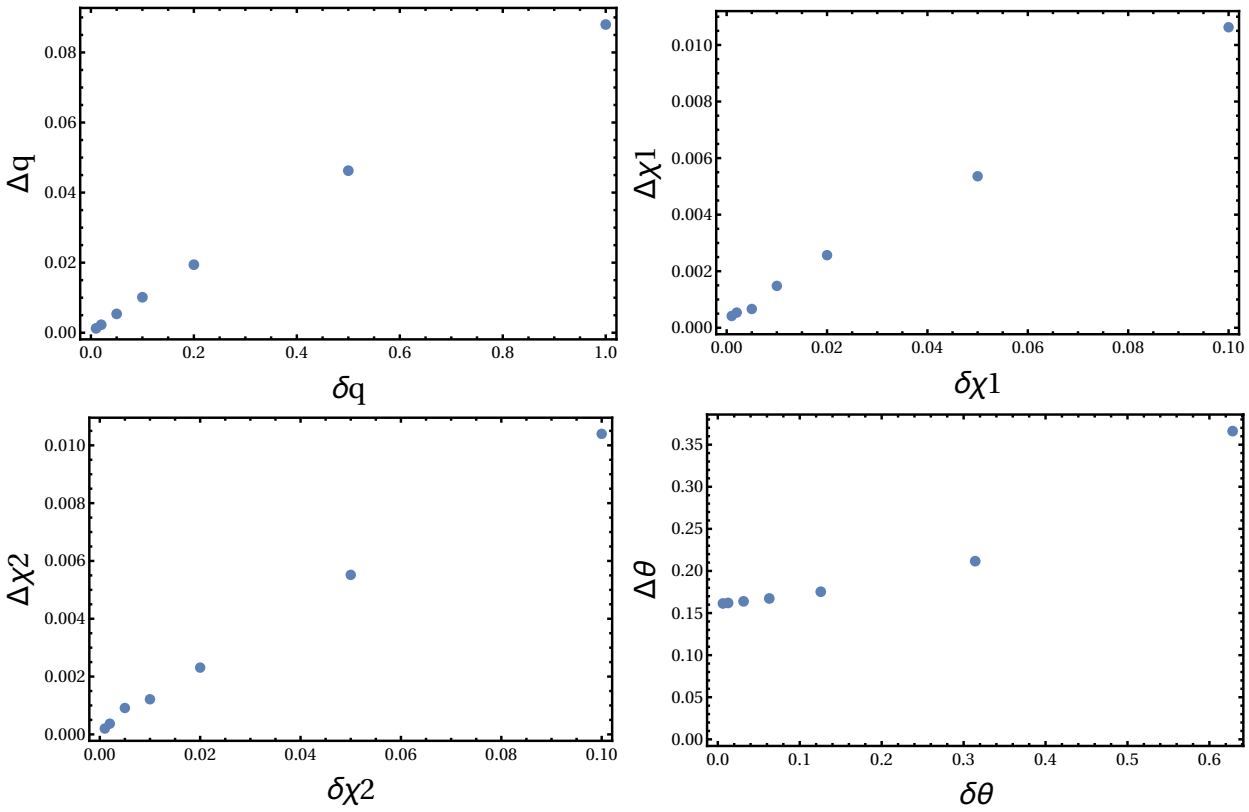

FIG. 11: The statistical uncertainty, or $\Delta\lambda_{i,stat}$, values for each parameter λ_i in the *aligned spin case*, with respect to the step size, $\delta\lambda_i$ values used to calculate the derivative of the waveform with respect to each parameter, during the generation of the Fisher matrix.

D. Systematic Error Analysis: Maximum Likelihood Estimator

The first step in implementing a *systematic* error analysis involves the numerical optimization of the likelihood function, or, equivalently, the computation of the fitting as given by Eq. (18), to determine the maximum likelihood estimator $\vec{\lambda}'$ for each binary configuration. We have yet to obtain satisfactory results from our implementation of a likelihood optimizer, but have so far obtained preliminary results which we report below.

In order to numerically maximize the fitting factor over the parameter space, we have been working with the function `NMaximize` in `Mathematica`. Because our parameter space has four dimensions in the aligned spin case and ten dimensions in the precessing case, and the likelihood function is likely to be multimodal, we must ensure that the numerical optimization algorithm we use is both able to deal with higher dimensional parameter spaces, and will not return a local maximum rather than a global maximum. However, we also want to ensure that the numerical optimization algorithms run on a reasonable time scale. Thus, we must choose a numerical optimization method that samples the parameter space densely enough to avoid converging to a local maximum, but also converges relatively quickly. Some optimization algorithms that are available in `Mathematica`'s `NMaximize` function include `SimulatedAnnealing`, `DifferentialEvolution`, and `NelderMead`.

So far, we have experimented primarily with the `NelderMead` method in `Mathematica`'s numerical optimization function. Although the Nelder-Mead method is sensitive to local maxima, its simplex method, described in Ref. [22], is well equipped to dealing with multi-dimensional parameter spaces, and by running the algorithm multiple times with different random seeds to initialize the parameter space sam-

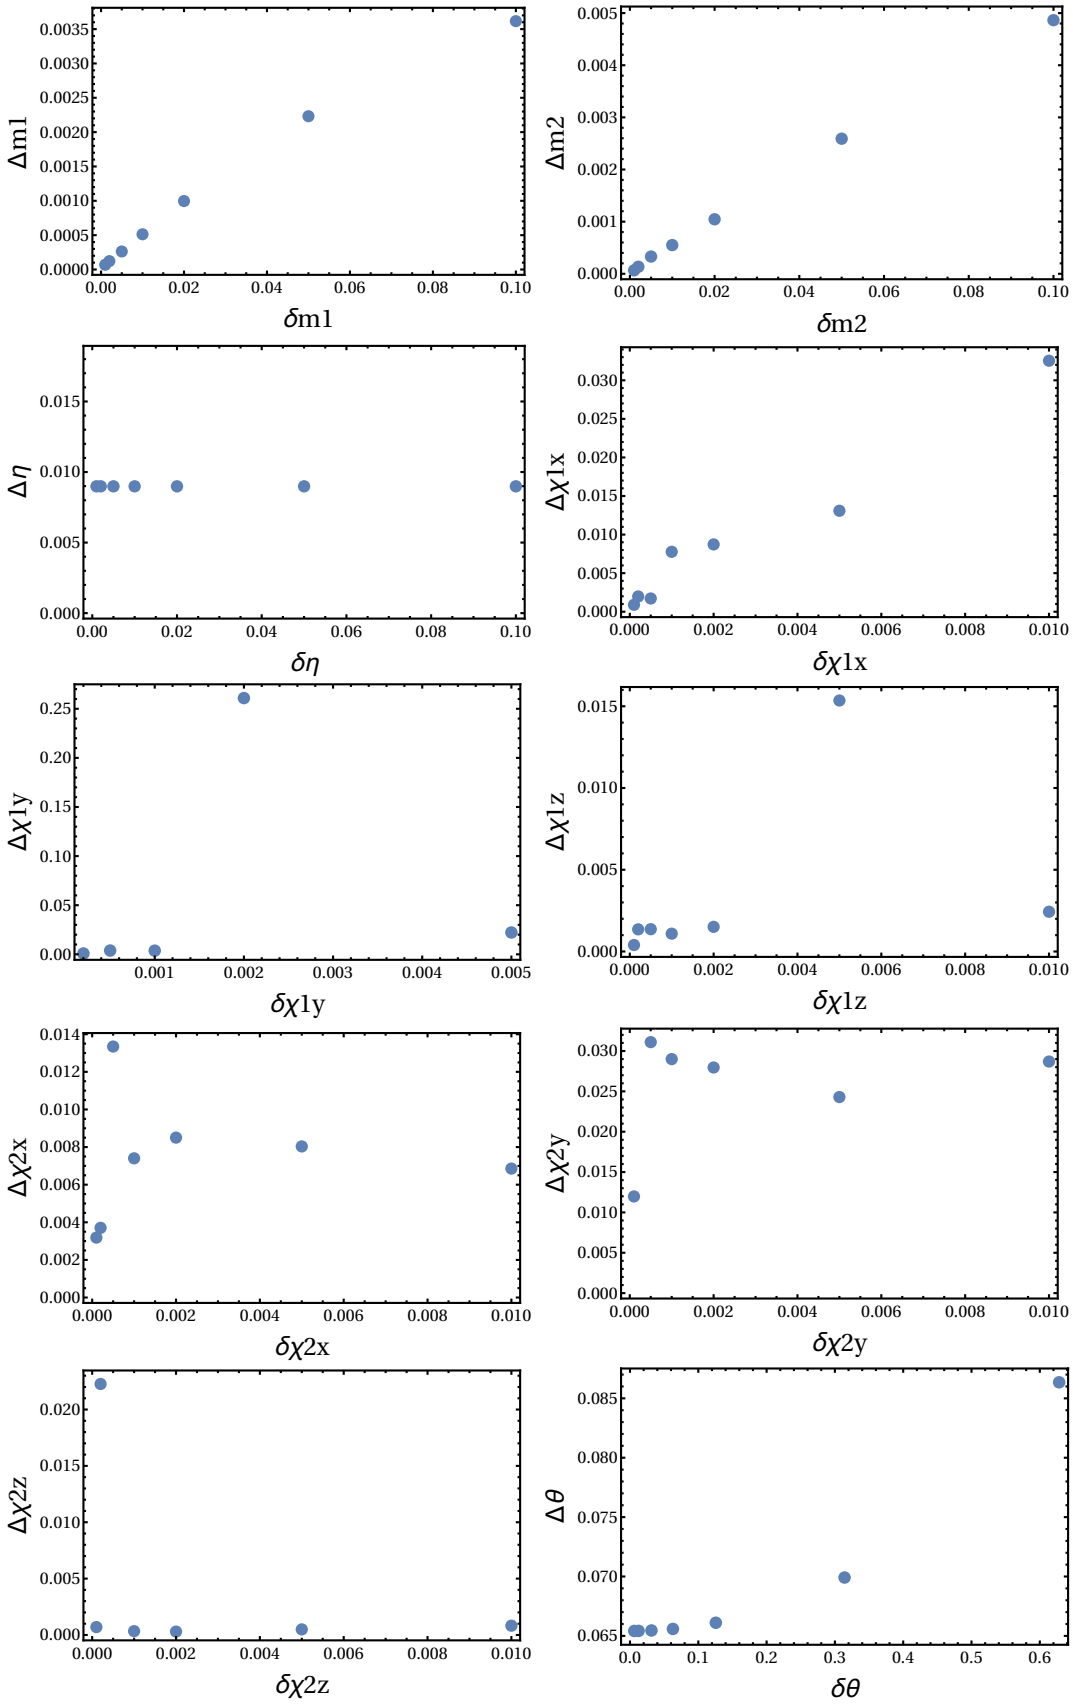

FIG. 12: The statistical uncertainty, or $\Delta\lambda_{i,stat}$, values for each parameter λ_i in the *precessing case*, with respect to the step size, $\delta\lambda_i$ values used to calculate the derivative of the waveform with respect to each parameter, during the generation of the Fisher matrix.

q'	χ_1'	χ_2'	θ'	FF
1.02649	0.130855	0.0649579	0.301209	0.999962
1.00909	0.0500649	0.148329	1.53802	0.999994
1.02851	0.0974825	0.0959615	0.479387	0.999868

TABLE I: The best fit values $\vec{\lambda}'$ for the equal-aligned spin case $\vec{\lambda} = (q,\ \chi_1,\ \chi_2,\ \theta) = (1,\ 0.1,\ 0.1,\ \pi/3)$ obtained via numerically maximizing the likelihood.

q'	χ_1'	χ_2'	θ'	FF
1.03693	0.191861	0.510272	0.239643	0.999955
1.01516	0.150654	0.554305	1.5434	0.9999995
1.00094	0.078191	0.621529	0.0239447	0.999593

TABLE II: The best fit values $\vec{\lambda}'$ for the unequal-aligned spin case $\vec{\lambda} = (q,\ \chi_1,\ \chi_2,\ \theta) = (1,\ 0.1,\ 0.6,\ \pi/3)$ obtained via numerically maximizing the likelihood.

pling, we were able to recover reasonable values for the maximum likelihood estimator for our test runs. We can also manipulate the `ShrinkRatio`, `ContractRatio`, and `ReflectRatio` values within the `NelderMead` method option to obtain better results for the maximum likelihood estimator. By increasing each of these values, we can avoid convergence to a local maximum, at the cost of increasing the runtime of the algorithm. However, increasing the `ShrinkRatio`, `ContractRatio`, and `ReflectRatio` from their default values of 0.5, 0.5, and 1, respectively, sometimes causes the algorithm to abort for yet unclear reasons. This can sometimes be avoided by manually changing the `RandomSeed`, but as of yet we cannot predict which random seed values will result in an algorithm failure.

We first tested our likelihood optimization on an aligned spin system where both the true waveform and the template waveform were $\ell = 2$ waveforms at $O(v^2)$, with parameters $\vec{\lambda} = (q,\ \chi_1,\ \chi_2,\ \theta) = (1,\ 0.1,\ 0.1,\ \pi/3)$. The system, as well as all the other systems we tested, also had $a_0 = 40M$ and $M_{total} = 12M_\odot$. Because the true and template waveforms were identical, the likelihood optimizer should return a maximum likelihood estimator of $\vec{\lambda}' = \vec{\lambda}$, for which the fitting factor is 1. After multiple runs at different random seeds, the returned maximum likelihood estimators with the highest fitting factors included the ones listed in Table I. We found that the maximum likelihood estimates for q, χ_1, and χ_2 were in the neighborhood of their expected values 1, 0.1, and 0.1, respectively, with the q' values being most consistently close to 1. The χ_1' and χ_2' values were less well constrained, as poor estimates for the dimensionless spin magnitudes can be compensated for by more accurate mass ratio values, resulting in high fitting factor values despite the incorrect spins magnitudes. The maximum likelihood estimates for θ were the most inconsistent and inaccurate, indicating issues with our likelihood optimizer's sensitivity to changes in the template waveform with the variance of θ'.

We then repeated the analysis for an aligned spin system with unequal spins, for which $\vec{\lambda} = (1,\ 0.1,\ 0.6,\ \pi/3)$. Once again, both the true and template waveforms were $\ell = 2$ waveforms at $O(v^2)$. Some of the maximum likelihood estimators with the highest fitting factor are listed in Table II. Again, the returned maximum likelihood estimator values are most reliable for q', somewhat reliable for χ_1' and χ_1', but very unreliable for θ'.

We also tested the `NelderMead` optimization for an aligned spin system with both unequal spins and unequal masses, where $\chi_1 = 0.1$, $\chi_2 = 0.6$, and $q = 4$, such that $\vec{\lambda} = (4,\ 0.1,\ 0.6,\ \pi/3)$. From our results, it would appear that systems with unequal masses in addition to unequal spins tend to be more multi-modal, as many of the returned values for the maximum likelihood estimator had fitting factor values of $\ll 1$, indicating the convergence of the algorithm to a local maximum rather than a global maximum. Using the default options of the `NelderMead` method, our optimizer func-

tion produced incorrect maximum likelihood estimators with fitting factors in the neighborhood of 0.5. Increasing `ShrinkRatio` and `ContractRatio` from 0.5 to 0.8 and `ReflectRatio` from 1 to 1.25, we obtained slightly better, but still unsatisfactory, maximum likelihood estimators, the best of which were $\vec{\lambda}' \approx (4.03736,\ 0.531547,\ 0.908874,\ 1.22131), (4.07033,\ 0.172382,\ 0.575718,\ 1.1287)$, and $(4.03414,\ 0.282369,\ 0.720981,\ 1.5708)$, with corresponding fitting factor values of 0.914212, 0.862859, and 0.98826, respectively. Again the maximum likelihood estimate for q' is closest to the expected value of $q = 4$. This time, the maximum likelihood estimates for the θ values appear to be a bit more stable than those for the χ_1 and χ_2 parameters, but this likely due to random chance rather some effect of an unequal mass system on the fitting factor's sensitivity to changes in θ' (see Section IV D 1).

We next attempted to test our likelihood optimizer on a case more relevant to our purposes, where the true and template waveforms were no longer equivalent. We once again used $\vec{\lambda} = (4,\ 0.1,\ 0.6,\ \pi/3)$, but we now used the $(\ell = 2,3,4)$ waveform of the system at $O(v^4)$ as the true waveform, and used template waveforms that included only $\ell = 2$ modes at $O(v^4)$. For this system, setting $\vec{\lambda}' = \vec{\lambda}$ results in a fitting factor of 0.982085, so while we can no longer expect the maximum likelihood estimator to be the same as the true parameter vector for the system, we can expect that the maximum likelihood estimator yields a fitting factor of $0.982085 \leq \mathrm{FF} \leq 1$. However, using the `NelderMead` method option with its default settings, the resulting maximum likelihood estimates lead to fitting factors in the range $0.5 \leq \mathrm{FF} \leq 0.75$, and increasing the `ShrinkRatio`, `ContractRatio`, and `Reflect Ratio` values had little to no effect. Thus, further fine tuning is necessary before we can directly apply our likelihood optimizer to cases relevant to our study.

Overall, we see that the fitting factor optimization is most likely to return reliable values for the maximum likelihood estimate for the mass ratio q, and is least likely to return reliable values for the maximum likelihood estimate for the inclination angle θ, indicating that the fitting factor is most sensitive to changes in q', and least sensitive to changes in θ'. This insensitivity to differences between template waveforms at different inclination angles is somewhat predicted by our results for the statistical uncertainty in different parameters for each parameter, which indicated that the statistical uncertainty in θ is far greater than for any other parameter, in both the aligned spin and the precessing case. However, the fitting factor's relatively high sensitivity to changes in q', as opposed to changes in χ_1' and χ_2', is not reflected in our statistical uncertainty results, which indicated, as shown in Fig. 7, that the statistical spread in q is greater than the corresponding spread in χ_1 and χ_2. Furthermore, while our statistical uncertainty results indicate that θ may not be well constrained, the spread in the maximum likelihood estimator results for θ' are far greater for the $\ell = 2$ waveforms than our statistical uncertainty results would suggest. A more thorough analysis of this issue is shown below in Section IV D 1. However, as can be seen from our visualization of the gravitational radiation field in Sec. IV A, higher order modes carry information about the angular energy distribution. We are therefore confident that the inclusion of higher order modes in the template waveforms will improve the constraints on θ.

1. *Stability of Fitting Factor Calculations*

As mentioned in the previous section, one problem we encountered while testing the maximum likelihood optimizer was our likelihood function's lack of sensitivity to changes in θ, resulting in maximum likelihood estimates for θ' that were both dramatically different from the θ value used to generate the true waveform, and inconsistent across runs of the optimizer. Calculating the fitting factors between waveforms of a system with $q = 1$ and $\chi_1 = \chi_2 = 0.1$ at different θ, we obtained a fitting factor of 1 every time, indicating identical waveforms (in phase evolution), regardless of the θ value.

One potential reason for this is that our calculation of the fitting factor FF, or the maximized overlap, between the true and template waveforms used a fixed polarization angle (just like our overlap function, described in Appendix A). For most calculations, we used a polarization angle of $\psi = 0$, meaning that

the fitting factor calculations took into account only the real component of each complex strain. Plotting only the real part of the strain produced by a system with $q = 1$, $\chi_1 = \chi_2 = 0.1$, and $a_0 = 40M$ at various inclination angles θ as shown in Fig. 13, we see that the waveforms are exactly in phase, regardless of the θ value, with the variations in θ resulting only in changes in the waveform amplitude.

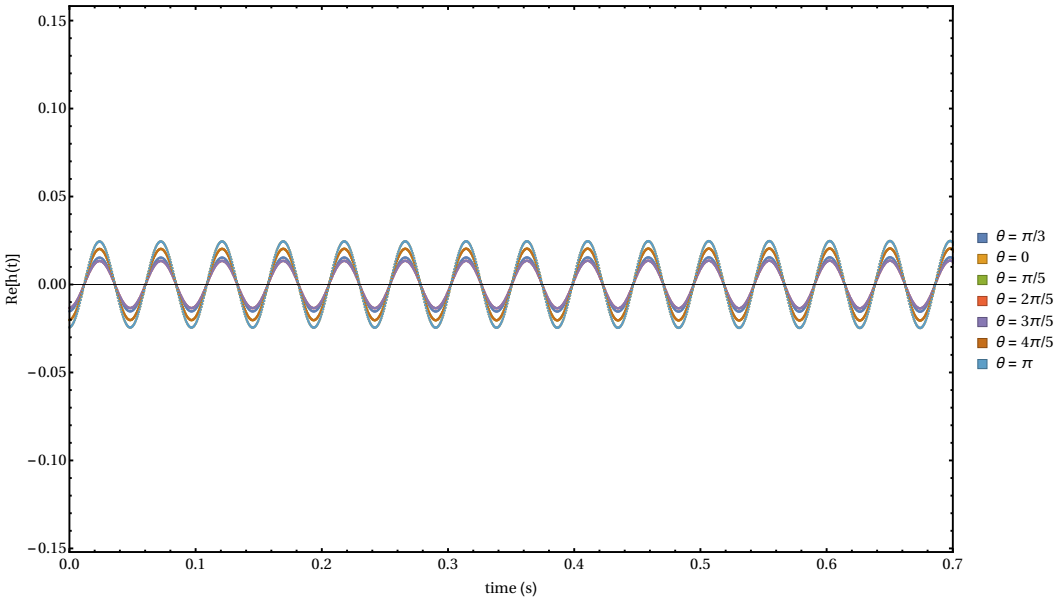

FIG. 13: The real part of the first 0.7 seconds of the strain produced by a system with $q = 1$, $\chi_1 = 0.1$, $\chi_2 = 0.6$, and $a_{init} = 40M$ at different inclination angles θ.

In order to determine if changing the polarization angle at which the fitting factor is evaluated improves the sensitivity of the fitting factor to changes in θ', we generated the first contour plot shown in Fig. 14, which shows the fitting factor between a "true" waveform, which is taken to be the $\ell = 2$ waveform at $O(v^2)$ of a system with $\vec{\lambda} = (q = 1, \chi_1 = 0.1, \chi_2 = 0.1, \theta = \pi/3)$, and a template waveform, which is the $\ell = 2$ waveform at $O(v^2)$ of a system with $\vec{\lambda}' = (q' = 1,\ \chi_1' = 0.1,\ \chi_2' = 0.1,\ \theta')$, at multiple different template inclination angles θ' and polarization angles ψ. We generated the same contour plot two more times with the true and template waveforms at $O(v^3)$ and $O(v^4)$ to see if the issue might be linked to our usage of waveforms truncated at $O(v^2)$.

From the plots in Fig. 14, it is clear that the polarization angle used to generate the fitting factor values does play a role in the fitting factor's sensitivity to changes in θ'. For example, while the fitting factor is constant at 1 regardless of the θ' value for certain polarization angles (i.e. $\psi = 0,\ \pi/2,\ \pi$ and $3\pi/2$), for other polarization angles (such as $\psi = \pi/4,\ 3\pi/4,\ 5\pi/4$ and $7\pi/4$), the fitting factor does drop below 1 for certain θ' values. However, even at those polarization angles, the fitting factor is 1 for all $0 < \theta' < \pi/2$, and only drops below 1 for $\theta' > \pi/2$, so even at those polarization angles the fitting factor is not sensitive enough to changes in θ' to return a reliable value for the maximum likelihood estimate for the inclination angle parameter. Furthermore, the three contour plots look largely identical, indicating that this problem exists regardless of the order at which the true and template waveforms are truncated. From the contour plots shown in Figs. 16 and 17, we see that this problem persists for systems with unequal spins and unequal masses as well.

For the same aligned spin system, we then generated the contour plots shown in Fig. 15, to examine how changing the polarization angle used in the fitting factor calculation effects the fitting factor between the true waveform of the system and the template waveform, as the q', χ_1', and χ_2' values used to generate the template waveform are varied.

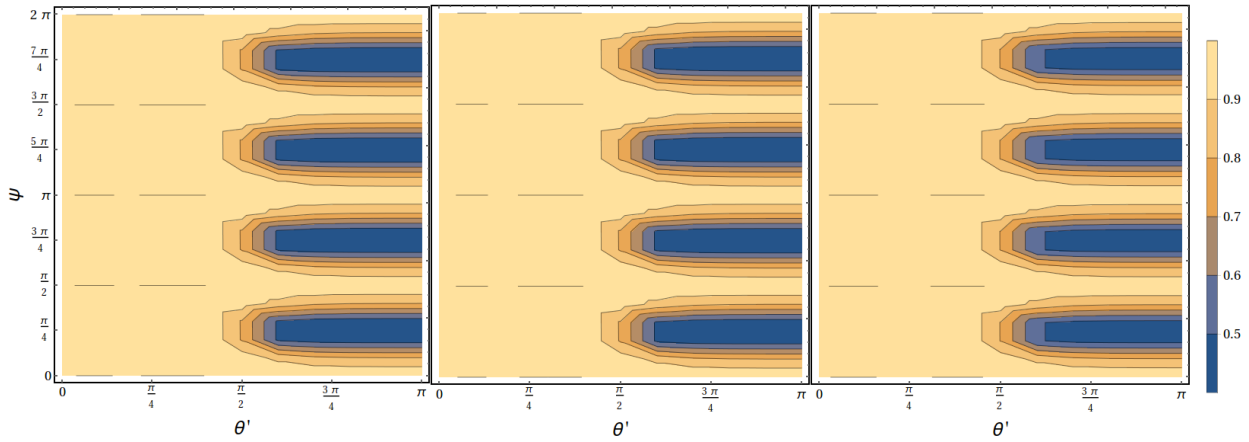

FIG. 14: Contour plots showing the behavior of the fitting factor FF with variation in both θ' and the polarization angle ψ used to calculate the fitting factor values, for a system where the true parameters were $\vec{\lambda} = (q = 1,\ \chi_1 = 0.1,\ \chi_2 = 0.1,\ \theta = \pi/3)$. From left to right, the true and template waveforms used to calculate the fitting factor values shown in the contour plots above were $\ell = 2$ waveforms at $O(v^2)$, $O(v^3)$, and $O(v^4)$, respectively. Lighter colors indicate a higher fitting factor, and darker colors indicate a lower fitting factor.

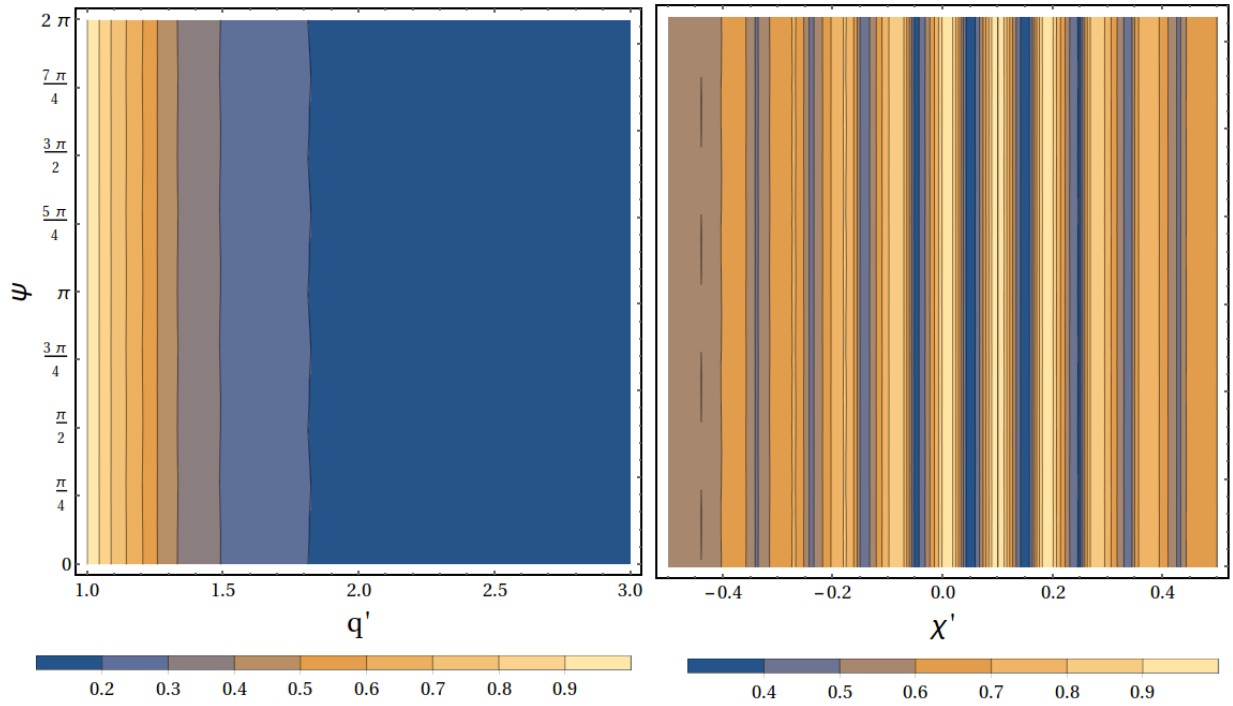

FIG. 15: Contour plots showing the behavior of the fitting factor FF with variation in both q' and χ' and the polarization angle ψ used to calculate the fit factor values, for a system where the true parameters were $\vec{\lambda} = (q = 1,\ \chi_1 = 0.1,\ \chi_2 = 0.1,\ \theta = \pi/3)$. Since this system has both equal masses and equal spins, the contour plot corresponding to χ_1' is identical to the one corresponding to χ_2', so only one is shown here.

We then did the same for the aligned spin system with unequal spins ($q = 1, \chi_1 = 0.1, \chi_2 = 0.6$, and $\theta = \pi/3$) as well as the aligned spin system with both unequal masses and unequal spins ($q = 4, \chi_1 = 0.1, \chi_2 = 0.6$, and $\theta = \pi/3$), as shown in Figs. 16 and 17, respectively.

With the exception of the inclination angle, the fitting factor calculation appears to be stable across different polarization angles with variation in all other parameters. This is shown by the fact that each

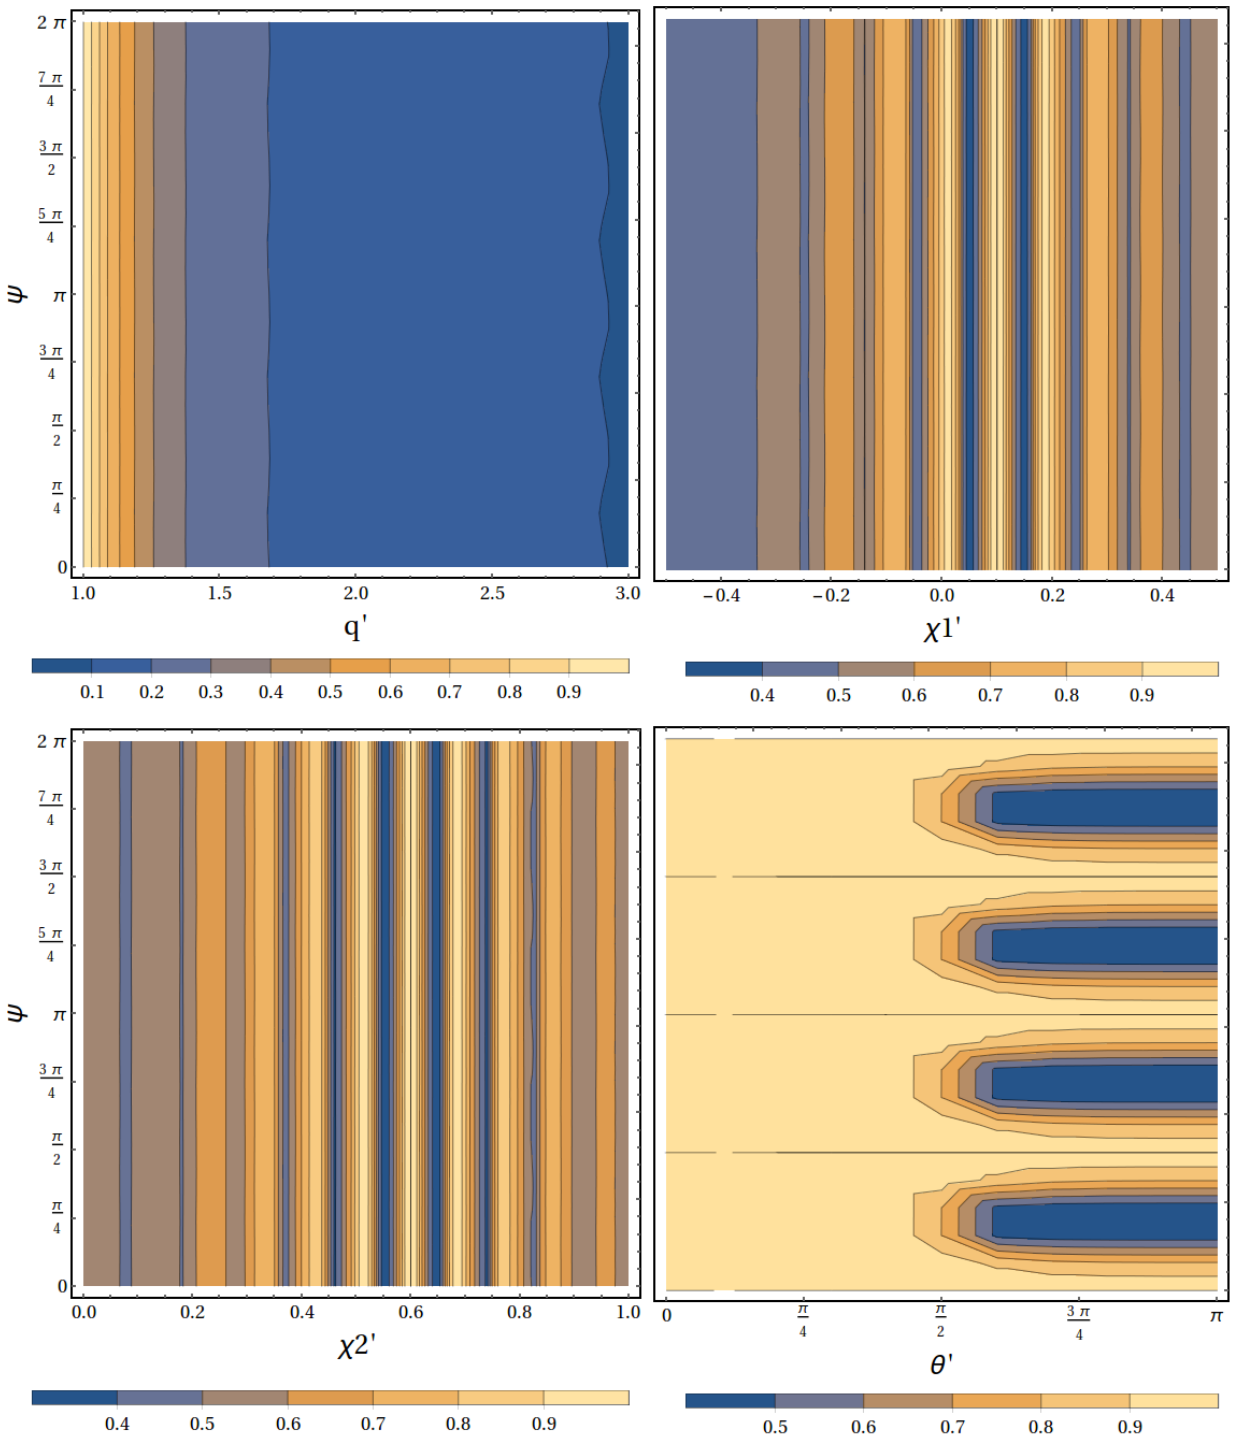

FIG. 16: Contour plots showing the behavior of the fitting factor with variations in both each individual parameter q', χ_1', χ_2', and θ', and the polarization angle ψ for a system with true parameters $\vec{\lambda} = (q = 1,\ \chi_1 = 0.1,\ \chi_2 = 0.6,\ \theta = \pi/3)$.

of the horizontal "slices" in the contour plots corresponding to q', χ_1', and χ_2' are more or less identical, indicating that the fitting factor variations that result from changes in these parameters are independent of the polarization angle.

Furthermore, in all the contour plots corresponding to the mass ratio parameter, there do not appear to

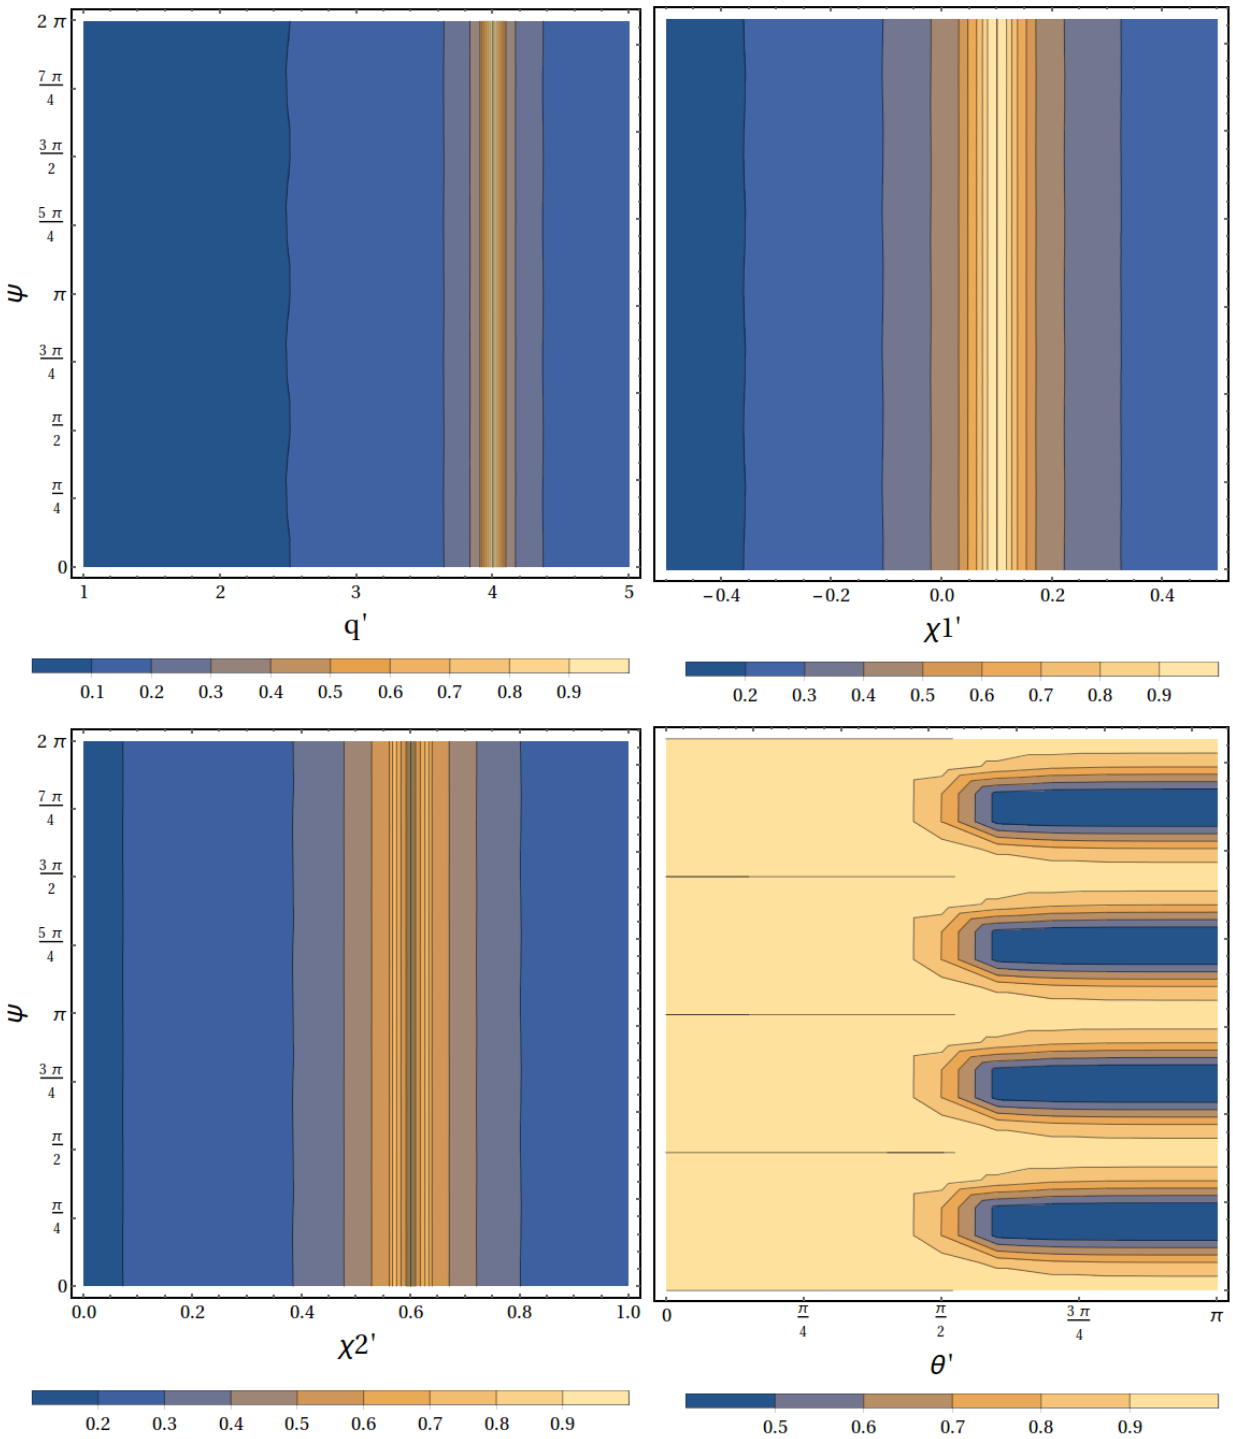

FIG. 17: Contour plots showing the behavior of the fitting factor with variations in both each individual parameter (q', χ_1', χ_2', and θ'), and the polarization angle ψ for a system with true parameters $\vec{\lambda} = (q = 4,\ \chi_1 = 0.1,\ \chi_2 = 0.6,\ \theta = \pi/3)$.

be any local maxima in the fitting factor values if q' alone is varied. Instead we see a single, strong peak in the fitting factor where $q' = q$ in both the equal and unequal mass cases. This is reflected in our preliminary maximum likelihood estimator results, which were consistently able to return relatively reliable values for q'.

By contrast, the fitting factor as a function of the individual spin magnitudes is highly multi-modal in the two equal mass cases. This can explain our likelihood optimizer's difficulty with consistently returning the correct values for the maximum likelihood estimates for χ_1 and χ_2. This multi-modality appears to decrease in the unequal mass case. In this case, the contour plot corresponding to χ_1' appears to have only a single, strong peak at the true dimensionless spin parameters of the system. Similarly, while contour plot corresponding to χ_2' has some multi-modal features in the neighborhood of $\chi_2' = \chi_2$, it has significantly less local maxima than in the equal mass cases. This would seem to contradict our preliminary results for the maximum likelihood estimator, which, as stated in Sec. IV D, implied that the likelihood function becomes more multi-modal when the true waveform corresponds with a system with unequal masses. However, these contour plots only indicate the general behavior of the fitting factor with the variation of a single variable, and do not adequately represent the multi-modal features that may arise due to the multi-dimensional nature of the parameter space.

V. DISCUSSION

This project has aimed to take a first step at assessing the effect of higher order modes on the parameter estimation of waveforms from precessing low-mass binaries. To do so, the two main contributions to parameter uncertainties need to be analyzed: Firstly, the statistical uncertainty which arises from the presence of noise, and secondly the systematic uncertainty, which occurs due to intrinsic modelling errors in the waveform families employed in GW data analysis. We have presented methods and results regarding the statistical uncertainty, but were only able to collect preliminary results regarding the systematic uncertainty. Therefore, our results regarding the effect of higher order modes on parameter estimation are not yet complete. Because it is entirely possible that the systematic error in each parameter that arises from using a higher-mode exclusive template waveform dominates over the statistical error due to the presence of noise, until we are able to complete our systematic error analysis, we have no means of assessing the significance of the statistical error results presented in Sec. IV C. Nonetheless, we were able to asses the effect of higher order modes on *statistical error* in each parameter for both an aligned spin black hole binary case, and a precessing black hole binary case.

In the aligned spin black hole binary case, which had parameters $q = 4$, $\chi_1 = 0.1$, $\chi_2 = 0.6$, and $\theta = \pi/3$, in addition to an azimuthal angle $\phi = 0$, a total mass $M_{total} = M_\odot$, and an initial separation $a_0 = 40M$, we found that while the inclusion of $\ell = 3$ modes has some effect on both the statistical spread in some parameters and the correlation between certain parameter pairs, the inclusion of $\ell = 4$ modes has little to no significant effect on statistical error. We found that the parameter with the greatest statistical error is the inclination angle, θ, and that this large statistical spread in θ is increased with the inclusion of $\ell = 3$ modes. The statistical spreads in the other parameters, q, χ_1, and χ_2 are largely unaffected by the inclusion of $\ell = 3$ and $\ell = 4$ modes, but the correlation between χ_1 and θ appears to decrease with the inclusion of $\ell = 3$ modes. Further, the parameters χ_1 and χ_2 are, unsurprisingly, heavily correlated, likely to the existence of a $\chi_{\rm eff}$ parameter, given by Eq. (27), that can be used to re-parameterize the waveform. Most other parameter pairs are negligibly correlated.

In the precessing black hole binary case, which had the parameters $m_1 = 2.4\ \mathrm{M}_\odot$, $m_2 = 9.6\ \mathrm{M}_\odot$, $\eta = 0.16$, $\chi_1 = \chi_2 = (0.6,\ 0,\ 0)$, $\theta = \pi/3$, $\phi = 0$, $M_{total} = 12\ \mathrm{M}_\odot$ and $a_{init} = 40\ \mathrm{M}$, we once again found that while the inclusion of $\ell = 3$ waveforms has an effect on the statistical uncertainties for some parameters and correlations between certain parameter pairs, the inclusion of $\ell = 4$ modes has no significant effect on statistical error and correlation. In general we found that the inclusion of $\ell = 3$ modes tends to decrease statistical uncertainty in various parameters, particularly for spin components, and most notably for the z components of the spin vectors for each black hole. The effect of including higher order modes on the correlations between parameters for this case is varied. In some cases, such as the $\chi_{1x} - \chi_{1y}$ parameter plane, the inclusion of higher modes increases parameter correlations, whereas in other cases, such as the

$\chi_{1x} - \chi_{1z}$ and $\chi_{1y} - \chi_{1z}$ parameter planes, the inclusion of higher modes decreases parameter correlations. In the precessing case, however, we also have to note that there exist some numerical instabilities in the statistical uncertainty analysis for the spin parameters, particularly the spin components of the more massive black hole.

Again, depending on the scale of *systematic* uncertainty, the effects of higher order modes on *statistical* uncertainty, as found thus far, may or may not be negligible. While we have not yet completed a systematic error analysis due to difficulties involving the numerical, nonlinear optimization of the likelihood function for various binary systems, our initial maximum likelihood estimator results have revealed some features regarding the likelihood function for aligned spin cases. For example, we have discovered that the our optimization of the fitting factor function is most likely to return reliable values for the maximum likelihood estimate q' for the mass ratio parameter q. This may be because, as shown in the contour plots in Section IV D 1, the fitting factor function is does not have many local maxima when q' alone is varied, making it easier for the numerical optimization algorithms to converge to a q' value that maximizes the fitting factor, and thus the likelihood function. We also found that when the optimization is run multiple times at many different random seeds, it can also return reliable values for χ_1' and χ_2' for some of our test cases. These estimates, however, are less consistent. This may be attributed to the fact that, as shown in the contour plots in Figs. 15, 16, and 17, the fitting factor function is highly multi-modal with variations in χ_1' and χ_2', making it easier for the optimization algorithms to get "trapped" near χ_1' and χ_2' values that correspond with a local maximum in the fitting factor function, resulting in a more unreliable maximum likelihood estimate. We also found that the optimization of the likelihood function does not, even with multiple random seeds and the manipulation of various settings within the `NMaximize` function in `Mathematica`, return reliable values for θ'. Not only are the returned values very different from the expected θ', but they are also very inconsistent across multiple runs of the optimization algorithm. This is explained by the contour plots shown in Section IV D 1, which indicate that the fitting factor function is very insensitive to changes in θ'. Further testing is necessary before our likelihood optimization function is ready to be applied to a systematic error analysis.

VI. NEXT STEPS

Once we have fine tuned our likelihood optimization function to consistently return satisfactory values for the maximum likelihood estimator for both the aligned spin and the precessing case, systematic error analysis can be performed on the cases for which we have already completed our statistical error analysis. From there, the statistical and systematic errors in each case can then be compared to analyze whether or not higher-mode inclusive waveforms are "indistinguishable" from higher-mode exclusive waveforms, in terms of parameter estimation results. Then, this error analysis can be extended to a larger ensemble of precessing systems, in order to more systematically identify whether or not there exist certain areas in the binary parameter space for which significant systematic error arises from using only quadrupole or $\ell = 2$ waveforms during parameter estimation. Bayesian parameter estimation, as described in Sec. III C can also be used to supplement the results of our error analyses using Fisher matrices.

Acknowledgements

Firstly, I would like to extend my gratitude to my mentor Patricia Schmidt, for her scientific guidance and patient support throughout this project, as well as my co-mentors Rory Smith and Yanbei Chen. I also thank Geraint Pratten, from the University of Sussex, for contributing his insight and expertise to this project as an external collaborator. I also would like to thank Alan Weinstein and his research group for enabling my participation in this program. This project was supported by the NSF, the Caltech SURF Office, and the LIGO Collaboration.

Appendix A: Mathematica Implementation of Overlap

The code to calculate the overlap between two waveforms h_1 and h_2 has been written in `Mathematica`. The overlap is calculated by fixing the polarization angles $\psi \equiv \sigma$ and providing time series data $h_1(t)$ and $h_2(t)$. The noise curve is provided as a frequency series for $S_n(f)$ in units of $\left[\mathrm{Hz}^{-1}\right]$. Algorithmically, the overlap is calculated in the following steps:

1. Compute the complex GW strains for a given orientation (θ, ϕ):

$$h_{i,\,\mathrm{strain}}(t;\theta,\phi) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h_{i,\ell m}(t)\, {}_{-2}Y_{\ell m},(\theta,\phi) \tag{A1}$$

   with $i = 1, 2$.

2. Convert the time values in each time series from units of M to units of seconds.

3. Calculate the dimensionful Fourier transforms, $\tilde{h}_1$ and $\tilde{h}_2$ using the `Fourier` function in Mathematica and multiplying the output by Δt, where Δt is the time interval in seconds. Note that we must set the `FourierParameters` flag equal to $(1,1)$ as we use non-unitary Fourier conventions in our analysis.

4. Set $f_{\mathrm{min}} = 10\,\mathrm{Hz}$ and $f_{\mathrm{max}} = M f_{\mathrm{ISCO}} * \texttt{scale}$, where $M f_{\mathrm{ISCO}}$ is the frequency in units of M of the innermost stable circular orbit (ISCO) and `scale` converts the frequency from units of M (the total mass of the system in solar masses) to units of Hz. We then take all values in $\tilde{h}_1$ and $\tilde{h}_2$ that correspond with frequencies $f_{\mathrm{min}} < f < f_{\mathrm{max}}$. $M f_{\mathrm{ISCO}}$ is given by:

$$M f_{\mathrm{ISCO}} = \frac{1}{\pi \sqrt{6}^3}, \tag{A2}$$

   which for the considered configurations corresponds to $\sim 368\mathrm{Hz}$.

5. Evaluate $S_n(f)$ at each frequency in the list.

6. Generate two lists of values:

$$L1 = \frac{\tilde{h}_1(f)\,\tilde{h}_2^*(f)}{S_n(f)} \qquad \text{and} \qquad L2 = \frac{\tilde{h}_1^*(-f)\,\tilde{h}_2^*(f)}{S_n(f)}. \tag{A3}$$

7. We can efficiently calculate the two integrals in Eq. (5) at various time shifts by taking the IFFT, given by the `InverseFourier` function in Mathematica, of $L1$ and $L2$, remember to set `FourierParameters` equal to $(-1,1)$.

8. As both waveforms have the same polarization angle ψ, we can now calculate the final overlap maximized over the time shifts by evaluating

$$Max@Re[\,\mathrm{InverseFourier[L1]} + \mathrm{InverseFourier[L2]}\, e^{4i\psi}\,]. \tag{A4}$$

Appendix B: Additional Error Ellipses

1. Error Ellipses for Inclination Angle θ

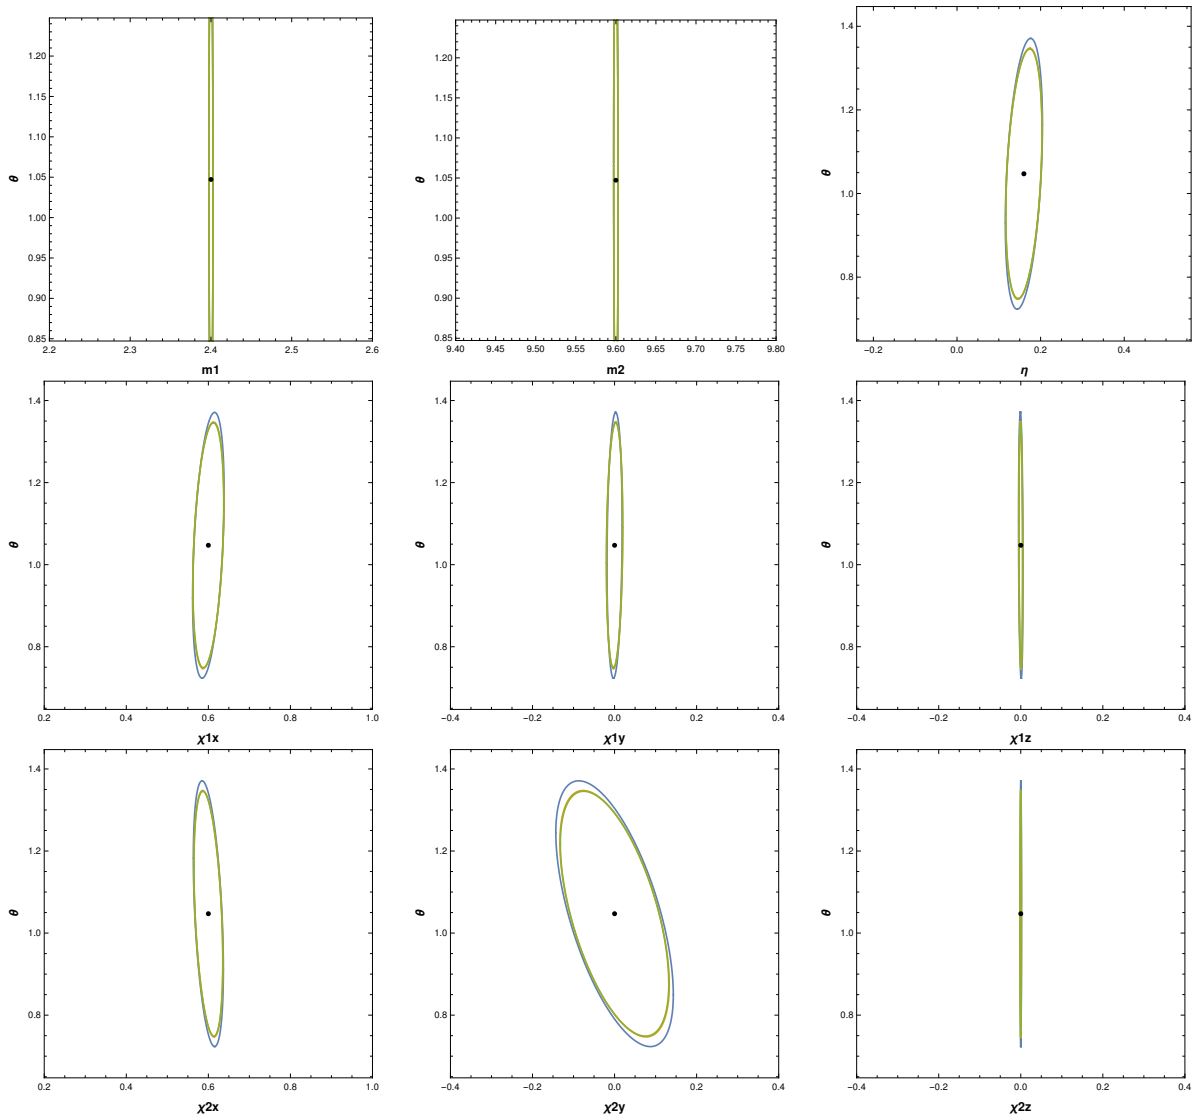

FIG. 18: All 2σ confidence ellipses corresponding with the inclination angle of the binary θ for the precessing system used in Sec. IV. The blue ellipses show the results for the $(\ell = 2)$ waveform, the orange ellipses for the $(\ell = 2, 3)$ waveform, and the green ellipses for the $(\ell = 2, 3, 4)$-waveform.

2. Error Ellipses for Spin Components and Mass Parameters

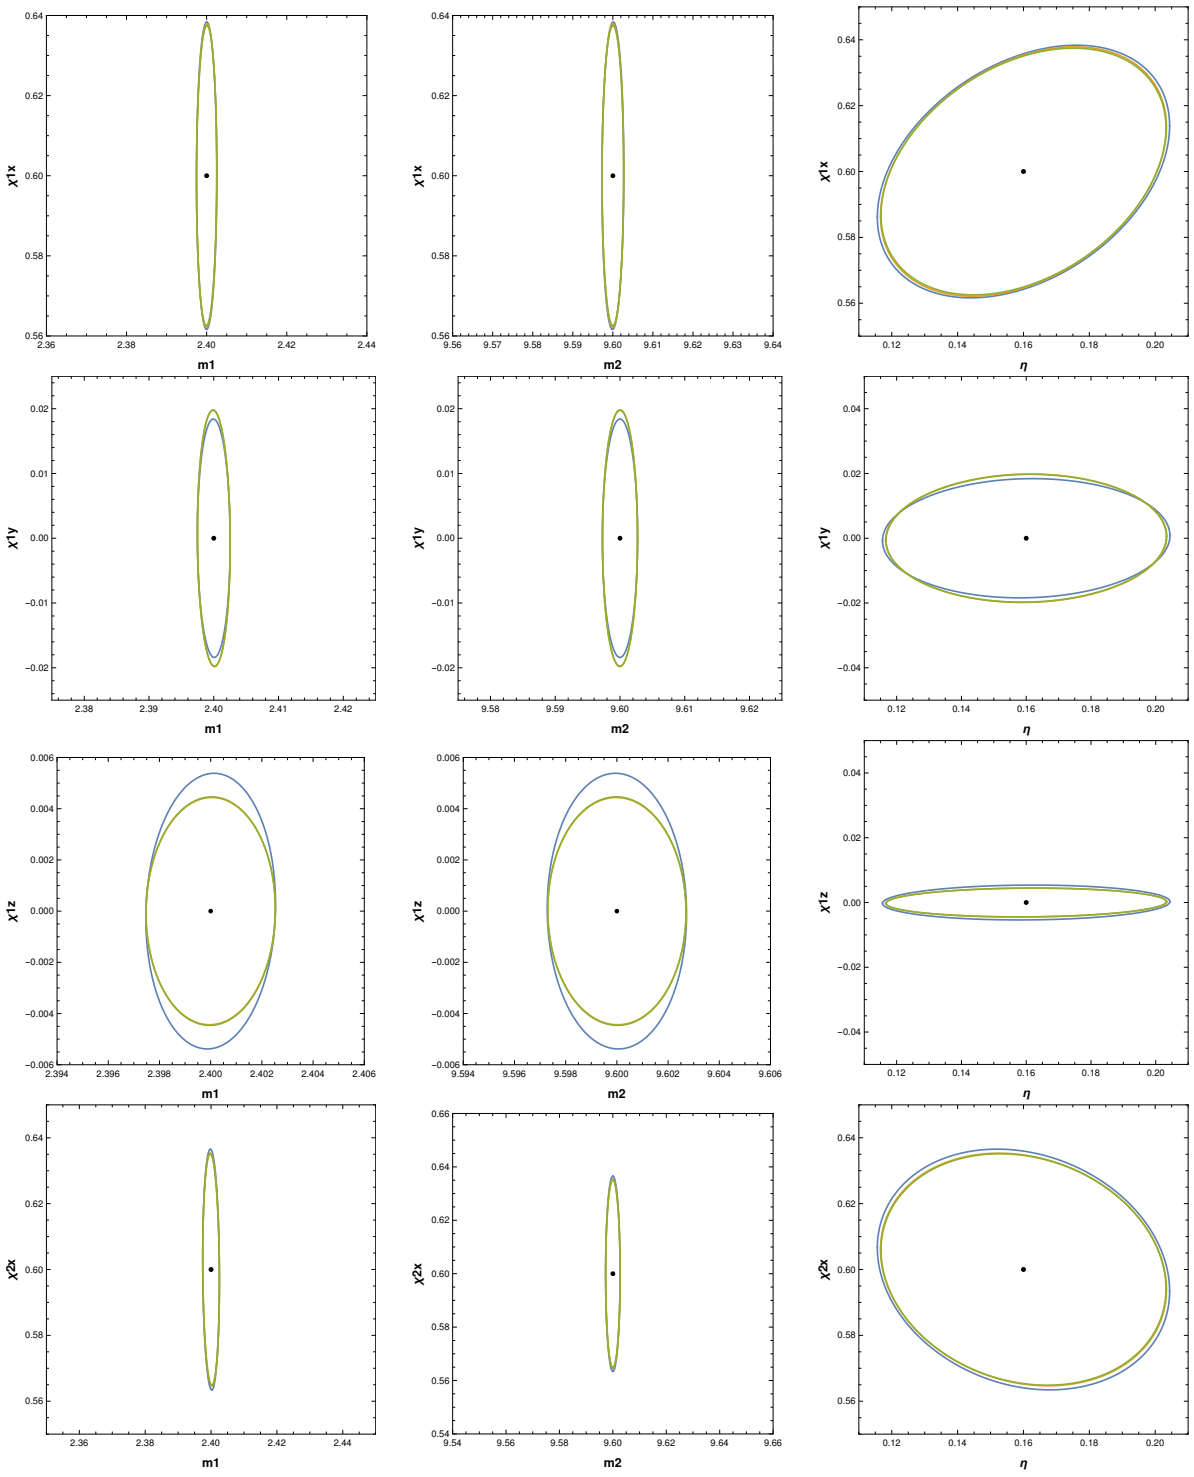

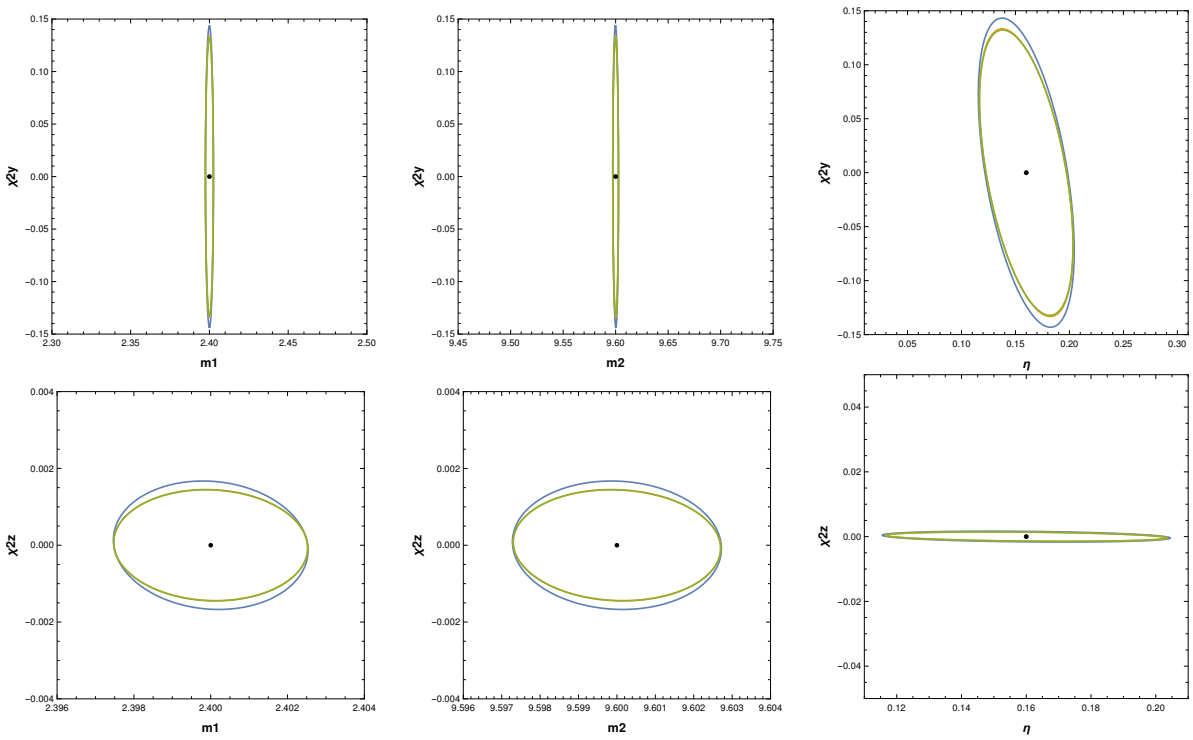

FIG. 20: All 2-σ confidence ellipses in the plane of a spin component and a mass parameter for the precessing system as discussed in Sec. IV. The blue ellipses show the results for the $(\ell = 2)$ waveform, the orange ellipses for the $(\ell = 2,3)$ waveform, and the green ellipses for the $(\ell = 2,3,4)$-waveform.

[1] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015.

[2] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016.

[3] B. P. Abbott et al. GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence. *Phys. Rev. Lett.*, 116(24):241103, 2016.

[4] Luc Blanchet. Gravitational radiation from post-newtonian sources and inspiralling compact binaries. *Living Reviews in Relativity*, 17(2), 2014.

[5] Kostas D. Kokkotas and Bernd G. Schmidt. Quasi-normal modes of stars and black holes. *Living Reviews in Relativity*, 2(2), 1999.

[6] Emily Baird, Stephen Fairhurst, Mark Hannam, and Patricia Murphy. Degeneracy between mass and spin in black-hole-binary waveforms. *Phys. Rev.*, D87(2):024035, 2013.

[7] R. O'Shaughnessy, B. Farr, E. Ochsner, Hee-Suk Cho, C. Kim, and Chang-Hwan Lee. Parameter estimation of gravitational waves from nonprecessing black hole-neutron star inspirals with higher harmonics: Comparing Markov-chain Monte Carlo posteriors to an effective Fisher matrix. *Phys. Rev.*, D89(6):064048, 2014.

[8] Philip B. Graff, Alessandra Buonanno, and B.?S. Sathyaprakash. Missing Link: Bayesian detection and measurement of intermediate-mass black-hole binaries. *Phys. Rev.*, D92(2):022002, 2015.

[9] Emanuele Berti, Vitor Cardoso, Jose A. Gonzalez, Ulrich Sperhake, Mark Hannam, Sascha Husa, and Bernd Bruegmann. Inspiral, merger and ringdown of unequal mass black hole binaries: A Multipolar analysis. *Phys. Rev.*, D76:064034, 2007.

[10] R. O'Shaughnessy, Benjamin Farr, E. Ochsner, Hee-Suk Cho, V. Raymond, Chunglee Kim, and Chang-Hwan Lee. Parameter estimation of gravitational waves from precessing black hole-neutron star inspirals with higher harmonics. *Phys. Rev.*, D89(10):102005, 2014.

[11] Ilya Mandel, Christopher P. L. Berry, Frank Ohme, Stephen Fairhurst, and Will M. Farr. Parameter estima-

tion on compact binary coalescences with abruptly terminating gravitational waveforms. *Class. Quant. Grav.*, 31:155005, 2014.

[12] K. G. Arun, Alessandra Buonanno, Guillaume Faye, and Evan Ochsner. Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms. *Phys. Rev.*, D79:104023, 2009. [Erratum: Phys. Rev.D84,049901(2011)].

[13] Patricia Schmidt, Mark Hannam, and Sascha Husa. Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and non-precessing inspiral signals. *Phys. Rev.*, D86:104063, 2012.

[14] Patricia Schmidt, Frank Ohme, and Mark Hannam. Towards models of gravitational waveforms from generic binaries: Ii. modelling precession effects with a single effective precession parameter. *Phys. Rev. D*, 91:024043, Jan 2015.

[15] Hee-Suk Cho, Evan Ochsner, Richard O'Shaughnessy, Chunglee Kim, and Chang-Hwan Lee. Gravitational waves from black hole-neutron star binaries: Effective Fisher matrices and parameter estimation using higher harmonics. *Phys. Rev.*, D87(2):024004, 2013.

[16] D. Coe. Fisher Matrices and Confidence Ellipses: A Quick-Start Guide and Software. *ArXiv e-prints*, June 2009.

[17] Frank Ohme, Alex B. Nielsen, Drew Keppel, and Andrew Lundgren. Statistical and systematic errors for gravitational-wave inspiral signals: A principal component analysis. *Phys. Rev.*, D88(4):042002, 2013.

[18] Frank Ohme. *Bridging the Gap between Post-Newtonian Theory and Numerical Relativity in Gravitational-Wave Data Analysis*. PhD thesis, Potsdam, Max Planck Inst., 2012.

[19] Vijay Varma, Parameswaran Ajith, Sascha Husa, Juan Calderon Bustillo, Mark Hannam, and Michael Pürrer. Gravitational-wave observations of binary black holes: Effect of nonquadrupole modes. *Phys. Rev. D*, 90:124004, Dec 2014.

[20] P. Ajith et al. Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins. *Phys. Rev. Lett.*, 106:241101, 2011.

[21] Mark Hannam, Patricia Schmidt, Alejandro Bohé, Leïla Haegel, Sascha Husa, Frank Ohme, Geraint Pratten, and Michael Pürrer. Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms. *Phys. Rev. Lett.*, 113(15):151101, 2014.

[22] Wolfram. Numerical nonlinear global optimization, 2016.

[23] The gravitational wave radiation field can be decomposed into spin weighted spherical harmonic modes, and, depending on the binary configuration, some modes may carry more energy and therefore accessible information than others. Many of the currently employed template waveforms include only a few, dominant modes of the gravitational wave radiation field, in what is called the quadrupole approximation.

[24] We refer to the angle between total angular momentum and the line-of-sight as the inclination angle θ.