Agnostic black hole spectroscopy: Quasinormal mode content of numerical relativity waveforms and limits of validity of linear perturbation theory

Vishal Baibhav[®], Mark Ho-Yeuk Cheung[®], Emanuele Berti[®], Vitor Cardoso[®], A Gregorio Carullo[®], 4.5.6.7 Roberto Cotesta[®], Walter Del Pozzo, and Francisco Duque[®]

¹Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA),

1800 Sherman Avenue, Evanston, Illinois 60201, USA

²William H. Miller III Department of Physics and Astronomy, Johns Hopkins University,

3400 North Charles Street, Baltimore, Maryland 21218, USA

³CENTRA, Departamento de Física, Instituto Superior Técnico—IST, Universidade de Lisboa—UL,

Avenida Rovisco Pais 1, 1049-001 Lisbon, Portugal

⁴Niels Bohr International Academy, Niels Bohr Institute, Blegdamsvej 17, 2100 Copenhagen, Denmark

⁵Theoretisch-Physikalisches Institut, Friedrich-Schiller-Universität Jena, 07743, Jena, Germany

⁶Dipartimento di Fisica "Enrico Fermi," Universitä di Pisa, Pisa I-56127, Italy

(Received 11 February 2023; accepted 8 September 2023; published 14 November 2023)

Black hole spectroscopy is the program to measure the complex gravitational wave frequencies of merger remnants, and to quantify their agreement with the characteristic frequencies of black holes computed at linear order in black hole perturbation theory. In a "weaker" (nonagnostic) version of this test, one assumes that the frequencies depend on the mass and spin of the final Kerr black hole as predicted in perturbation theory. Linear perturbation theory is expected to be a good approximation only at late times, when the remnant is close enough to a stationary Kerr black hole. However, it has been claimed that a superposition of overtones with frequencies fixed at their asymptotic values in linear perturbation theory can reproduce the waveform strain even at the peak. Is this overfitting, or are the overtones physically present in the signal? To answer this question, we fit toy models of increasing complexity, waveforms produced within linear perturbation theory, and full numerical relativity waveforms using both agnostic and nonagnostic ringdown models. We find that higher overtones are unphysical; their role is mainly to "fit away" features such as initial data effects, power-law tails, and (when present) nonlinearities. We then identify physical quasinormal modes by fitting numerical waveforms in the original, agnostic spirit of the no-hair test. We find that a physically meaningful ringdown model requires the inclusion of higher multipoles, quasinormal mode frequencies induced by spherical-spheroidal mode mixing, and nonlinear quasinormal modes. Even in this "infinite signal-to-noise ratio" version of the original spectroscopy test, there is convincing evidence for the first overtone of the dominant multipole only well after the peak of the radiation.

DOI: 10.1103/PhysRevD.108.104020

I. INTRODUCTION

A striking aspect of black hole (BH) perturbation theory is its formal analogy with quantum mechanics. This analogy follows from the fact that after separation of the angular variables using tensor spherical harmonics with angular indices (ℓ, m) , the scattering of gravitational waves (GWs) off a Schwarzschild BH becomes formally equivalent to a Schrödinger-like equation with a potential barrier. This is true for both odd-parity [1] and even-parity [2,3] perturbations, which fully characterize the linear dynamics of the Schwarzschild spacetime.

Once the boundary conditions for the scattering problem were understood, Vishveshwara realized that the response of the BH to an incoming pulse of radiation is characterized by a superposition of damped exponentials with discrete frequencies and damping times, now commonly known as the "ringdown" by analogy with the dying tones of a vibrating bell [4]. The damping occurs because (unlike many textbook problems in quantum mechanics) the BH scattering problem is not self-adjoint; BH spacetimes absorb gravitational radiation at the horizon and emit radiation at spatial infinity—hence the name "quasinormal" modes (QNMs), as opposed to the "normal" modes of self-adjoint physical systems [5–8].

The correspondence between BH spectra and atomic spectra was repeatedly used at the formal level in the development of QNM theory during the 1970s. Press [9] used the analogy to prove that the BH would not "divest itself of the unwanted perturbations in a single large belch,"

but rather radiate initial perturbations with multipole index ℓ "only gradually, yielding a long and nearly sinusoidal wave train of gravitational radiation." The QNMs identified by Press were shown to play an important role in physical processes producing gravitational radiation—for example, when particles fall radially into the BH [10].

Chandrasekhar and Detweiler [11] used again the quantum mechanical analogy to prove the isospectrality of even and odd perturbations, and to compute the so-called "overtones" of a Schwarzschild BH. For given angular indices (ℓ, m) , there is a whole "tower" of QNM frequencies $\omega_{\ell mn}$ that can be sorted by the magnitude of their imaginary part. Typically n = 0 denotes the fundamental mode, and increasing values of n correspond to larger imaginary parts and shorter damping times. Deeper connections between the quantum mechanical scattering problem and the gravitational scattering problem emerged in the work by Ferrari et al. [12–15]. After Teukolsky's breakthrough proof of the separability of the perturbation equations for rotating (Kerr) BHs [16-19], Detweiler pointed out that the fundamental QNM frequency of a Kerr BH (the one with the smallest imaginary part and longest damping time) depends only on its mass and spin [20], so—at least conceptually—the relation can be inverted to identify the Kerr BH parameters from a knowledge of the frequency and damping time.

While GW astronomy was a long time coming, the potential observational implications were clear to the first theorists studying the gravitational spectrum of Kerr BHs. Detweiler's landmark calculation of the Kerr QNM spectrum ends with a remarkably prophetic statement; "After the advent of GW astronomy, the observation of [the BH's] resonant frequencies might finally provide direct evidence of BHs with the same certainty as, say, the 21 cm line identifies interstellar hydrogen" [21]. An even deeper analogy between Kerr perturbations and the quantum mechanical treatment of the H₂ ion [22,23] was exploited by Leaver to compute the Kerr spectrum with high accuracy using continued fraction techniques [24].

Practical attempts to implement the spectroscopy program in GW data analysis started later. Echeverria quantified the measurability of the frequency and damping time of the fundamental mode [25] (see also [26]). The merger of two comparable-mass BHs was identified early on as one of the most promising LIGO-Virgo sources [27], but predicting which QNMs would be excited as a result of the merger was essentially a matter of guesswork before the first numerical BH merger simulations.

In the late 1990s, a conjectured correspondence between large-*n* QNMs and BH area quantization [28,29] triggered several studies of overtones in various BH spacetimes. By the time Dreyer and collaborators introduced the term "BH spectroscopy" [30], the idea had been explored by the GW community for decades (there was also a flourishing industry of research on "GW asteroseismology," trying

to infer the properties of ultradense matter from the analogous problem for compact stars, see e.g. [6,31]). The fact that overtones should generically contribute to the GW signal had been proven in many different contexts. For example, the classic study of Oppenheimer-Snyder collapse to a Schwarzschild BH by Cunningham, Price, and Moncrief clearly identified the first overtone (see Fig. 12 of Ref. [32]). Inclusion of one overtone was shown to better fit the waveforms from infalling particles with finite kinetic energy in the Schwarzschild case [33], as well as the waveforms from the first simulations of rotating collapse to a Kerr BH [34].

The fact that multiple modes contribute to the ringdown offers the opportunity to characterize the remnant as a Kerr BH. The idea of BH spectroscopy is quite simple (see [6,7,35,36] for reviews). In general relativity (GR), the two GW polarizations $h_{+\times}$ can be decomposed as

$$h_{+} - ih_{\times} \equiv \sum_{\ell m} h_{\ell m}(t)_{-2} Y_{\ell m}(\iota, \phi), \tag{1}$$

where the spin-weighted spherical harmonics $_{-2}Y_{\ell m}(\iota,\phi)$ depend on two angles that characterize the direction from the source to the observer.¹

Within linearized BH perturbation theory, Leaver [39] proved that each multipolar component of the waveform at intermediate times—after the "prompt response," and before the onset of power-law tails—is described by a superposition of QNMs,

$$h_{\ell m}(t) \equiv \sum_{n} A_{\ell m n} e^{-i[\omega_{\ell m n}(t - t_{\text{start}}) + \phi_{\ell m n}]}.$$
 (2)

Here, $t_{\rm start}$ is an arbitrary starting time. In linearized GR, the complex Kerr QNM frequencies $\omega_{\ell mn}$ depend only on the remnant BH mass M_f and dimensionless spin χ_f , but not on the nature of the perturbation, and are known to very high accuracy [40,41]. On the contrary, the QNM amplitudes $A_{\ell mn}$ and phases $\phi_{\ell mn}$ depend on the astrophysical process causing the perturbation.

Green's function techniques imply that the QNM amplitudes $A_{\ell mn}$ can be factorized as a product of complex "excitation factors" $B_{\ell mn}$ that depend only on the remnant's mass and spin and complex-valued, initial data-dependent integrals $I_{\ell mn}$ [39,42–47]. However, the excitation amplitudes of the overtones in a binary merger were unknown before the first numerical BH merger simulations. Heuristic arguments suggested that comparable-mass BH mergers may have ringdown signal-to-noise ratio (SNR) roughly comparable to the inspiral SNR [48], while other astrophysical processes would be much less efficient at exciting QNMs [49]. Early work on BH spectroscopy trying to

¹As we discuss below, spin-weighted spheroidal harmonics ${}_{-2}S_{\ell m}(\iota,\phi)$ are more appropriate to study perturbations of Kerr BHs. This produces spherical-spheroidal mode mixing [37,38].

quantify the measurability of different multipoles and different overtones had to rely on educated guesses [40].

Our understanding of ringdown excitation changed dramatically after the 2005 numerical relativity (NR) breakthrough [50-52]. Fits of NR simulations revealed that the radiation from a BBH merger is dominated by the $\ell = |m| = 2$ spherical harmonic multipole, while higher multipoles are subdominant [53,54]. A superposition of overtones with frequencies fixed at the values predicted for the asymptotic Kerr remnant can fit the waveform even before the peak, but Ref. [53] questioned the physical meaning of extending these fits to the peak of the radiation: "The Kerr QNM frequencies and decay constants are computed assuming that the mass and angular momentum they carry away constitute a negligible perturbation on the system. This raises the question as to whether or not the radiated energy and angular momentum are affecting the QNM fits. This issue will, of course, become more significant as the fits are pushed to earlier times."

As our understanding of QNM excitation has improved since 2005, so have GW data analysis techniques. The first detection of GWs from the binary black hole (BBH) merger GW150914 [55] marked the beginning of a new era in astronomy. Since then, the LIGO-Virgo-KAGRA (LVK) Collaboration [56–58] has reported 90 events of probable astrophysical origin during the first three observing runs [59–62]. These GW signals, in combination with those detected by independent groups [63–67], can be used for various tests of GR in the strong-field regime [68–70].

While in principle a perfect knowledge of the dominant QNM frequency in a BH binary merger (which we now know to be the mode with $\ell = m = 2$, n = 0) could be used to infer both the mass and the spin of the remnant, in practice a single mode is not sufficient to get accurate and unbiased values for these quantities; mass and spin estimates can and should be improved by combining different multipolar components [71] and by including overtones [72]. Since multiple modes will always be excited to some extent, we must first understand which combination of modes will dominate the signal [40]. Are we going to observe a combination of low-n modes for different multipoles, or are higher overtones of the $\ell=m=2$ component dominant? Can we measure the frequencies and damping times accurately enough to resolve the modes? The answers to these questions depend on the properties of the merger remnant progenitors and on the sensitivity of the detectors [54,72–85].

Initial estimates of the detectability and resolvability of different modes used a Fisher matrix approximation, which is only valid for large SNRs [25,26,40]. SNR estimates based on NR-calibrated amplitudes were instead presented in [71,73,74]. In the first "modern" Bayesian treatment of BH spectroscopy [86], the ringdown was parametrized in terms of mass, spin and a single deviation parameter, reducing the number of free parameters and related

correlations. A search for beyond-Kerr signatures based on this model is more sensitive, but less generic than one where all the modes are measured independently (i.e., the "classical" formulation of BH spectroscopic tests; see the discussion in Sec. VII below). The "resolvability" of the modes was quantified in terms of Bayes factors. Reference [87] introduced an improved model for ringdown amplitudes and sources, focusing on the Einstein Telescope (ET) detector [88]. After the first detections, updated forecasts of our ability to observe the fundamental modes with $\ell = m = 3$ and $\ell = m = 4$ were computed within a Bayesian framework. Reference [89] concluded that subdominant fundamental modes with an amplitude of 0.1 (0.3) relative to the fundamental mode with $\ell=m=2$ could be detected with SNR of 30 (15) in the late-time ringdown without assuming NR constraints on the amplitudes. Relying on a model calibrated to NR [90], these estimates were revised in Ref. [91]. This work confirmed previous Fisher-based estimates [92] (see also [83]), concluding that Voyager-class detectors would be necessary to have "decisive" Bayesian evidence for the presence of two modes in a single detection.

All of these studies focused on ringdown-only parameter estimation, ignoring the preringdown signal. The first study of ringdown signals from complete inspiral-merger-ringdown simulations windowed the signals at the ringdown start time to avoid spurious preringdown contributions in the frequency domain, and showed that percent-level "nohair" tests are possible by combining multiple loud sources detected by the LIGO-Virgo network [93]. The need to tune specific windowing parameters can be overcome by formulating the test directly in the time domain. The uncorrelated case was considered in Ref. [94], while the case of a nondiagonal autocovariance matrix was tackled in Ref. [95], applying the time-domain method to search for multiple (fundamental) modes in GW150914 and constraining parametric deviations from the GR spectrum. This formalism was extended by adopting a truncated likelihood formulation with a fixed ringdown starting time in Ref. [96], and applied to the search of the first ringdown overtone in GW150914 data. These works led to the construction of a full ringdown pipeline built around the pyRing package [95,97], which was used by the LVK Collaboration to produce a catalog of ringdown observations and to search for beyond-GR signatures [69,70,98,99]. The searches employed models of increasing complexity, ranging from agnostic superpositions of damped sinusoids to templates calibrated against BBH simulations [78]. The pyRing pipeline was later applied to search for BH charges [100], set bounds on the BH information emission mechanism [101], investigate heuristic models of "area-quantized" BHs [102], and constrain modified gravity corrections in a perturbative framework [103,104]. The truncated formulation has also been implemented in the RINGDOWN package [105].

In data analysis, the omission of overtones may lead to biases in the remnant mass and spin estimates [72]. However, QNM tests often relied only on fundamental modes with different values of (ℓ, m) for two main reasons. The first reason is practical; overtones are short-lived and difficult to confidently identify in the data [75]. The second reason (and the focus of this work) is conceptual; it is unclear whether multiple overtones have physical meaning, or they just happen to phenomenologically fit the nonlinear part of the merger signal [53].

Giesler *et al.* [106] focused on the $\ell = m = 2$ multipole of the radiation, and showed that including overtones up to n = 7 in the ringdown model improves the mismatch with NR simulations for all times $t \ge t_{\text{peak}}$. Here the "peak time" $t_{\rm peak}$ is defined as the time at which $|h_+^2 + h_\times^2|$ has a maximum. According to Ref. [106], the inclusion of higher overtones "provides an unbiased estimate of the true remnant parameters" and the low mismatch "implies that the spacetime is well-described as a linearly perturbed BH with a fixed mass and spin as early as the peak." This emphasis on linearity prompted a sequence of additional investigations, both on the modeling and on the observational side [81-83,107-116]. Some works extended the idea to counter-rotating modes and higher multipoles [85], included an even larger numbers of overtones [117], and proposed possible explanations for the apparent simplicity of the signal [110,118].

If higher overtones could be measured by starting at the peak, the larger ringdown SNR would open the door to more precise tests of GR. This argument (that we will challenge below) motivated a reanalysis of GW150914 [96] where the postpeak waveform was fitted with a QNM superposition including overtones, claiming evidence for "at least one overtone [...] with 3.6σ confidence." This claim is at odds with Ref. [109] and with the subsequent LVK analysis [69], both reporting weak evidence (with a log-Bayes factor of ~ 0.6) in favor of the model including two modes relative to the model including only one. A recent reanalysis of GW150914 found no evidence in favor of an overtone in data after the peak [119]. Around the peak, the log-Bayes factor does not indicate the presence of an overtone, while the support for a nonzero amplitude is sensitive to changes in the starting time much smaller than the overtone damping time. GW150914-like injections in neighboring segments of the real detector noise suggest that noise can artificially enhance evidence for an overtone. The matter was further debated in Refs. [120-123], using different choices for the likelihood, noise estimation, sampling rate and analysis duration.

In this paper we set aside the controversial issue of identifying overtones in real data (but we discuss the implications of our findings on this debate in Sec. VII), and we carefully reanalyze the main conclusion of Ref. [106]; can we reliably conclude from an analysis of NR simulations that the entire postpeak waveform is

described by a superposition of overtones consistent with linear BH perturbation theory on a fixed Kerr background? To improve readability, we present our main conclusions in the following executive summary.

A. Executive summary

The main goal of the paper is to understand which physical modes can be extracted from GW signals and at which point in time, with a special (but nonexclusive) focus on overtones. To this end, we will often compare two broad classes of models: one in which all complex frequencies $\omega_{\ell mn}$ are fixed to match theoretical predictions from BH perturbation theory [40,41], and one in which some (or all) of these frequencies are left free to vary. The first class leads to a "weak" version of BH spectroscopy tests (often used in recent investigations), while the second corresponds to the original, "agnostic" spectroscopy proposal.

While the nonagnostic method leads to an easier extraction of the modes, it must be used with great care; by a priori assuming the presence of short-lived modes, it can often lead to false evidence for unphysical contributions, as in the case of higher overtones. Generically, when looking for overtones, we will show that it is very important to take into account subdominant contributions to the waveform, such as tails in linear theory or spherical-spheroidal mode mixing in full NR simulations. The notation used for our fitting models is introduced in Sec. II, and it is summarized in Table I.

1. Extracting overtones in linearized theory

In Sec. III we show that the "original" BH spectroscopy test, where one agnostically searches for multiple QNMs in the data, is hard to carry out using overtones even within linear perturbation theory (and in the absence of noise). To understand why, we introduce three toy models of increasing complexity, such that linear perturbation theory is valid by construction: (i) a hypothetical "pure ringdown" (i.e., a pure superposition of damped exponentials, which is never expected to describe a real signal)²; (ii) a "pure ringdown" model to which we add a power-law tail; and (iii) a numerical waveform constructed by a replica of the original Vishveshwara scattering experiment [4]. These toy models illustrate that even within linear perturbation theory, it is not sensible to fit the waveform with ~7 overtones starting at the peak of the radiation.

²In Appendix A we use Green's function techniques to make two important conceptual points. First, a "pure ringdown" never exists: other components (including prompt emission and back-scattering) will always affect the waveform and cause amplitudes to change in time. Second, even in the linear regime (and even for a Minkowski background) the "ringdown starting time" is ill-defined; it depends on which of the radiation properties (amplitude or energy) that we monitor, and there is no mathematical or physical basis to claim a well-defined instant as "the" ringdown starting time.

All frequencies free

4N + 6

Model Description Number of modes Number of parameters Mode mixing $Q_N(t)$ All frequencies fixed No N+12(N+1) N_f free frequencies N+1 $2(N + N_f + 1)$ $Q_{N,N_f}(t)$ No $2(N + N_f + 2)$ $Q_{N,N_{\varepsilon}}^{(m)}(t)$ N_f free frequencies Yes N+2 $Q_N^{(f)} \equiv Q_{N,N+1}$ All frequencies free N+14N + 4No

Yes

TABLE I. Fitting models used in this paper. All models include N overtones (i.e., N+1 QNMs when mode mixing is not included, and N+2 QNMs when mode mixing is included). When present, the mode-mixing contribution has fixed frequency, but free amplitude and phase.

We prove this point in two parts. We first ask can we recover the known frequencies by fitting? As a test of the fitting algorithm, we show that this is possible in the "pure ringdown" model by using a "bootstrap" procedure: i.e., we first identify longer-lived modes in an agnostic way, and then we fix their complex frequencies when searching for shorter-lived ones. In the "ringdown + tail" model, however, as we increase the number of fitting modes, the minimum of the mismatch \mathcal{M} between the fitting model and the waveform keeps decreasing as we get closer to the waveform peak, even when the free mode does not approach the expected overtone frequency. Adding modes to the fitting model can reduce the mismatch even if the mode frequency is unphysical. Therefore, a small mismatch is not sufficient evidence to claim the presence of an overtone. The individual modes stop converging towards their expected values when their mismatch drops below the mismatch induced by the tail; at that point, the mode is effectively trying to fit the tail, and the mismatch saturates. Even if the overtones are physically present in the waveform (as in this toy model), fitting with many overtones is not optimal from a data-analysis point of view, because the fit is not robust against small contaminations. This is an important limitation for agnostic tests of GR; tiny contaminations can hinder our ability to extract higherovertone frequencies, even if the modes are physically present in the signal. The toy models also allow us to better understand the behavior of the fitted QNM amplitudes. For the (unrealistic) "pure ringdown" waveforms, all of the overtone amplitudes converge to constant values at late times. When we add a tail, there is one remarkable difference: the fitted amplitudes blow up exponentially at some critical time—more specifically, when the highest (fastest-decaying) overtone in the fitting model starts to pick up the contamination due to the tail.

 $Q_N^{(f,m)} \equiv Q_{NN+1}^{(m)}$

The conclusion of this exercise is clear; unless additional physics is taken into account the original, agnostic BH spectroscopy test is unfeasible for all overtones (including n=1), and only possible at late times for the fundamental mode, even within linear perturbation theory. When we consider ringdown waveforms resulting from the scattering of a Gaussian wave packet in linearized gravity, an agnostic

damped-sinusoid fit cannot recover the correct frequencies for any of the overtones. Even if we fix the frequencies to their known values, there is no convincing evidence that overtones with n > 2 are present in the signal. If we insist to use multiple overtones to test GR, we should start fitting the waveform at times significantly after the peak. Many of the lessons learned in linearized theory carry over to the full GR case.

N+2

2. Are postpeak BBH waveforms linear?

Reference [106] claimed that the waveform resulting from the merger of two comparable-mass BHs can be (i) adequately described by linear perturbation theory starting from the peak of the strain, (ii) well-modeled by a combination of QNM overtones, and (iii) used to test GR by identifying the overtones in the signal. In Sec. IV we revisit these claims. We present three different arguments against the validity of the linear approximation.

First, we show that a constant-amplitude overtone superposition does not work in the BH merger case, and the amplitudes of the overtones change significantly when we change the fitting window. This mode-amplitude evolution has two important implications: (i) overtone models with $N \ge 2$ are unphysical, because they try to overfit other features of the waveform; and (ii) models with at most one overtone ($N \le 1$) are physical, but they can only be used for meaningful spectroscopy tests at late times.

These conclusions are reinforced in three appendixes. In Appendix B we demonstrate that using NR waveforms with different resolutions and different extrapolation orders makes almost no difference in the amplitude fits, and therefore NR errors cannot explain the time variation of the fitted amplitudes. In Appendix C we investigate the effect of a spurious late-time constant observed in the SXS simulations. We find that this spurious constant has only a small impact at very late times, and that it does not affect our conclusion that it is not possible to extract the first overtone at the peak. In Appendix D we examine a complex-exponential toy model for SXS:BBH:0305 which includes an estimate of the numerical noise, confirming and strengthening the main conclusions of Sec. III.

A second issue with a linear perturbation theory interpretation of postpeak ringdown concerns overfitting. How many overtones are really necessary to minimize the mismatch between ringdown waveforms and the full inspiral-merger-ringdown waveform? Which QNMs are most effective at minimizing the mismatch and reproducing the correct values of the remnant mass and spin?

Recent studies claim that the inclusion of the fundamental mode and seven overtones provides a very accurate description of the ringdown up to the peak strain amplitude, and significantly reduces the uncertainty in the extracted remnant properties [106]. However, we show that the higher overtones lead to very small mismatches by merely overfitting the waveforms. Furthermore, we argue that these higher overtones try to fit other physics (such as time variation in the QNM amplitudes due to initial data, an evolving spacetime background, and nonlinearities) close to the merger. The addition of several overtones allows for better extraction of the fundamental mode and the first overtone, which carry most of the information about the remnant properties, because they effectively "fit away" poorly understood physics.

In Appendix E we construct an unphysical postpeak BBH waveform, and show that the overtones can still fit it with similar accuracy. The fact that using the overtones allows us to improve the measurement of χ_f and M_f even for this unphysical hybrid waveform supports the conclusion that overtones can match any early postpeak waveform portion, thus allowing the dominant mode to correctly fit the late postpeak waveform; it is the fundamental mode that really carries most of the information about the remnant BH properties.

In Sec. IV C we go beyond mismatches and ask, in the same vein—which overtones are necessary to correctly extract the remnant's properties? We swap individual modes with random damped exponentials. If the "fake" random-frequency mode still fits the waveform with similar or better accuracy, or if it still extracts the remnant properties accurately, we can conclude that the originally "swapped" overtone was not really necessary. We use this argument to show that higher overtones do not play a significant role in extracting the remnant's properties either. Overtones with $n \ge 2$ do not significantly contribute to the extraction of the remnant's parameters, and therefore there is no motivation to include them in the modeling.

3. Agnostic BH spectroscopy: Extracting complex frequencies from the waveform

Since the inclusion of several overtones leads to overfitting, in Sec. V we adopt a different strategy. Rather than imposing a priori that the known overtones associated to a given (ℓ, m) spin-weighted spheroidal harmonic are present in the (ℓ, m) spin-weighted spherical harmonic multipole of the NR data, we consider the complex frequencies as free fitting parameters. This is a much

stronger test (in fact, it is the original BH spectroscopy proposal at infinite SNR), and therefore it should lead to more robust conclusions about which modes are truly present in the data.

We first focus on the $\ell = m = 2$ spherical harmonic multipole and we ask how many QNM frequencies can we extract without assuming any (no-hair theorem enforced) relation between them? We show that (i) in general, the mismatch between the fitting model and the waveform is lower when we keep the frequencies free; (ii) many of the fitted damped exponentials robustly converge towards known QNM frequencies, naturally selecting the physical modes that contribute to the ringdown signal; and (iii) it is essential to include spherical-spheroidal mode mixing to identify the correct modes. In fact, including modes due to spherical-spheroidal mixing is essential to extract the first overtone (at late times) from the dominant (2, 2) multipole. At least three free modes are required to extract the first overtone, and it is easier to extract the long-lived fundamental modes than the fast decaying overtones, even if the latter have a much larger amplitude.

We then study a variety of different fitting models. In some of these fitting models we include the dominant spherical-spheroidal mode at the expected frequency; in others, we do not. We adopt both a "fully agnostic" strategy, in which we include more and more free modes, and a "bootstrap" strategy in which we identify modes, fix them, and then search for an additional free frequency.

The study of these different fitting models supports an important conclusion; the only identifiable physical modes in the $\ell=m=2$ multipole of the radiation are (2, 2, 0), (2, 2, 1), and (3, 2, 0). Higher overtones (N>1) cannot be robustly identified by free-frequency fitting. Furthermore, once the "physical" ringdown modes (typically the fundamental mode, the first overtone, and the mode-mixing contribution) have been fitted for, additional free modes have a tendency to simply track the characteristic "GW frequency chirp" at early times. In fact, this observation has already been used in the gravitational waveform modeling community, where "pseudo-QNMs" were introduced in the context of the effective-one-body framework to model the rapid transition from the inspiral GW frequency to the postmerger QNM frequency "plateau" [124–126].

How do the frequencies inferred by fitting free damped exponentials translate into mass and spin estimates? In Sec. V C we limit attention to the $\ell=m=2$ spin-weighted spherical multipole. We show that a free fit with two modes near the peak would give significantly biased mass and spin values (whether we include spherical-spheroidal mode mixing or not), and therefore it cannot be used for BH spectroscopy test. Even in the infinite-SNR limit, including the first overtone never allows for percent-level estimates of the mass and spin: it may allow for $\sim 10\%$ estimates of the mass and spin at late times for loud enough signals, but only if we carefully take into account mode mixing and higher multipoles.

In fact, an analysis of the $\ell=m=4$ multipole (Sec. V D) shows that nonlinearites are also important; the nonlinear (2, 2, 0)(2, 2, 0) QNM is easier to recover than the first linear overtone (4, 4, 1) in a free-frequency fit. In Appendix F we study other subdominant multipoles.

4. Can the first overtone be extracted in the presence of subdominant multipoles?

In real-world data analysis problems, the waveform will in general be a superposition of several multipoles. In Sec. VI we ask: how many free modes would be necessary to extract the first overtone (2, 2, 1) in this more realistic scenario?

We find that even in the relatively optimistic case of a face-on binary, when only the (2, 2), (3, 2), and (4, 2) multipoles significantly contribute to the strain, the extraction of the first overtone (2, 2, 1) requires four free frequencies, and it relies on the successful extraction of the (long-lived) fundamental modes of subdominant multipoles. The extraction of overtones becomes significantly more difficult when the binary is not face-on. Therefore, in practical data analysis applications it would be hard to extract the first overtone from GW150914-like signals. Even when the overtone can be confidently identified, the fundamental modes of subdominant multipoles generally yield more reliable BH spectroscopy tests.

In Sec. VII we discuss the observational implications of our work; the possible role of nonlinearities and pseudospectral instabilities in destabilizing higher overtones; how parametrizing the QNM amplitudes may facilitate BH spectroscopy, and the Occam penalties associated with the inclusion of several modes; the difference between "weak" and "strong" spectroscopy tests, and the role of beyond-Kerr parametrizations of the QNM spectrum as effective detection templates or tools to constrain modified gravity theories.

II. FITTING MODELS AND NOTATION

Let us begin by introducing some notation for the fitting models adopted in the rest of this paper. We will often compare two broad classes of models—one in which the real and imaginary parts of the complex frequencies $\omega_{\ell mn}$ are fixed to match theoretical predictions from BH perturbation theory [40,41], and one in which most (or all) of these frequencies are left free to vary. The first scenario corresponds to the "weak" version of the no-hair test used in recent investigations, while the second reflects the original "stronger," i.e., more agnostic, proposal.

In the first class of models, we fit the waveform by N+1 damped exponentials with complex frequencies fixed to the QNMs $\omega_{\ell mn}$ of a remnant with known mass and spin,

$$Q_N(t) \equiv \sum_{n=0}^{N} A_{\ell mn} e^{-i[\omega_{\ell mn}(t-t_{\text{peak}})+\phi_{\ell mn}]}, \qquad (3)$$

where the only unknowns are $A_{\ell mn}$ and $\phi_{\ell mn}$. In this case, we assume that the complex frequencies $\omega_{\ell mn}$ of all modes present in the signal are known $a\ priori$. In this sum n is the overtone number. Since $Q_0(t)$ corresponds to a model with only one mode—the fundamental mode (n=0), or "zeroth overtone"—the $Q_N(t)$ model contains N overtones, and N+1 modes in total. This model has been employed, for example, in Refs. [53,106].

In a more general category of models, we will assume that the complex frequencies of some modes are known a priori, while the complex frequencies of others are unknown. In this case we fit the waveform by N+1 damped exponentials where the complex frequencies ω_{n_f} of N_f "free" modes are unknown, while the frequencies of the remaining $N-N_f+1$ modes are fixed to the QNMs $\omega_{\ell mn}$ of a remnant with known mass and spin,

$$Q_{N,N_f}(t) \equiv \sum_{n=0}^{N-N_f} A_{\ell m n} e^{-i[\omega_{\ell m n}(t-t_{\text{peak}})+\phi_{\ell m n}]} + \sum_{n_f=0}^{N_f-1} A_{n_f} e^{-i[\omega_{n_f}(t-t_{\text{peak}})+\phi_{n_f}]}.$$
 (4)

This second class of models has $2(N+N_f+1)$ fitting parameters; $2(N-N_f+1)$ real amplitudes and phases $A_{\ell mn}, \phi_{\ell mn}, 2N_f$ real amplitudes and phases A_{n_f}, ϕ_{n_f} , and the $2N_f$ complex free frequencies ω_{n_f} . The notation $Q_{N,N_f}(t)$ is a reminder that we are fitting the waveform by an N-overtone model, where N_f modes are "free" (in the sense that their frequencies are not fixed).

In order to take into account spherical-spheroidal mixing, we will sometimes need to add another mode to $Q_{N,N_\ell}(t)$,

$$Q_{N,N_f}^{(m)}(t) = \sum_{n=0}^{N-N_f} A_{\ell mn} e^{-i[\omega_{\ell mn}(t-t_{\text{peak}}) + \phi_{\ell mn}]} + \sum_{n_f=0}^{N_f-1} A_{n_f} e^{-i[\omega_{n_f}(t-t_{\text{peak}}) + \phi_{n_f}]} + A_{\ell'm0} e^{-i[\omega_{\ell'm0}(t-t_{\text{peak}}) + \phi_{\ell'm0}]}.$$
 (5)

The superscript "(m)" in $Q_{N,N_f}^{(m)}(t)$ is a reminder that we also include the mode-mixing contribution $A_{\ell'm0}e^{i\omega_{\ell'm0}(t-t_{\mathrm{peak}})+i\phi_{\ell'm0}}$. For simplicity, we will only consider the dominant mode-mixing contribution. This comes from $\ell'=\ell+1$ when $\ell=m$, and from $\ell'=\ell-1$ when $\ell>m$ and $\ell>2$. The model now has $2(N+N_f+2)$ unknowns; besides $A_{\ell mn}, \phi_{\ell mn}, A_{n_f}, \phi_{n_f}$, and the complex frequencies ω_{n_f} , we also have the amplitude and phase $A_{\ell'm0}, \phi_{\ell'm0}$ of the mode-mixing contribution on the last line. The $Q_{N,N_f}^{(m)}(t)$ model contains N+2 QNMs.

Two special subclasses of models will be of special interest below. In one subclass, all modes have free complex frequencies, and there are no fixed modes. In another subclass, we will only fix the complex frequency of the spherical-spheroidal mixing mode, while the rest of the complex frequencies are kept free. These cases correspond to setting $N_f = N+1$ in the $Q_{N,N_f+1}(t)$ and $Q_{N,N_f+1}^{(m)}(t)$ models, respectively, and we will denote them with the following shorthand notation:

$$Q_N^{(f)} \equiv Q_{N,N+1}, \qquad Q_N^{(f,m)} \equiv Q_{N,N+1}^{(m)}.$$
 (6)

For reference, the fitting models we will consider below are summarized in Table I.

As a final note, while the postmerger waveforms computed in numerical relativity that will be the main focus of this paper are complex, the linear waveform model Ψ_{num} that we will consider as a warm-up problem below is real, because we specify real initial conditions for the time evolution. When the waveform is real, we will simply consider the real part of the fitting models, but otherwise we will use the more general complex models.

III. EXTRACTING OVERTONES IN LINEARIZED THEORY

Before turning to BBH mergers in full nonlinear GR, in this section we investigate linear perturbations in GR, as well as toy models built to elucidate the main features observed in the linear regime. This will allow us to build some understanding of what to expect in the full GR nonlinear case while working in a controlled setting, and highlight the limitations of waveform fits by a superposition of damped exponentials. In fact, we will show that even when linear perturbation theory is valid by construction, it is not sensible to model the ringdown by fitting the waveform with ~7 overtones starting at the peak. More generally, although fitting the waveform at the peak with more overtones yields smaller fit residues, the model fails to pass further basic consistency checks.

We will also see that extracting high-overtone frequencies at the peak of the waveform to test GR does not yield robust results. Even when the waveform is by construction a combination of QNMs with a small contamination from other components (such as power-law tails), the high-overtone frequencies estimated by fitting are easily biased by these subdominant contributions. Many of the lessons learned in these simple settings will carry over to the full GR case.

A. Preliminaries

As we recall in Appendix A, the starting time of the ringdown regime is an ill-defined quantity even within linear perturbation theory, because the Green's function always contains additional contributions (most notably, a

prompt response, a tail due to backscattering of radiation, and effects coming from the build-up of initial data). What we want to understand now—insisting on modeling the waveform as a superposition of damped exponentials—is the precision to which one can recover a hypothetical "pure ringdown" waveform, and the physical grounds for claiming the presence or absence of overtones.

For a linearly perturbed Schwarzschild BH geometry, after separation of the angular variables and working in the time domain, the linearized Einstein field equations imply that odd-parity (or axial) perturbations are governed by the Regge-Wheeler equation [1]

$$\frac{\partial^2 \mathbf{\Psi}_{\text{num}}}{\partial r_*^2} - \frac{\partial^2 \mathbf{\Psi}_{\text{num}}}{\partial t^2} - V_{\text{RW}} \mathbf{\Psi}_{\text{num}} = 0, \tag{7}$$

where the Regge-Wheeler potential is

$$V_{\rm RW} = \left(1 - \frac{2M}{r}\right) \left[\frac{\ell(\ell+1)}{r^2} - \frac{6M}{r^3}\right],$$
 (8)

and the tortoise coordinate r_* is defined by the relation $dr/dr_* = 1 - 2M/r$. For the purpose of this discussion we focus on the dominant, quadrupolar component of the radiation ($\ell = 2$) and we denote the wavefunction as Ψ_{num} to emphasize that it is computed numerically by solving the Regge-Wheeler equation within linear BH perturbation theory. Even-parity (or polar) perturbations, governed by the Zerilli equation [2,3], are known to be isospectral to odd-parity perturbations and behave in a qualitatively similar way [127].

We compute the waveform $\Psi_{\text{num}}(t, r)$ by imposing the following initial conditions for Eq. (7):

$$\Psi_{\text{num}}(0,r) = 0, \qquad \frac{\partial \Psi_{\text{num}}}{\partial t}(0,r) = \exp\left[-\frac{(r_* - r_0)^2}{2\sigma^2}\right],\tag{9}$$

where $r_0 = 4M$ and $\sigma = M/2$. We then perform a time evolution and extract the time-domain waveform $\Psi_{\text{num}}(t)$ at future null infinity, following Refs. [128,129]. We evolve the initial data using a two-step Lax-Wendroff method with second-order finite differences [130,131], which has been extensively used in the past in the study of late-time tails in Kerr [132,133] and extreme-mass ratio inspirals [134,135].

Our goal is to analyze $\Psi_{\text{num}}(t)$ and see how well it can be fitted by a QNM model $Q_N(t)$ (here and below we will adopt the convention that the waveform we fit is denoted in boldface, while the fitting model is not).

In general, a full time evolution of the linear equation (7) will contain a "prompt response" component that depends on the initial data, as well as a late-time power law tail [39]. We will try to quantify the impact of these additional components below. For now, we just try to fit $\Psi_{\text{num}}(t, r)$ using a QNM model comprising a finite number of

exponentially damped sinusoids. Since for the moment we focus on $\ell=2$ and on a nonrotating BH background, we can drop the index ℓ and set m=0 without loss of generality. For economy of notation, in this section we denote the amplitude $A_{\ell mn}$ of a generic mode by $A_n \equiv A_{20n}$, and similarly for the phases ϕ_n and the complex frequencies ω_n . Therefore, our fitting model is

$$Q_N(t) = \sum_{n=0}^{N} A_n e^{-i[\omega_n(t - t_{\text{peak}}) + \phi_n]}, \qquad t \in (t_0, t_{\text{end}}).$$
 (10)

The range of the fit is chosen to be $(t_0, t_{\rm end})$. For all fits in this section we set $t_{\rm end} = 200M$, but we have verified that the results would not change significantly if we used a larger value for $t_{\rm end}$.

Given a best-fit model ψ to a waveform Ψ , we can quantify the goodness of fit by computing the mismatch

$$\mathcal{M} \equiv 1 - \frac{\langle \Psi | \psi \rangle}{\sqrt{\langle \Psi | \Psi \rangle \langle \psi | \psi \rangle}},\tag{11}$$

where the scalar product is defined as

$$\langle f|g\rangle = \int_{t_0}^{t_{\text{end}}} f(t)g^*(t)dt, \qquad (12)$$

and an asterisk denotes complex conjugation. In this section both $\Psi = \Psi_{\text{num}}(t)$ and $\psi = Q_N(t)$ are real because we specified real initial data, but complex conjugation will be important later on, when we will consider complex waveforms from binary BH merger simulations in full GR.

In our fits we use a Levenberg-Marquardt nonlinear least-squares algorithm. We have cross-checked our results by comparing two different implementations, using either the Python package SciPy [136] or the NonlinearModelFit function in *Mathematica*.

B. Some considerations on overfitting

By fitting Eq. (10) to the solution of Eq. (7), we found that the linear waveform $\Psi_{\text{num}}(t)$ can indeed be fitted "well" by a model with seven overtones Q_7 if the fitting range starts at the peak of $|\Psi_{\text{num}}|$, i.e. $t_0 = t_{\text{peak}}$; the mismatch between the best fit waveform and the data is small, $\mathcal{M} \sim 10^{-8}$. This result agrees with Ref. [106].

However, a problem immediately arises. By including more fitting parameters we can (in principle) decrease \mathcal{M} indefinitely, even if the new parameters are not physically well-motivated and simply overfit the signal. Specifically, when adding more overtones to our fit model, we risk overfitting the early part of the ringdown, "fitting away" any contamination close to the peak of the waveform (due e.g., to the prompt response) with the rapidly decaying overtones. Indeed, as will be clear from the following discussion, a small \mathcal{M} is a necessary but not sufficient

condition to conclude that the fitting model is consistent with the actual waveform or well-motivated.

To showcase overfitting issues, we consider a toy waveform \mathbf{Q}_7 consisting of a fundamental mode and seven overtones, where the amplitude and phase of each mode is obtained by fitting $\mathbf{\Psi}_{\text{num}}$ at $t_0 - t_{\text{peak}} = 0M$. In other words, \mathbf{Q}_7 is a reconstruction of $\mathbf{\Psi}_{\text{num}}$ with the Q_7 model (recall that we use bold symbols for the waveforms to be fitted, and normal symbols for the fitting model; in this case \mathbf{Q}_7 denotes the waveform and Q_7 is the fitting model).

By comparing the black and cyan lines in the left panel of Fig. 1 we see that the waveform is indeed well-approximated by the \mathbf{Q}_7 model as early as the peak. In the right panel we plot the fitted individual overtone modes from n=0 to n=7 of \mathbf{Q}_7 . The first few overtones ($n \leq 3$) consist, as expected, of sinusoidal oscillations modulated by an exponential decay. However, higher overtones ($3 \leq n \leq 7$) have a larger oscillation period and a shorter exponential decay time. If we focus on the early part of the waveform, the high-order overtones look like exponentials rather than damped oscillators because of their low quality factor.³

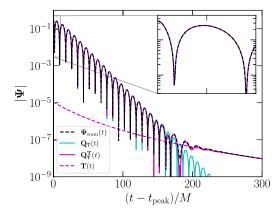
When we fit an exponentially decaying waveform with a ringdown model using a nonlinear least-squares algorithm, or when we compute the mismatch (12), the squared residue is dominated by the earlier part of the waveform. Then the higher overtones act effectively as "bumps," removing early-time parts of the waveform that could not be well-fitted by the lower overtones.

In linear perturbation theory, an important early-time contribution is the "prompt response" due to the initial wave packet that propagates directly towards null infinity without scattering off the light ring. For the BBH merger waveforms considered in later sections, the early-time waveform includes the merger phase and any nonlinearities that may be present close to the peak.

The higher overtone phases can—and often do—(anti) align to produce destructive interference between modes; this (not neessarily physical) destructive interference reduces the mismatch, allowing ringdown models with many overtones to fit the early part of the waveform. We see hints of this destructive interference in Fig. 1; note that the higher overtones can have amplitudes as high as $\sim 10^4$, much larger than the peak amplitude of the actual waveform (which is $\lesssim 1$).

The bottom line is that fitting the ringdown with many overtones requires great care because of their exponentially decaying nature. A low mismatch with a multiple-overtone fitting model can easily mislead us into believing that the overtones are physically present when, in fact, they are just an unphysical artifact that produces good fits to other

 $^{^3} The quality factor <math display="inline">\mathcal{Q}_{\ell mn} = - Re(\omega_{\ell mn})/(2 Im(\omega_{\ell mn}))$ is essentially the ratio between the decay timescale τ and the oscillation period.



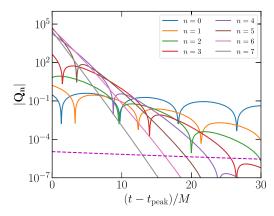


FIG. 1. Left panel; the three waveforms considered in this section. $\Psi_{\text{num}}(t)$ is the waveform extracted at future null infinity by solving the time-domain Regge-Wheeler equation (7) with the initial conditions of Eq. (9). $\mathbf{Q}_7(t)$ is a "pure ringdown" fit of $\Psi_{\text{num}}(t)$ starting at $t_0 = t_{\text{peak}}$, with a fundamental mode and seven overtones fixed at the Schwarzschild frequencies computed in perturbation theory. $\mathbf{Q}_7^{\mathbf{T}}(t)$ is the same as $\mathbf{Q}_7(t)$, but we also add a power-law tail to mimic the contamination due to backscattering of radiation. The tail amplitude is obtained by fitting $\Psi_{\text{num}}(t)$ at times t > 200. The inset shows a zoomed-in view near the peak. Right panel: the individual modes (including the fundamental mode and 7 overtones) that make up $\mathbf{Q}_7(t)$. For reference, the extrapolated tail added to $\mathbf{Q}_7^{\mathbf{T}}(t)$ (purple dashed line) is plotted in both panels. Note that in the right panel, the overtones with $n \geq 5$ drop below the extrapolated tail before completing even half of an oscillation cycle.

components of the signal. Therefore, a small mismatch is not sufficient to argue that the fitting model is a good representation of the waveform. At the very least, we should test whether the complex overtone frequencies assumed in our fitting model are actually those that best fit the waveform. Moreover, the fits should not be very sensitive to the choice of t_0 and $t_{\rm end}$. The fitted amplitudes and phases should be consistent within about one period of oscillation of the corresponding QNMs, to exclude the possibility that overtones are just "bumps" fitting nonperiodic components of the waveform with their first half-a-period.

In fact, even in linear perturbation theory (where we neglect any possible nonlinearities associated to BBH mergers), the amplitudes of the QNMs are expected to vary close to the peak of the signal due to initial data effects (see Appendix A). The Q_7 model (a superposition of constant-amplitude damped sinusoids) is insufficient to consistently fit the whole ringdown even in this simple case, and we will observe amplitude modulations near the peak as we vary the starting time of the fit.

C. Controlled fitting experiments

Let us perform some controlled experiments to test whether model Q_7 (with fixed frequencies) is a good representation of the linear ringdown waveform found by solving Eq. (7). Here we build analytical toy-model waveforms to mimic the main features of the actual linear waveform, and to understand what our fitting procedure would return if the simulated waveform were exactly described by a combination of damped sinusoids. We consider three such toy-models: (a) a "pure" ringdown waveform consisting only of damped sinusoids; (b) a waveform consisting of damped sinusoids plus a Price

power-law tail [137], as expected in linear perturbation theory; and (c) the actual linear waveform found by solving Eq. (7).

The details of the waveforms are as follows:

(a) $\mathbf{Q}_7(t)$, pure damped sinusoids. This waveform is a particular realization of the (real-valued) fitting model Q_7 , as specified in Eq. (10), whose mismatch with $\Psi_{\text{num}}(t)$ is $\sim 10^{-8}$,

$$\mathbf{Q}_{7}(t) \equiv \operatorname{Re} \sum_{n=0}^{7} A_{n} e^{-i[\omega_{n}(t-t_{\text{peak}})+\phi_{n}]}.$$
 (13)

The frequencies ω_n are the overtone frequencies of the $\ell=2$ multipole of a Schwarzschild BH, i.e., $\omega_n \equiv \omega_{20n}$. The amplitudes A_n and phases ϕ_n are fixed by fitting $\Psi_{\text{num}}(t)$ with the Q_7 model at $t-t_{\text{peak}}=0M$, as explained earlier and shown in the left panel of Fig. 1. The toy waveform is denoted in bold fonts to distinguish it from the fitting model.

(b) $\mathbf{Q_7^T}(t)$, damped sinusoids with contamination from a power-law tail. The waveform in linear theory must contain contributions from a Price power-law tail due to backscattering of GWs, which scales as $t^{-2\ell+3}$ at late times for perturbing fields of any spin, and for momentarily static initial data such as those used here [137]. At late times, the power-law tail dominates the signal. To understand the dominant ($\ell=2$) mode, we model the tail of the numerical waveform $\Psi_{\text{num}}(t)$ by

$$\mathbf{T}(t) \equiv A_{\text{tail}}(t - t_{0,\text{tail}})^{-5}.$$
 (14)

A fit of $\Psi_{\text{num}}(t)$ with this model at $t - t_{\text{peak}} > 200M$ yields $A_{\text{tail}} \approx 9.26 \times 10^4$ and $t_{0,\text{tail}} \approx -96.9$.

This backscattering is expected to affect also the earlier part of the ringdown, and we approximate its effects by extrapolating the fitted tail to early times $t \in (0M, 200M)$. Then, we can construct a more realistic toy model by adding the extrapolated tail to $\mathbf{Q}_7(t)$,

$$\mathbf{Q}_{7}^{\mathbf{T}}(t) = \mathbf{Q}_{7}(t) + \mathbf{T}(t). \tag{15}$$

This toy waveform is shown as the solid purple curve in the left panel of Fig. 1, while the extrapolated tail is shown by dashed purple lines in both the left and right panels. Our approximation of the backscattering effects is admittedly crude. Nonetheless, this toy waveform should be sufficient to understand whether backscattering can affect the fits. In fact, we have found qualitatively similar results when we replace the power-law tail by a constant shift in amplitude, by a Gaussian noise floor, or by the estimated numerical noise error intrinsic in $\Psi_{\text{num}}(t)$. We decided to show results for a power-law tail (instead of numerical noise contributions) because the extrapolated tail has larger amplitude than the noise in the time range of interest.

(c) $\Psi_{\text{num}}(t)$, linear perturbation theory waveform. This is the waveform found by solving Eq. (7) with the initial conditions Eq. (9), extracted at null infinity.

D. Recovering QNM frequencies through an agnostic fit

We now test whether we can recover the QNMs in the above toy waveforms in a frequency-agnostic manner; we consider a fitting model Q_{N,N_f} in which some of the QNM frequencies are free parameters to be determined by the fit, instead of fixing all of the frequencies to the theoretically predicted overtone frequencies. If the waveform is well described by a superposition of QNMs with certain overtone frequencies, we should be able (at least in principle) to recover these frequencies with our fitting procedure.

1. Pure damped sinusoids

Consider first the simplest toy model waveform $\mathbf{Q}_7(t)$, a superposition of damped sinusoids that mimic the full linear waveform $\mathbf{\Psi}_{\text{num}}(t)$. Recall that the frequencies of the QNMs used to construct $\mathbf{Q}_7(t)$ are those of the first seven overtones of a Schwarzschild BH.

We start by fitting the waveform with a single damped sinusoid, leaving the real and imaginary parts of the QNM frequency as two free parameters to be recovered by the fit. In our notation, then, the fitting model is $Q_{0,1}(t)$, and the single free mode should converge to the fundamental mode when the fit starts at a sufficiently late time. As shown by the blue curve in the top left panel of Fig. 2, the fitted frequency converges to the theoretically predicted fundamental QNM frequency of Schwarzschild BHs. In the second row we plot the deviation

$$\delta\omega = M|\omega_r + i\omega_i - \omega_{\text{ref}}|,\tag{16}$$

where $\omega_{\rm ref}$ denotes the reference value of the complex QNM frequency that we expect to find (ω_{200} in this case). The blue curve in the left panel of the second row shows that $\delta\omega \sim 10^{-4}$ when $t_0 - t_{\rm peak} \sim 30M$.

We can now add one more QNM to our fitting model. The least demanding procedure to look for the first overtone is to fix the frequency of the fundamental QNM in the fit. Now our fitting model is $Q_{1,1}(t)$; a model with one overtone (plus the fundamental mode, so two modes in total), in which one complex frequency (the frequency of the first overtone) is a free fitting parameter. As shown by the orange curve in Fig. 2, the free mode converges to the first overtone frequency as expected, this time with $\delta\omega \sim 10^{-3}$ when $t_0-t_{\rm peak} \sim 30M$.

We repeat this procedure by adding more modes to our model. Each time we fix the frequencies of the modes we have already recovered, leaving only one mode frequency free. When we use the fitting models $Q_{2,1}(t), Q_{3,1}(t), ...$, we find that we can recover the frequencies to an accuracy $\delta\omega\lesssim 10^{-2}$ up to the third overtone, even when the actual waveform always contains seven overtones. For overtone numbers $N\geq 4$, however, the fitted frequency becomes noisy and diverges before we reach $t_0-t_{\rm peak}\sim 30M$. The $|\delta\omega|$ curves are truncated at the point where the divergence occurs. For overtone numbers $N\geq 6$, the fits become very computational costly and they do not converge well, so we do not include the results in the plot.

Each additional mode in the fitting model increases the number of fitting parameters by two (the amplitude and the phase of each mode). This makes it difficult to locate the global minima in nonlinear least-squares fitting when the waveform is contaminated by unaccounted for higher overtones (i.e., close to the peak) and when the overtones have decayed significantly (i.e., at late times). At late times, the higher-overtone fits only fail once the mismatch between the fitting model and the actual waveform reaches machine precision. Even when our toy waveform contains seven overtones (more than the number of overtones included in the fitting model), we can still recover the $N \leq 3$ overtones at sufficiently late times.

These results indicate that if the lowest $(N \le 3)$ overtones are dominant in full GR waveforms, our fitting procedure should return their correct frequencies. For full GR BBH waveforms, we will find very weak evidence of the third overtone even under the most lenient requirement

⁴Here we are using a particular realization of the waveform $\mathbf{Q}_7(t)$. We could use other realizations by choosing a different set of amplitudes and phases. For example, if we fit the amplitudes and phases of $\Psi_{\text{num}}(t)$ at a time $t_0 - t_{\text{peak}} \neq 0$, the agnostic frequency fit could recover the correct frequency up to N=5 for some cases, and the "amplitude constancy" test explained later in this section could also work slightly better.

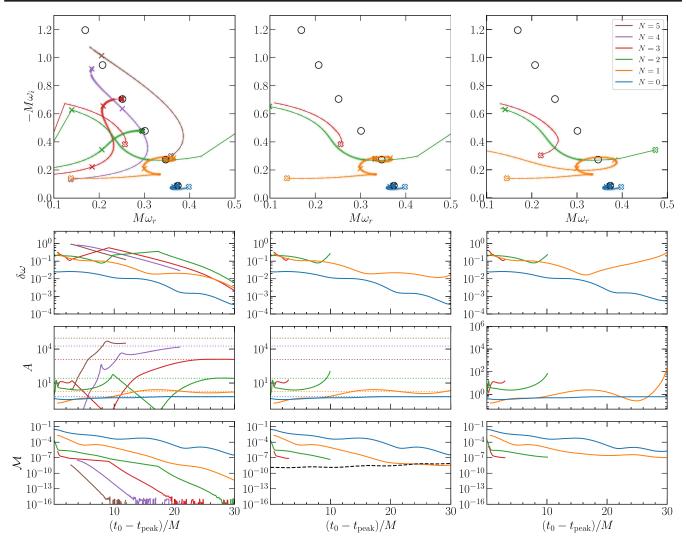


FIG. 2. Frequency-agnostic fits of the $Q_7(t)$ (left column), $Q_7^T(t)$ (middle column), and Ψ_{num} (right column) waveforms with the $Q_{0,1}(t), Q_{1,1}(t), \dots Q_{5,1}(t)$ models, each consisting of 0, 1, ...5 overtones, with the QNM frequency of the highest overtone in the model left as a free (complex) parameter for the fit. Top row: Recovered frequency of the free-frequency mode in each model. Each curve corresponds to a fitting model $Q_{N,1}(t)$ with a different number of overtones. Each point on the curve corresponds to a fit starting at a different t_0 , where t_0 is increased in 0.1M steps. The empty thick cross corresponds to $t_0 - t_{\text{peak}} = 0M$, the two thin crosses correspond to 10M and 20M, and the solid thick cross corresponds to 30M. The empty black circles are the expected overtone frequencies $\omega_0, \omega_1, ...\omega_5$ (from bottom to top) of the $\ell m = 20$ harmonic for a Schwarzschild BH, used to construct the $\mathbf{Q}_7(t)$ and $\mathbf{Q}_7^{\mathsf{T}}(t)$ toy waveforms. Spurious data points (e.g., points that move away from neighboring points by a disproportionate distance, or points that "jump" to a location outside of the plotting domain) are removed, hence some curves start later than 0M, terminate prematurely before 30M, or jump discontinuously at intermediate times, especially for the higher overtones. We also remove the same spurious points in the bottom three rows of this figure. We do not show the $N \ge 4$ curves for $\mathbf{Q}_{\mathbf{T}}^{\mathbf{T}}(t)$ and $\mathbf{\Psi}_{\text{num}}$ (center and right column) because they do not approach any of the expected overtone frequencies meaning fully. Second row: Deviation of the free fitted frequency from the expected overtone frequency, as defined in Eq. (16). Third row: Amplitude of the free-frequency mode. Although the starting time t_0 of the fit is varied, we define the amplitude of the mode to be the amplitude that we would measure by extrapolating back to t_{neak} . In the left and central panels, we plot the injected amplitudes of each mode of the toy waveforms by horizontal dotted lines. Bottom row: Mismatch between the best-fit model and the actual waveform, as defined in Eq. (12). In the central panel, we also plot the mismatch between $\mathbf{Q}_7(t)$ and $\mathbf{Q}_7^{\mathbf{T}}(t)$ (black dashed line), i.e., the mismatch induced by the tail.

(i.e., when we impose the weak "consistency test" that the mode amplitude should be approximately constant for a brief period in time $\sim 2M$). Therefore, including $N \gtrsim 4$ overtones in the fitting model is not useful anyway. We will confirm and reinforce these conclusions below.

2. Damped sinusoids with tail

Next, we test whether adding even the simplest subdominant contamination to the waveform would hinder our ability to recover the QNM frequencies. In the middle panels of Fig. 2 we repeat the agnostic QNM fitting procedure using the same fitting model $Q_{N,N_f}(t)$, but on the toy waveform $\mathbf{Q_7^T}(t)$. We find that the recovery of the fundamental mode and first overtone are almost as good as for the toy waveform $\mathbf{Q_7}(t)$, but the higher overtones $(n \ge 2)$ are not recovered.

Perhaps the most instructive outcome of this experiment is shown in the bottom central panel of Fig. 2. In that panel, the horizontal black dashed line corresponds to the mismatch between $\mathbf{Q}_{7}(t)$ and $\mathbf{Q}_{7}^{\mathrm{T}}(t)$, i.e., the mismatch induced by the tail. It is clear that the individual modes stop converging towards their expected values around the time where their mismatch drops below the mismatch induced by the contamination. At that point, the mode is effectively trying to fit the contamination, and the mismatch saturates. The results are qualitatively similar when we replace the tail with other types of injected subdominant contaminations (e.g., Gaussian noise). Clearly, if we increase the amplitude of the injected contamination, the recovery of the overtones becomes even worse. In summary, this toy waveform illustrates that the presence of an expected subdominant contamination (beyond the "pure ringdown" signal) can prevent a robust extraction of the frequencies of the higher overtones.

What is worse, the free-mode frequency for fitting models with $N \ge 2$ does not converge to any particular value. This is an indication that the fit does not decisively "pick up" any QNM in the waveform. However, as we increase N the minimum mismatch \mathcal{M} keeps decreasing and getting closer to the waveform peak, even when the free mode does not approach the expected overtone frequency. This shows that adding modes to the fitting model can reduce the mismatch even if the mode frequency is unphysical. Therefore a small mismatch is not sufficient evidence to claim the presence of an overtone. Numerical waveforms with high accuracy are necessary to confidently identify overtones in the data, and it is preferable to have a good model of all the sources of non-QNM contamination (including tails, noise, and nonlinearities).

On the other hand, the results for this toy waveform seem to imply that a failure to identify overtones with a frequency-agnostic search is not a proof that the overtones are absent, either. The overtones could be physically present (as they are in this toy model), but the fits could be failing simply because of small, subdominant contaminations. However, even if the overtones were physically present, including many overtones in our ringdown waveform model might not be optimal from a data-analysis point of view, because the model is not robust against even small contaminations. This is an important point if we want to test GR by extracting different frequencies in an agnostic manner; any small contamination would hinder our ability to extract higher-overtone frequencies even if the modes are physically present in the signal.

3. Linear perturbation theory

Finally, in the right panels of Fig. 2 we consider the more realistic case of fitting $\Psi_{\text{num}}(t)$, the time-domain solution of

Eq. (7) with initial conditions given by Eq. (9), extracted at null infinity. The left panel of Fig. 1 shows that $\Psi_{\text{num}}(t)$ (black solid line) results from a combination of QNMs, contamination due to backscattering, and direct propagation of the initial wave packet. Moreover, for reasons explained with a toy model in Appendix A, the initial data contamination means that the "effective" QNM amplitude should typically increase continuously until it approaches a constant. In this more realistic model, extracting the overtone frequencies should be even harder than in the toy models examined previously.

This expectation is validated by our fits. While we can recover the fundamental mode ($\delta\omega\lesssim 10^{-3}$) with a free-frequency fit, none of the overtone frequencies converge to the theoretically expected values at late times. In fact, only the first overtone (orange line) passes relatively close to the "correct" theoretical value ($\delta\omega\lesssim 10^{-1}$) for a short time interval $\sim 1M$ around $t_0-t_{\rm peak}\sim 16M$. This test shows that, in the presence of unmodeled subdominant contributions, we can never recover the correct overtone frequencies by an agnostic fit to test GR.

The original BH spectroscopy test is unfeasible for all overtones, and only possible at late times for the fundamental mode, even within linear perturbation theory.

E. Recovering QNM amplitudes at fixed frequencies

Since the original, agnostic spectroscopy test seems too ambitious even within linear perturbation theory, let us consider a more modest goal. In this "weak" version of the BH spectroscopy test, we assume the frequencies of the QNMs that should be present in the waveform to be known. In other words, we fit our waveforms with the $Q_7(t)$ model (see Table I); the complex frequencies in the model are fixed at the injected values for each QNM, and only the mode amplitudes and phases are free-fitting parameters.

Our toy models will illustrate two main points:

- (i) If the waveform could really be described by a superposition of QNMs with the "right" frequencies, the fitted QNM amplitude should not change sensitively when we choose to fit the waveform at different starting times.
- (ii) The linear waveform $\Psi_{\text{num}}(t)$ gives inconclusive results when we consider $N \gtrsim 4$ (or arguably even fewer) overtones in the model. This is further evidence that we cannot reliably identify overtones with $N \gtrsim 4$, even within linear theory and in a "weak" formulation of BH spectroscopy tests.

1. Pure damped sinusoids

In Fig. 3 we plot QNM amplitudes fitted to various models when we change the starting time of the fit t_0 . As in Fig. 2, for any t_0 we extrapolate the amplitude back to t_{peak} to "unfold" its exponential decay, so the plotted quantity is

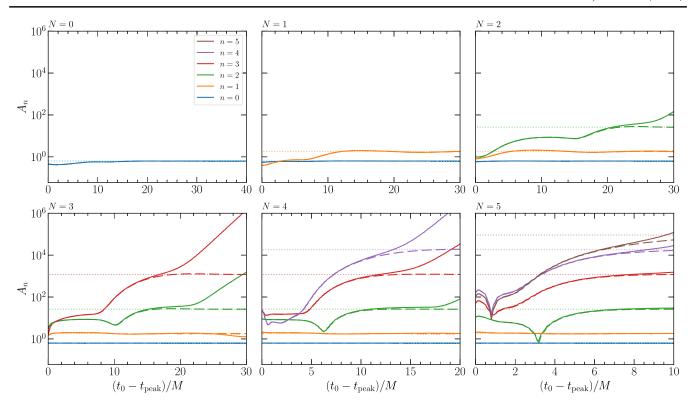


FIG. 3. The amplitudes of QNMs fitted to the $\mathbf{Q}_7(t)$ (long dashed lines) and $\mathbf{Q}_7^{\mathbf{T}}(t)$ (solid lines) toy waveforms, when the frequencies of all modes are fixed to the "true" frequencies of the corresponding overtones. Each panel corresponds to the fit results from a model $Q_N(t)$ with a different number of overtones N. Different colors in each panel correspond to the fitted amplitudes of the different overtones (n=0,1,...,N). While the starting time of the fit t_0 is varied, we always plot the value of the amplitudes extrapolated back to t_{peak} . The injected values of the amplitude in the toy models $\mathbf{Q}_7(t)$ and $\mathbf{Q}_7^{\mathbf{T}}(t)$ are shown as faint, horizontal dotted lines. Note the different time ranges in different panels.

really the amplitude at t_{peak} , as defined in Eq. (3). In other words: if the plotted curve is a flat horizontal line, the amplitudes extracted at different values of t_0 are consistent with each other.

For the "pure QNM" toy waveform $\mathbf{Q}_7(t)$ (long dashed lines in Fig. 3), we find that if we start the fit late enough, the recovered amplitudes "flatten out" to their injected values (shown as faint dotted lines), even when our fitting model contains fewer QNMs than the "true" pure ringdown waveform (i.e., when $N \le 7$).

The different panels show a clear trend: as we add more overtones, the lower overtones converge to a flat line earlier and earlier. Note that all of the overtone amplitudes (not only a subset) converge towards a constant at late times. This behavior is evidence that the model is a complete representation of the true waveform, and that the fit has good convergence properties.

2. Damped sinusoids with tail

Next, we apply the same procedure to the toy waveform $\mathbf{Q}_{7}^{\mathbf{T}}(t)$. The results are shown as solid lines in Fig. 3. The fitted amplitudes are practically the same found in the pure damped sinusoid case for early starting times.

However, there is one remarkable difference: the amplitudes fitted to waveforms contaminated by power-law tails blow up exponentially at some critical time. When the injected contamination increases in amplitude, this "critical blow up" occurs at earlier times. In fact, the exponential blow up occurs when the highest (fastest-decaying) overtone in the fitting model starts to pick up the contamination, and tries to fit the power-law tail with an exponential even at late starting times t_0 , instead of following the expected exponential decay. This results in an exponential blow-up in the figure, since we always plot the amplitude as defined at t_{neak} . We will observe an analogous behavior when fitting NR waveforms in the full, nonlinear theory. Nonetheless, when the fitting model contains many overtones (e.g., N=5), the lower overtone amplitudes (e.g., those with $n \le 3$) are still roughly constant at late times. From Fig. 3, we also conclude that a nonconstant amplitude at early times cannot be attributed to backscattering or some other subdominant contamination, because the early variation of the amplitude is similar with or without the contamination.

3. Linear perturbation theory

In Fig. 4 we consider $\Psi_{\text{num}}(t)$, the waveform computed within linear perturbation theory, and we fix the

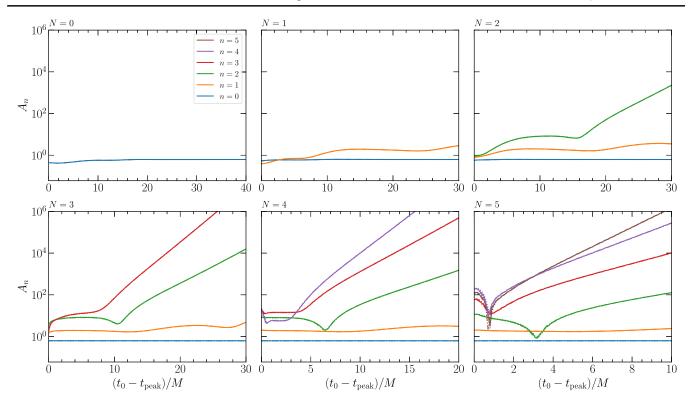


FIG. 4. Same as Fig. 3, but for the linear waveform $\Psi_{\text{num}}(t)$. The dotted and dash-dot lines show the results for the simulations with grid spacing $dt = dr = 1.25 \times 10^{-2} M$ and $6.25 \times 10^{-3} M$, while solid lines show the results obtained from Richardson extrapolation. The three curves are similar, especially at late times, indicating that the variation of the amplitude and the exponential blow up at later times is not primarily driven by numerical error. When compared to Fig. 3, the duration in which the amplitudes look flat is significantly shorter for the $N \le 4$ fit models. For N = 5, only the n = 0, 1 overtones show signs of constancy. Note the different time ranges in different panels.

frequencies to the standard QNM frequencies of a Schwarzschild BH. While the amplitudes are flat within certain time ranges, their consistency is not as apparent as in the simpler toy models we considered earlier.

Consider e.g. the model with N = 4. For the toy waveform $\mathbf{Q}_{7}^{\mathrm{T}}(t)$ in Fig. 3, the amplitude of the n=2 overtone is consistent across a range of $\sim 10M$ at $t_0 - t_{\text{peak}} \gtrsim 7M$. However, the corresponding amplitude when fitting $\Psi_{\text{num}}(t)$ is only constant at times $t_0 - t_{\text{peak}} < 5M$, over a time range shorter than half the period of the fundamental mode $(P_{n=0}/2 \sim 8.4M)$. When we go to N=5, none of the amplitudes is ever constant, except for the n = 0 and n = 1modes. This is different from the previous toy models, where all overtone amplitudes with $n \le 3$ are stable at late times. Also, for the previous toy models, the amplitudes stabilize earlier for fits with higher N, and across different N the individual amplitudes all stabilize to the same value (faint dotted lines). None of these behaviors is observed for $\Psi_{\text{num}}(t)$; compare e.g., the N=4 and N=5 panels in Fig. 4. This implies that for $\Psi_{\text{num}}(t)$, the model breaks down at N = 5, and we should certainly not go further than that. A cautious reader may even argue that we can only identify overtones with $n \le 2$, because for n = 3 and higher the amplitude is flat over times smaller than $P_{n=0}/2.5$

In fact, the $\mathbf{Q}_7^{\mathrm{T}}(t)$ template considered above has been engineered to mimic $\Psi_{\mathrm{num}}(t)$, so the flatness of the amplitude between these two waveforms should have been similar. Differences could arise either because the back-scattering effect has been poorly modeled by extrapolating the power-law tail to early times in $\mathbf{Q}_7^{\mathrm{T}}(t)$, or perhaps because there are additional contamination in the linear theory waveform. For example, while the QNMs correspond to poles in the Green's function for the Regge-Wheeler potential, the "prompt response" (i.e., the direct propagation of the initial wave packet towards spatial infinity; see Appendix A), is expected to contaminate the waveform. Even without such contamination, the QNM amplitudes are varying close to the peak of the

⁵Moreover, a constant amplitude is necessary but not sufficient to infer that QNMs are present in a waveform. If in our fitting model we were to assume the presence of a QNM whose real frequency is very similar to a QNM that is actually present in the waveform, but whose decay time is significantly different, the fitted amplitudes could also turn out to be approximately flat for a brief period close to the peak. This is why agnostic frequency fits represent the most robust method to identify an overtone.

waveform; they build up continuously according to the shape of the initial data, so constant-amplitude QNMs should not be used to model the waveform too close to the peak (see Appendix A and [47]). Pseudospectral instabilities may also cause alterations to the ringdown waveform, especially for higher overtones [138–141] (perturbations of the fundamental mode have been shown to affect the timedomain ringdown waveform only minimally [141], but in principle the instability of the overtones could be observable [139]).

We have estimated the noise in the numerical solution of the Regge-Wheeler equation $\Psi_{\text{num}}(t)$, and we have found it to be subdominant when compared to the extrapolated tail used for constructing $\mathbf{Q}_{7}^{\mathbf{T}}(t)$. Consistently with this finding, if we use a lower-resolution calculation of Ψ_{num} , the fitting results in Fig. 4 do not change significantly. In other words, the exponential blow-up is likely driven by physical effects (power-law tails) rather than by numerical noise. This may not be true for comparable-mass BBH mergers. As far as we know, a power-law tail has not yet been confidently identified in BBH simulations in full GR.

F. Take-home messages

The lessons learned from fitting the three toy waveforms (Fig. 2) can be summarized as follows:

- (a) If the waveform consisted of a pure superposition of damped QNMs, as in model $\mathbf{Q}_7(t)$, we should be able to recover the QNM frequencies by fitting the waveform up to at least n=3, and the amplitude of each QNM would be consistent across different starting times of the fit. This applies to both the case where the frequencies are free, and to the case where they are fixed to their "exact" values.
- (b) Contaminations that are subdominant close to the waveform peak, such as the power-law tails injected in $\mathbf{Q}_7^{\mathbf{T}}(t)$, limit our ability to agnostically recover the overtone frequencies as required for testing GR, even if the overtones exist in the waveform and are dominant. However, if we assume that overtones are present and we fix their frequencies in the fitting model, while the amplitudes will blow up exponentially at a later time, they should be approximately constant at intermediate times for overtones up to n=3. The constancy of the lower overtones should improve when we add more overtones to the model.
- (c) An agnostic damped-sinusoid fit cannot recover the correct frequencies for any of the overtones when we consider "true" ringdown waveforms computed within linearized gravity $\Psi_{\text{num}}(t)$. A fit of the linearized waveform with frequencies fixed to their known values does not show convincing evidence that overtones with $n \gtrsim 2$ are present in the signal. In fact, the identification of the overtones is significantly more problematic than in the toy waveforms $\mathbf{Q}_{7}(t)$ and $\mathbf{Q}_{7}^{T}(t)$.

The conclusions of this exercise are quite clear.

First and foremost, using a small-mismatch criterion is not sufficient to conclude that overtones are present in the waveform. Overtones are prone to overfit the early part of the waveform, because the rapidly decaying higher overtones, combined together, are just fine-tuned "bumps" that can fit away other sources of contamination.

Even if the ringdown following a BBH merger were precisely described by linear theory, this would not imply that a superposition of multiple QNMs is sufficient for waveform modeling. Even the linearized waveforms are plagued by physical contamination from the prompt response and tails, and this makes it hard to conclude whether higher overtones are present, even if we assume the overtone frequencies to be known.

Moreover, the Gaussian scattering example implies that it is difficult to use more than n=2 overtones to test GR. In fact, Fig. 2 shows that—even in linear perturbation theory and for nonrotating BHs—recovering the theoretically predicted QNM frequencies is difficult even for the first overtone, unless we fine-tune the starting time t_0 of the fit. A "blind" (percent-level) precision measurement of QNM frequencies is only feasible for the fundamental mode.

If we insist to use $0 < n \lesssim 2$ overtones to test GR, we should start fitting the waveform at times significantly after the peak (e.g., $\gtrsim 10M$ after the peak for a model with one overtone). This is because the overtone amplitudes are roughly constant only at late times, whether or not the waveform is linear starting from the peak.⁶

While it is true that the fitting model we use for overtone extraction is incomplete (because a tail is clearly present in the numerical linear waveform), the key point is that similar unmodeled linear and nonlinear contributions will certainly be present in full GR. Hence we can expect overtone recovery to be affected by similar issues in more realistic cases, in the absence of extremely accurate analytical models. (Incidentally, we have also tried to fit Ψ_{num} with a power-law tail in addition to QNMs, but we could not confidently identify the portion of the waveform where the tail starts being dominant.) Additional contributionsincluding nonlinear effects [85,142] and nonlinear QNMs [143-147]—have indeed been found in BBH merger waveforms in full GR. As a linear superposition of QNMs cannot fit a linear waveform in a self-consistent manner, we would expect the model to perform even worse when fitting the postmerger waveform of two comparablemass BHs. The next sections will confirm these expectations. In our analysis of the ringdown of BBH mergers

⁶The time at which we should start the fit depends on the initial conditions used when solving for the linear waveform, and on the error tolerance we are willing to accept when we fit the frequency and amplitude. Note also that the Regge-Wheeler waveform is related to \ddot{h} , the second derivative of the GW strain, so the time delay needed for fitting h in BBH merger waveforms might be significantly different. We will return to this topic below.

simulated in full GR, the $n \ge 2$ overtones cannot be confidently identified, and the first overtone can only be identified at times $\approx 10M$ after the waveform peak.

In Appendix D we repeat some of the present analysis on a complex-valued toy waveform constructed to mimic the NR postmerger waveform. The conclusions are qualitatively similar, if not stronger. When we consider a complex toy model consisting of seven overtones, the frequency-agnostic fits work better than those presented here, and the fitted amplitudes are even more stable. In other words, a failure of these test for a complex NR waveform is even stronger indication that the waveform cannot be modeled by a superposition of overtones.

IV. ARE POSTPEAK BBH WAVEFORMS LINEAR?

Let us now turn to the real problem of interest—fitting waveforms in full GR. Fitting overtones in ringdown signals is a notoriously difficult problem even in the absence of instrumental noise (see, e.g., [53,75,81,85,148,149]). Besides the physical effects discussed so far, time variations in the inferred mode amplitudes in full GR can occur because the mass and spin of the remnant extracted from numerical simulations vary significantly close to the peak of the radiation [53,54,72,119,142]. Let us illustrate this point more concretely.

To facilitate comparison with previous work, we will follow Ref. [106] and focus on the GW150914-like NR waveform SXS:BBH:0305 in the Simulating eXtreme Spacetimes (SXS) catalog [150,151]. The waveform represents a BH binary with mass ratio of 1.22, primary dimensionless spin $\chi_1 = 0.33$ aligned with the orbital angular momentum, and secondary dimensionless spin $\chi_2 = -0.44$ antialigned with the orbital angular momentum. The merger remnant in this simulation has final mass $M_f = 0.9520M$ and dimensionless spin $\chi_f = 0.6921$.

Reference [106] suggested that the addition of several overtones is necessary to appropriately model ringdown and to infer the final mass and spin. When starting to fit at $t \ge t_{\text{peak}}$, where t_{peak} is defined as the time where the $(\ell, m) = (2, 2)$ component of the strain has a maximum, overtones up to n = 7 were included to obtain an unbiased estimate. Earlier work had indeed found that using a single mode can lead to large systematic errors on the inferred mass and spin of the remnant [53,71,72].

As pointed out in early systematic studies of ringdown from nonspinning BH merger simulations [53,54], the fact that a linear superposition of damped exponentials can reproduce the merger waveform for $t \ge t_{\text{peak}}$ does not necessarily imply that the time evolution of the background and nonlinearities can be ignored. A significant fraction of the mass and angular momentum is being radiated away from the system postmerger, while the Kerr QNM frequencies are computed assuming a fixed background.

In Fig. 5 we show the difference between the BH mass and angular momentum and their asymptotic values,

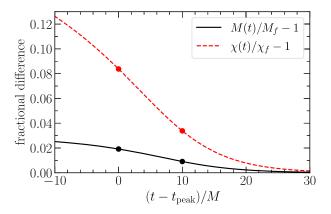


FIG. 5. Evolution of the remnant mass and dimensionless spin for the GW150914-compatible simulation SXS:BBH:0305. The background is rapidly evolving at least until $t-t_{\rm peak} \simeq 20M$, where $t_{\rm peak}$ is defined as the time at which the strain h has a maximum.

computed for the SXS:BBH:0305 simulation following the procedure outlined in Ref. [152]. We find that the remnant mass and dimensionless spin differ from their asymptotic value by 2% (1%) and 8% (4%) at $t=t_{\rm peak}$ ($t=t_{\rm peak}+10M$), respectively. Such large variability in the background spacetime can significantly complicate the analysis, and there is no reason *a priori* why the simple model of a linearly perturbed BH with a fixed mass and spin should work around the peak. In fact, several authors pointed out that modeling waveforms close to the peak of the radiation by a linearly perturbed BH with a fixed mass and spin leads to conceptual issues [82,108,142,153,154].

As noted in Ref. [53], the large amount of radiation in a BBH merger "raises the question as to whether or not the radiated energy and angular momentum are affecting the QNM fits. This issue will, of course, become more significant as the fits are pushed to earlier times."

In this section we ask two questions: (1) is it really legitimate to describe the whole postpeak waveform as a linear perturbation of the final, stationary Kerr BH? and (2) how many overtones can be reliably used to obtain unbiased estimates of the remnant's mass and spin?

If we could indeed ignore the time-evolving background and describe the whole postpeak waveform as a superposition of QNMs from a fixed Kerr background then the overtone amplitudes should be constant, by definition. We have seen that this expectation is questionable even in linear theory. In Sec. IV A we confirm, perhaps at this point unsurprisingly, that a constant-amplitude overtone superposition does not work in the BH merger case either; the amplitudes of the overtones change significantly when we change the fitting window.

The nonconstancy of the amplitudes is not the only issue with a linear perturbation theory interpretation of postpeak ringdown. Reference [106] claims that (i) the inclusion of the fundamental mode and seven overtones provides a very

accurate description of the ringdown up to the peak strain amplitude, and (ii) including seven overtones significantly reduces the uncertainty in the extracted remnant properties.

In Sec. IV B we show that the higher overtones lead to very small mismatches by merely overfitting the waveforms. Furthermore, we argue that these higher overtones try to fit other physics (such as prompt response effects, time variations of the background, and nonlinearities) close to the merger. The addition of several overtones allows for better extraction of the fundamental mode and the first overtone, which are mainly responsible for better estimating remnant properties. This improved extraction arises because high overtones contributions act as effective terms, "fitting away" poorly understood physics.

The fact that higher overtones lead to overfitting is confirmed in Sec. IV C, where we show that these modes do not play a significant role in extracting the remnant's properties, either. To show this, we swap a particular mode with a random damped exponential. If the "fake" mode still fits the waveform with similar or better accuracy, or if it still extracts the remnant properties accurately, we can conclude that the originally "swapped" overtone was not really necessary. We use this argument to conclude that overtones with $n \ge 2$ do not significantly contribute to the extraction of the remnant's parameters, and therefore there is no motivation to include them in the modeling.

A. Are overtone amplitudes consistent with the linear model?

We start by discussing the inconsistency of the amplitudes obtained by fitting a simple overtone model with fixed frequencies, the Q_7 model of Eq. (3), to numerical BH merger simulations. This inconsistency confirms, from a different perspective, that the asymptotic values of the frequencies should only be used at late times.

In linear BH perturbation theory, the ringdown part of the strain h is given (by definition) by the Q_N model of Eq. (10). In this section we focus on the dominant, $\ell = m = 2$ component of the radiation. To avoid cluttering, we simplify the notation by setting $A_n = A_{22n}$, $\phi_n = \phi_{22n}$, and $\omega_n = \omega_{22n}$.

Each QNM depends on four real parameters; a real amplitude A_n , a real phase ϕ_n , and a complex QNM frequency $\omega_n = \omega_r^n + i\omega_i^n$. The complex frequency is determined by the remnant Kerr BH's mass and spin, while the nature of the perturbations (i.e., the parameters of the binary progenitor) dictates the amplitudes and phases. In linear perturbation theory, all of these quantities are time-independent constants [39,42–46].

In the Q_7 model, the entire postpeak waveform is fitted by this simple model at the cost of including a large number of overtones (N=7). Note that in this model, the complex frequencies are forced to their asymptotic values; only the amplitudes and phases are free parameters, extracted by fitting the waveform between t_0 and $t_f=t_{\rm peak}+90M$. We want to explore the behavior of amplitudes and phases in the Q_7 model we vary t_0 . Time variations of the amplitudes or phases with t_0 imply departures from the bare-bones ringdown model in linear BH perturbation theory, which can happen for various reasons—e.g., because the mode frequency has not relaxed to its true value, because the QNM included in the fit is not present in the waveform, or because damped exponentials are actually overfitting some other feature (such as power-law tails, the prompt response, nonlinearities, or numerical noise). We will mostly focus on the constancy of the amplitudes, but similar arguments can be made for the phases.

In Fig. 6, we plot the amplitudes of different modes as a function of the starting time of the fit. Each panel uses a Q_N fitting model with a different number of overtones, from N=0 (top left) to N=7 (bottom right). As usual, we extrapolate the amplitudes to their values at $t_0=t_{\rm peak}$. Let us define (somewhat arbitrarily) the amplitude of a mode to be "constant" if it varies by less than 10%. Vertical colored bands in each panel show the time range in which the amplitude of the highest overtone (n=N) is constant according to this criterion.

Consider first the simplest case in which we fit the waveform using only the fundamental mode (n = N = 0, top left). The amplitude of the mode is not constant at early times, and it only saturates to a constant value when the starting time of the fit $t_0 \gtrsim t_{\rm peak} + 12M$. Adding the first overtone (N = 1) has the effect of further stabilizing the amplitude A_0 of the fundamental mode, which now is constant as soon as $t_0 \gtrsim t_{\text{peak}} + 1M$. This doesn't mean that the two-mode model is a good description close to the peak, because the amplitude of the first overtone is now rapidly changing in time close to the peak: if we require both mode amplitudes to be constant, we are again forced to consider the late-time portion of the waveform, where $t_0 \gtrsim t_{\rm peak}$ + 13M for N = 1. This confirms previous findings in Ref. [108] (see the top panel in their Fig. 3). Note also that the two-QNM model seems to fail at very late times $(t_0 \gtrsim t_{\text{peak}} + 28M)$, where A_1 blows up very rapidly. We interpret this behavior in light of the findings of Sec. III. In this particular case, the blow-up is due to the free overtone amplitude latching onto a different QNM; as we will show in Sec. VA, the late-time fit is dominated by the (3, 2, 0) QNM due to spherical-spheroidal mode mixing, and not by the (2, 2, 1) QNM. The exponential blow-up seen here is quite generic. Since we show the mode amplitude rescaled to its peak value (thus incorporating an exponential factor) any term with a time-dependence slower than the expected exponential (e.g., a different QNM, a power-law tail, or numerical artifacts) would give rise to this behavior.

We should also try to systematically quantify fitting errors on the overtone amplitudes. Without error estimates, one may argue that amplitude variations could be ascribed to the accuracy of the numerical simulations, rather than being due to physical effects (such as the prompt response,

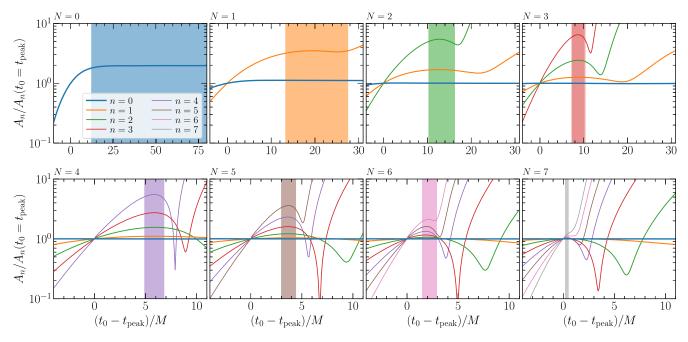


FIG. 6. Amplitude $A_n^N(t_0)/A_n^N(t_0=t_{\rm peak})$ of QNMs as a function of the starting time t_0 for the SXS:BBH:0305 simulation. The shaded regions show the largest time range such that the amplitude of the highest overtone (n=N) is constant within 10%.

tails, or a time-evolving background). In Appendixes B and C we exclude this possibility. We show that the finite resolution of the simulation, the extrapolation procedure used to extract the radiation, and the presence of a spurious constant in the signal due to a suboptimal frame choice are subdominant effects that do not significantly affect the inferred amplitudes. In addition, we will see below that the qualitative behavior of the mode amplitudes is unchanged when we include all known ringdown contributions (such as those coming from spherical-spheroidal mode mixing). Therefore we conclude that the blow-up is triggered by unmodeled postmerger components, including possible residual nonlinearities from the merger phase.

As we keep adding more and more overtones, we observe some important trends:

- (i) The fundamental mode amplitude A_0 stabilizes to constant values at earlier and earlier times; for example, in the Q_7 model (N=7) A_0 is already constant at $t_0 \approx t_{\rm peak} 15M$. This was also noted in Ref. [106]. The fundamental mode having constant amplitude before the peak does not necessarily mean that it is excited at times before the peak. In fact, there is a more economic explanation: since the fundamental mode is the longest-lived, the amplitude inferred by fitting is largely fixed by the late-time behavior, where the mode is actually present and the amplitude constant.
- (ii) For $N \ge 2$, the highest-overtones amplitudes are "constant" only close to a local maximum, and never saturate to a constant value at late times. This suggests that these modes are not actually contributing to the

- signal close to the waveform peak, but they are rather overfitting other features of the waveform beyond the simple linear ringdown model (possibly, nonlinear contributions).
- (iii) For N = 1, A_1 is constant (according to our definition) at times $13 \lesssim (t_0 - t_{\text{peak}})/M \lesssim 28$. For N = 2, A_1 is constant in the region $7 \lesssim (t_0 - t_{\text{peak}})/M \lesssim 25$. As N increases, two things happen. First of all, the constant- A_1 time band moves to earlier times. Secondly, the region where the highest overtone has approximately constant amplitude keeps shrinking, so that higherovertone amplitudes are roughly constant in smaller and smaller regions. This makes sense, because higher overtones are very short-lived. Consider for example the Q_7 model (N=7): the width of the time range over which A_2 , A_3 , A_4 , A_5 , A_6 , A_7 are constant is 5.1, 2.8, 1.8, 1.6, 0.9, 0.3, respectively. Once the amplitude and phase of the fundamental mode are fixed by the late-time behavior of the waveform, the amplitudes and phases of the higher overtones can be adjusted to fit the rest of the waveform close to the merger.

In our opinion the mode-amplitude evolution shown in Fig. 6 is among the main results of this paper. It has two important implications:

- (i) Overtone models with $N \ge 2$ are unphysical, because they try to overfit other features of the waveform.
- (ii) Models with at most one overtone $(N \le 1)$ are physical, but they can only be used for meaningful spectroscopy tests at late times.

Below, we show that the constancy of the amplitude is not the only issue. Higher-overtone models do not

necessarily yield the smallest mismatch with numerical relativity waveforms (Sec. IV B), nor the best estimate of the remnant's mass and spin (Sec. IV C).

Finally, we remark that most investigations so far have been devoted to the "detection problem"—if overtones are present, can we extract them from the waveform? In Appendix E we look at the problem of "false positives," i.e., can a linear-overtone model find spurious evidence for the overtones even when they are not present? Somewhat concerningly, we find that the answer is affirmative.

The consistency checks performed here and below are necessary, because ill-posed models can lead us to extracting wrong physics from the numerical simulations.

B. Which overtones are necessary to correctly fit the waveform?

In the absence of an analytical description of the postmerger phase it is important to avoid overfitting issues when we try to assess the physical contribution of different overtones. For example, we can check that the contribution of overtones cannot be replaced, or even be improved upon, by possibly unphysical low-frequency contributions which are effective at fitting the data. Indeed, here we show that the mismatch achieved by a superposition of overtones as predicted in GR can be matched, or even surpassed, by adding "unphysical" low-frequency damped exponentials with frequency and damping time that do not appear in the spectrum predicted in GR for the final BH [108,111]. It should be noted that similar "pseudo-ONMs" were introduced in the context of the effective-one-body framework precisely to model the rapid transition from the inspiral GW frequency to the postmerger QNM frequency "plateau" [124-126].

To understand the contribution of each overtone, we fit the waveform using a modified version of the Q_7 model,

$$Q_7^{n_{\rm rnd}} = A_{n_{\rm rnd}} e^{-i[\omega_{n_{\rm rnd}}(t - t_{\rm peak}) + \phi_{n_{\rm rnd}}]} + \sum_{n=0 \atop n \neq n_{\rm rnd}}^{7} A_n e^{-i[\omega_n(t - t_{\rm peak}) + \phi_n]}.$$
(17)

This is similar to the Q_7 model, except that the complex frequency of the $n=n_{\rm md}$ overtone is set to random (and, in general, incorrect) values uniformly sampled in the range $[0,2{\rm Re}(\omega_{22n})]+i[0,2{\rm Im}(\omega_{22n})]$. The aim of this exercise is to estimate which overtones are overfitting the waveform, if any. If a random damped exponentials can fit the waveform better than the Q_7 model, we can conclude that the overtone being randomized is not "special," and it is merely overfitting the waveform.

In each panel of Fig. 7 we randomly draw 10³ complex frequencies for the randomized overtone, and we show the 90% confidence level of the corresponding mismatches using dark-shaded bands. The thick black line shows the

mismatches in the (unrandomized) Q_7 model. We also show the 90% confidence level of mismatches computed by randomizing all overtones with $n > n_{\rm rnd}$ —as opposed to Eq. (17), where only $n = n_{\text{rnd}}$ is randomized—using lighter-shaded bands. The top-left panel shows that the fundamental mode plays a key role in fitting the ringdown; the fit becomes very poor if we substitute it with a damped exponential with random frequency. The first overtone is also crucial in fitting the early part of the waveform, confirming the findings of Ref. [72]. However, as we randomize the frequencies of the higher overtones, we observe that the gap between the colored bands and the black line (the Q_7 model) decreases, and the two start overlapping. The inset below each panel shows the fraction of random samples that outperform the Q_7 model. For example, at $t_0 = t_{\text{peak}} + 6M$, more than 90% of randomly drawn damped sinusoids that replace n = 2 yield smaller mismatches than the Q_7 model. If all overtones with $n \ge$ $n_{\rm rnd}=2$ were replaced by random frequencies, 2–20% of the samples would still yield better fits than the Q_7 model for $t_0 - t_{\text{peak}} = 0-5M$. The performance of higher overtones is even worse. If n = 3 is replaced, as many as 75% of random samples fit better than the Q_7 model at $t_0 \approx t_{\rm peak} + 3M$. Similarly, for $n_{\rm rnd} = 4-7$, 15% to 66% of the random samples perform better than the Q_7 model for $t_0 - t_{\text{peak}} = 0-5M$.

In conclusion, random frequencies often yield better results than higher overtones for $N \ge 2$. This is further evidence that higher overtones play the role of unphysical low-frequency components, which are nonetheless effective at fitting the early-time signal. Recall that the real part of the frequency of the nth overtone, ω_r^n , decreases with n and that, at the end of the inspiral, the $\ell=m=2$ component of the signal has frequency $\omega_{\rm peak} \simeq 0.65\omega_{220}$, increasing monotonically towards ω_{220} postmerger. This is why overtones naturally latch onto lower-frequency components of the signal (see Sec. V B for a discussion).

C. Which overtones are necessary to correctly extract the remnant's properties?

It has been claimed that the inclusion of higher overtones not only fits the waveform better, but also yields a more accurate estimation of the remnant (dimensionless) spin χ_f and mass M_f . What is the reason for this improvement in mass and spin estimation? Do higher overtones really contribute to the mass and spin measurement, as often claimed in the literature? One less explored alternative is the possibility that higher overtones are instead just "fitting out" nonlinear effects, thus allowing for more accurate extraction of the fundamental mode and of the first overtone. This would also lead to a better estimation of the remnant's mass and spin.

To address this question, we perform a gedanken experiment similar to the one in Sec. IV B. We fit the waveform

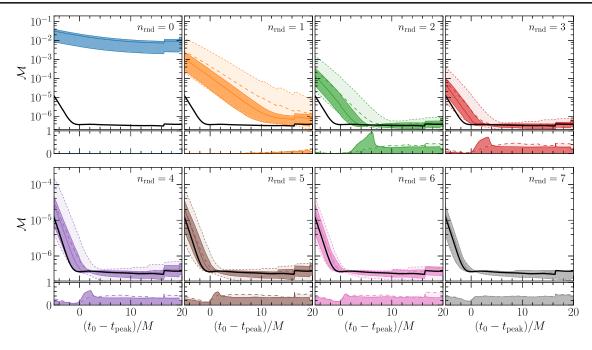


FIG. 7. Median and 90% confidence interval of the mismatch \mathcal{M} between ringdown models and the SXS:BBH:0305 waveform as a function of the starting time of the fit t_0 . Solid lines and darker shaded areas are found by randomly varying the frequency of a single overtone with $n = n_{\rm rnd}$ from a uniform distribution, as described in the main text. Dashed lines and lighter shaded areas correspond to random variations of all overtones with $n \ge n_{\rm rnd}$. The solid black line shows the mismatch for the Q_7 model. In the insets below each panel we show the fraction of random samples that outperform the N = 7 model.

assuming that the complex frequencies of each overtone depend on the remnant's mass and spin as predicted by linear perturbation theory within GR. Once again, we select one specific overtone $n_{\rm md}$ and replace it by (incorrect) complex frequencies drawn randomly in the range $[0, 2{\rm Re}(\omega_{22n})] + i[0, 2{\rm Im}(\omega_{22n})],$

$$Q_{7}^{n_{\text{rnd}}}(\chi_{f}, M_{f}) = A_{n_{\text{rnd}}} e^{-i[\omega_{n_{\text{rnd}}}(t - t_{\text{peak}}) + \phi_{n_{\text{rnd}}}]} + \sum_{n=0 \atop n \neq n}^{7} A_{n} e^{-i[\omega_{n}(\chi_{f}, M_{f})(t - t_{\text{peak}}) + \phi_{n}]}.$$
(18)

In this model the mode with index $n = n_{\text{rnd}}$ is not used in the estimation of χ_f and M_f . To quantify the accuracy in extracting the remnant's spin and mass, following Ref. [106], we compute the quantity

$$\epsilon = \sqrt{(\delta \chi_f)^2 + \left(\frac{\delta M_f}{M}\right)^2},\tag{19}$$

where $\delta \chi_f$ and δM_f are the differences between the inferred value of these quantities and their asymptotic, "exact" Kerr values.

In Fig. 8 we show the median and 90% confidence interval of ϵ as a function of the starting time of the fit t_0 . The solid black line uses all N=7 overtones to estimate the remnant's properties. In each panel, solid lines and darker colors refer to randomly varying the $n=n_{\rm md}$

overtone frequency from a uniform distribution and ignoring the corresponding overtone in the estimation of χ_f and M_f . Dashed lines and lighter colors were computed by randomizing all overtones with $n \ge n_{\rm rnd}$.

The conclusions of this exercise are completely consistent with the ones drawn from Fig. 7. In fact, the top-left panel shows that the fundamental mode carries most of the information. The second top panel from the left confirms that the first overtone (n = 1) is crucial to get a good estimate of χ_f and M_f close to the peak, again confirming the findings of Ref. [72]. At later times the first overtone does not matter as much, and we observe more overlap between the orange bands and the black line. The role of higher overtones in parameter estimation is much more marginal. For example, at $t_0 = t_{\text{peak}} + 2M$, 15% of the random complex frequencies that replace n = 2 yield better results than the second overtone. As we randomize higher overtones $(n_{\rm rnd} = 3-7)$, the number of samples that do better than the Q_7 model ranges between 55% and 70%. Note also that the value of ϵ from Q_7 model has a maximum when $t_0 - t_{\text{peak}} \approx 7 - 8M$; at those times, the Q_7 model is worse than most random samples. If all overtones $n \ge$ $n_{\rm rnd} = 2$ are randomized, nearly 35% of the random samples do better than the Q_7 model at $t_0 = t_{\text{peak}} + 5M$; this fraction rises to 66% at $t_0 = t_{\text{peak}} + 8.4M$. The fact that random frequencies often do better than high-overtone models is further evidence that most of those overtones are unlikely to be physical.

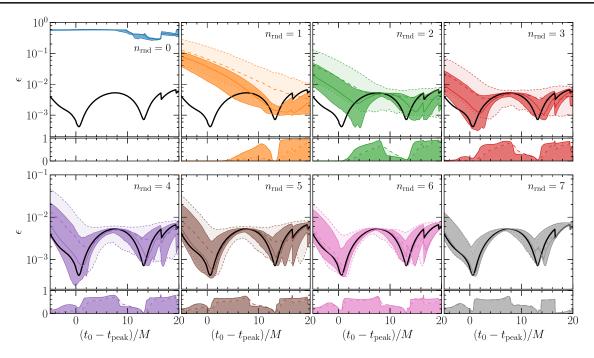


FIG. 8. Median and 90% confidence interval of the combined error on the remnant mass and spin ϵ [see Eq. (19)] as a function of the starting time of the fit t_0 . Solid lines and darker shaded areas are found by randomly varying the frequency of a single overtone with $n = n_{\rm rnd}$ from a uniform distribution, and ignoring the corresponding overtone in the estimation of χ_f and M_f . Dashed lines and lighter shaded areas correspond to doing the same for all overtones with $n \ge n_{\rm rnd}$. The solid black line shows the value of ϵ found when all N = 7 overtones are used for the estimation of the remnant's properties, as in the Q_7 model. In the insets below each panel we show the fraction of random samples that outperform the Q_7 model.

The results above can be understood considering that each overtone introduces two additional fitting parameters, and that ω_r^n decreases with n. These properties allow to easily "fit away" the earlier nonlinear and low-frequency part of the postmerger signal, making the fundamental mode and the first overtone more easily resolvable. This, in turn, improves the estimation of the remnant mass and spin, which are still constrained to the correct Kerr value by the late-time behavior of the signal. It is not a surprise that models with several parameters are prone to overfitting (in fact, this concern was also raised by Ref. [106]). However, Ref. [106] concluded that all overtones are physical by varying all of them while keeping the fundamental mode fixed. This approach does not estimate the significance of each individual overtone. Our analysis shows that as we increase the overtone number, the relevance of each overtone in estimating the remnant's mass and spin decreases. Another difference between our results and those of Ref. [106] is that they varied the complex frequencies by less than 20% in one dimension; they set $\omega_{\rm rnd} = \omega(1+\delta)$ with $\delta \in \mathbb{R}$. Our experiment is more conservative and general, because the real and imaginary parts of the complex QNM frequencies are sampled independently and allowed to vary by 100%.

To further our understanding of which overtones are physical, we perform one more test. We first compute the mismatch at $t_0 = t_{\rm peak}$ for the Q_7 model. Then we vary the frequency of a single QNM by assuming that it depends on (M_f,χ_f) as predicted in GR and changing (M_f,χ_f) , while we keep all other modes fixed to their "correct" asymptotic values.

In Fig. 9 we show the mismatch at the peak as a function of M_f and χ_f . Each panel was computed by varying the corresponding overtone. For n = 0 the mismatch has a sharp minimum at the true value of the remnant's mass and spin (top left panel), as expected. As n increases, the minimum of the mismatch extends over a much larger region in the (M_f, χ_f) plane. More remarkably, the mismatch for higher overtones has a minimum at values of (M_f, χ_f) that do not coincide with the true remnant quantities. For n = 0, 1 the true (M_f, χ_f) coincide with the minimum in the figure with an accuracy of $\epsilon \lesssim 0.01$. For n = 2, the remnant mass and spin at the minimum mismatch deviate from the true values by $\epsilon \approx 0.05$, and the error is dominated by the inaccuracy in M_f . The minimum mismatches for n > 2 occur for values of (M_f, χ_f) which differ by $\epsilon > 0.1$ from the true values, and these minima tend to cluster at near-extremal values of the spin $(\chi_f \to 1)$. Furthermore, as expected, the variations of the mismatch

⁷As Enrico Fermi told Freeman Dyson in a famous encounter [155]: "I remember my friend Johnny von Neumann used to say, with four parameters I can fit an elephant, and with five I can make him wiggle his trunk."

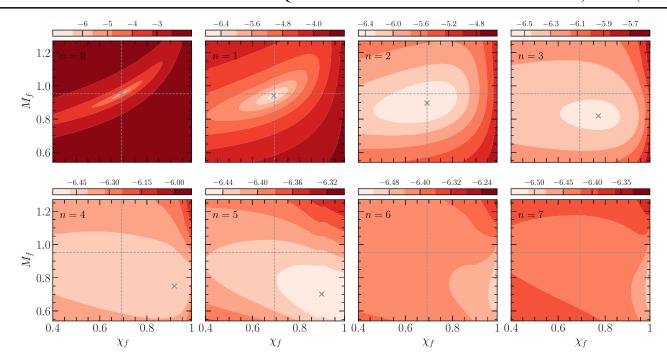


FIG. 9. Logarithm of the mismatch, log \mathcal{M} , evaluated at $t_0 = t_{\text{peak}}$ in the (M_f, χ_f) plane. Starting from the Q_7 model, we vary the mass and spin of overtone n, while all other overtones are set to their "true" asymptotic values. The cross-hair identifies the true value of the remnant's mass and spin. Note that the color maps in each panel correspond to very different ranges. The cross marks the (M_f, χ_f) values minimizing the mismatch.

with (M_f, χ_f) are much milder as we move to higher overtones. This is consistent with Fig. 7, where the 90% confidence levels in the mismatch are very narrow for higher overtones, even though their complex frequencies are allowed to vary by as much as 100%.

Reference [106] used a similar test to claim support for the Q_7 model, because the minimum mismatch for N=7 occurs at the true remnant spin and mass. However, Fig. 9 implies that the minimization of the mismatch is almost entirely due to the modes with n=0 and n=1. When considered individually, the higher overtones play a negligible role. Even worse, they minimize the mismatch at the wrong values of the remnant's parameters.

In this section we have used the Q_7 fitting model to facilitate comparison with Ref. [106]. This model does not take into account spherical-spheroidal mode mixing. As we will see below, this is a substantial effect. By repeating the experiments discussed in this section with mode mixing taken into account, we have verified that the extraction of the first overtone can significantly improve at late times; for example, the amplitude A_{221} would not diverge at late times (as it does in Fig. 6 of Sec. IVA), but remain constant. Similarly, the mismatch \mathcal{M} (Sec. IV B) and the mass/spin measurement error ϵ (Sec. IV C) associated with n=1 would improve at late times. However, including the spherical-spheroidal mixing does not yield visible improvements for overtones with n>1.

Although we have demonstrated the presence of several pathologies in the linear model with N=7 overtones, and

therefore shown that it is unphysical, this does not imply that overtones are absent altogether. Lack of evidence is not evidence of absence—the overtones may still be present in the signal, but they may be hard to extract because additional physics (including prompt response, tails, and nonlinear effects) is present in the waveform. By accounting for these effects, and in particular for postmerger nonlinearities, it may still be possible to extract the overtones. In Appendix D we show that, indeed, the overtones can be correctly extracted from a pure linear model with seven overtones, at least when numerical errors are negligible. Numerical errors increase the minimum mismatch to $O(10^{-9})$, and compromise our ability to extract the overtones. Thus, it may still be possible to successfully extract the overtones with improved NR waveforms and a better understanding of nonlinearities.

V. EXTRACTING COMPLEX FREQUENCIES FROM THE WAVEFORM

One approach to look for the presence of a QNM is to fix the mode's complex frequency (or assume that it depends on mass and spin as predicted in GR), and then check if adding that mode improves the mismatch with the numerical waveform (Sec. IV B) or leads to better accuracy in the remnant parameters (Sec. IV C). As we discussed, this approach can lead to overfitting.

In this section we take a more agnostic route in the spirit of the original BH spectroscopy proposal. To prevent overfitting, we do not impose *a priori* that modes are present in the numerical data. Instead of fixing the complex frequencies of the damped exponentials, we keep them free (see the analogous treatment of the linearized case in Sec. III). This is a much stronger test (it is effectively the original BH spectroscopy proposal at infinite SNR), and therefore it should lead to more robust conclusions about which modes are truly present in the data.

We will show that many of the fitted damped exponentials robustly converge towards known QNM frequencies, naturally selecting the physical modes contributing to the ringdown signal; that is essential to include spherical-spheroidal mode mixing to identify the correct modes; and that the agnostic fits occasionally lead to some surprises. To improve readability, in this section we focus on the dominant (2, 2) multipole. We defer a discussion of higher multipoles to Appendix F.

A. An agnostic fit

We start off by asking the question how many QNM frequencies can we extract without assuming any (no-hair theorem enforced) relation between them?

We fit the complex frequencies and amplitudes using a standard Levenberg-Marquardt algorithm. The large number of fitting parameters implies the possibility of multiple local minima for the parameter values resulting from the fit, and therefore different initial values may result in different solutions. To address this issue, we employ an iterative approach.

- (1) In the initial iteration, we perform the fit using 1000 different random initial parameter values for every value of t_0 . The initial values of the amplitudes are drawn from a log-uniform distribution in the range $[10^{-5}, 10^3]$, while phases are drawn uniformly in $[0, 2\pi]$, $M\omega_r$ is drawn uniformly in $[0, 5\text{Re}(M\omega_{\ell,m,0})]$, and $M\omega_i$ is drawn uniformly in $[0, 5\text{Im}(M\omega_{\ell,m,0})]$. This yields 1000 solutions, out of which we select those that exhibit the smallest mismatch with the NR waveform, for each t_0 .
- (2) In the subsequent iteration, we perform the fitting process again for each value of t_0 , just like in the initial iteration. However, this time, instead of starting with random initial parameter values, we use the initial conditions that correspond to all solutions obtained in the previous iteration. To clarify this step, suppose the previous iteration resulted in unique solutions for 300 different values of t_0 . Now, for the current iteration, we fit the parameters again for each specific t_0 , using initial guesses that equal the solutions obtained in the previous step. In other words, those 300 distinct solutions are employed as initial guesses for every single t_0 . Consequently, we will often encounter different solutions for each of these values of t_0 .

- Once more, we select the solutions that have the smallest mismatches with the NR waveform.
- (3) We repeat this procedure in the same range of t_0 until the mismatch for all points is within 1% of their values from the previous iteration. The number of iterations necessary to achieve convergence is typically less than three for small values of N_f ; when N_f is large (e.g., $N_f > 5$), as many as eight iterations might be required.

In the top row of Fig. 10 we fit the waveform with model $Q_N^{(f)}$ (see Table I), i.e., we extract the complex frequencies of all modes from the fit. We plot the complex frequencies ω_{n_f} (where the subscript "f" stands for "free") extracted from the numerical waveforms for N=0 (left panel), N=1 (middle panel), and N=2 (right panel). In the second row we use model $Q_N^{(f,m)}$ instead: in addition to the N free-frequency modes, we also include one mode with frequency fixed at the value of the fundamental mode with $(\ell=3, m=2)$, that could (and indeed does) contaminate the $\ell=m=2$ mode because of spherical-spheroidal mode mixing.

In each of the six panels of the top two rows, we show for reference: (i) the known (2,2,n) overtone frequencies as black, hollow circles; (ii) the "mirror modes" corresponding to the known (2,-2,n) overtone frequencies as gray, hollow triangles; and (iii) the known (3,2,n) overtone frequencies as gray, hollow squares. When the fundamental mode with $(\ell=3,m=2)$ is fixed (i.e., in model $Q_N^{(f,m)}$) the (3,2,0) is shown as a gray, filled square, as a reminder that it is already included in the waveform and we are not fitting for its complex frequency.

In the panels below we plot various diagnostic quantities; namely, the "frequency error" $\delta\omega$ relative to the expected QNM frequency $\omega_{\rm ref}$, as defined in Eq. (16) (third row), the mode amplitudes A_{n_f} (fourth row), and the mismatches $\mathcal M$ of the fits (fifth and bottom row).

Let us analyze these results starting from the left column. If we fit the waveform with only one damped exponential (N=0), we see that the extracted mode saturates to the fundamental mode n=0. At the peak $(t_0-t_{\rm peak}=0)$ the free mode $n_f=0$ deviates by $\delta\omega<0.1$ from $\omega_{\rm ref}=\omega_{220}$. The disagreement improves as the BH relaxes to its stationary end state, crossing $\delta\omega=10^{-2}$ at $t_0-t_{\rm peak}\simeq15M$, and reaching $\delta\omega<10^{-3}$ at $t_0-t_{\rm peak}=30M$. If we add one more free mode (model $Q_1^{(f)}$, central

If we add one more free mode (model $Q_1^{(J)}$, central panels) we observe that the fundamental mode is extracted much more accurately; the solid blue lines in the third row show that at $t_0 - t_{\rm peak} = 5M$ it deviates from the expected value by $\delta\omega < 0.01$, with $\delta\omega \simeq 3\times 10^{-4}$ at $t_0 - t_{\rm peak} = 17M$. The first overtone (solid orange lines) is extracted much more poorly, with $\delta\omega = 0.2$ at $t_0 = t_{\rm peak}$ and $\delta\omega < 0.1$ after $\approx 6M$. Interestingly, there is a minimum ($\delta\omega \approx 0.003$) at about $t_0 - t_{\rm peak} = 18M$, after which the free mode starts deviating again from the "true" n = 1

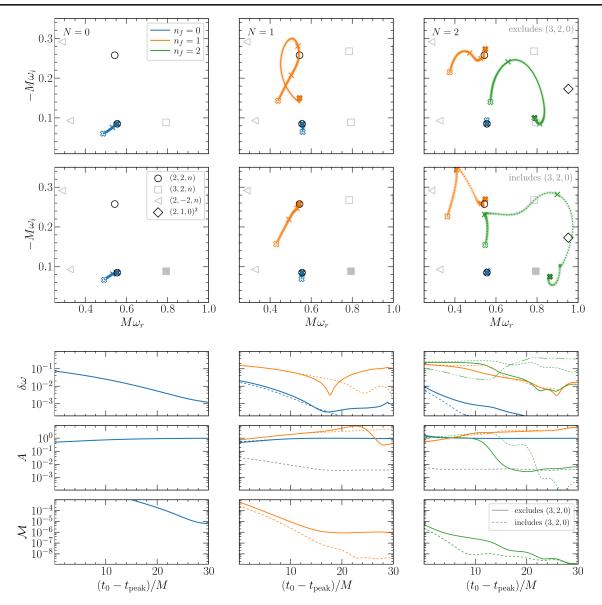


FIG. 10. Fits of the SXS:BBH:0305 waveform with the $Q_N^{(f)}$ and $Q_N^{(f,m)}$ models. First panel: the extracted frequencies $(\omega_r, -\omega_i)$ for N=0 (left), N=1 (middle) and N=2 (right), with t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross). We also mark $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$ by crosses. Second panel: same as first panel, but when adding a mode with complex frequency fixed to ω_{320} , indicated by a filled gray square. Third, fourth, and fifth panels show, respectively, $\delta\omega$, A and M as a function of t_0 . For $n_f=0,1$, $\delta\omega$ is determined compared to $\omega_{\rm ref}=\omega_{22n_f}$, while for $n_f=2$, $\omega_{\rm ref}=\omega_{320}$. Solid lines represent the $Q_N^{(f)}$ model while dotted lines represent the $Q_N^{(f,m)}$ model. With the green dot-dot-dashed line, we also plot $\delta\omega$ with $\omega_{\rm ref}=\omega_{221}$ when fitted by the $Q_N^{(f,m)}$ model. We also plot the A_{320} with a dashed gray line.

overtone. After 27M the deviation is $\delta \omega > 0.1$. This seems to indicate that some other component is spoiling the recovery of the n=1 overtone.

To confirm this conjecture, we add one more free mode (model $Q_2^{(f)}$, right panels). Now the fundamental mode is always extracted extremely well, with $\delta\omega < 0.01$, and even $\delta\omega < 10^{-5}$ at late times. Both of the remaining free frequencies ($n_f = 1, 2$) start off in the vicinity of n = 1 at $t_0 = t_{\text{peak}}$. At later times, $n_f = 1$ (solid orange line) moves towards n = 1, with a minimum of $\delta\omega = 6 \times 10^{-4}$ at 26M. However,

 $n_f=2$ (solid green line) behaves very differently, and at late times it converges towards the (3,2,0) mode, with a relative deviation of $\delta\omega=0.01$ at 27M. This contamination, expected because of spherical-spheroidal mode mixing [37,38,53,156], explains why the $Q_1^{(f)}$ model failed to converge to the n=1 overtone. Since the n=1 is dominant at early times, the free mode $n_f=1$ latches onto it. The (3,2,0) mode, despite having $\mathcal{O}(10^3)$ smaller amplitude than the (2,2,1) mode, decays more slowly and becomes dominant at later times.

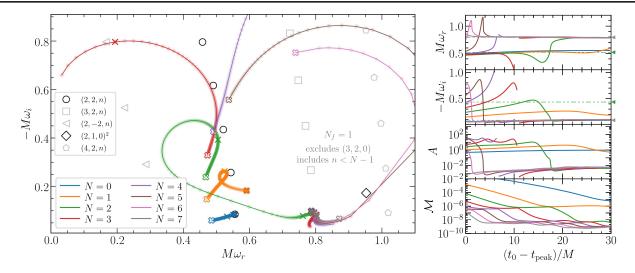


FIG. 11. Fits of the SXS:BBH:0305 waveform with the $Q_{N,1}$ models, with $n=0,\ldots,N-1$ fixed-frequency modes, plus one free-frequency mode. Left panel: extracted frequencies $(M\omega_r,-M\omega_i)$ for N=0-7, with t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross). We also mark $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$ by crosses. We mark the QNMs (2,2,n) (black circles), (2,-2,n) (gray triangles), (3,2,n) (gray squares), (4,2,n) (gray pentagon) and (2,1,0)(2,1,0) (black cross). Right panels plot $M\omega_r,-M\omega_i,A$, and M as a function of t_0 . We also plot $M\omega_{320}$ as a gray dot-dot-dashed line and gray triangle marker.

In the second row of Fig. 10 we "subtract" mode mixing by fixing one damped exponential to the frequency of the (3, 2, 0) mode. Dashed lines in the bottom three rows show the results of fitting the waveform with this "fixed mode mixing," $Q_N^{(f,m)}$ model.

The left column shows that the addition of the (3, 2, 0) mode has a negligible effect on $Q_0^{(f,m)}$. For $Q_1^{(f,m)}$ (central column), including mode mixing further improves the extraction of the fundamental mode at late times, with $\delta\omega = 10^{-5}$ at $t_0 - t_{\rm peak} = 25M$. More importantly, the free mode with $n_f = 1$ now robustly converges towards ω_{221} , crossing the predicted value around 24M, with $\delta\omega \in [10^{-3}, 10^{-2}]$ for $t_0 - t_{\text{peak}} \in [24, 30]M$. With the $Q_2^{(f,m)}$ model (right column), we still observe that $n_f = 1$ saturates toward ω_{221} at late times, with $\delta\omega < 0.01$ between 22–27M. At early times, somewhat surprisingly, the ω_{221} overtone is "picked up" by the second free mode $(n_f = 2)$, which at later times drifts off towards the right of the complex plane, passing by the ω_{321} mode and even showing nonconclusive hints of a possible presence of the nonlinear mode sourced by the square of the $\omega_{2,1,0}$ mode. The presence of these two modes in the simulations is plausible, but none of them has amplitude large enough to be confidently identified.

The fixed-frequency component of $Q_1^{(f,m)}$ has a consistent amplitude $A_{320} \sim 4 \times 10^{-3}$ for $t_0 - t_{\rm peak} \ge 16M$. A similar value is recovered with $Q_2^{(f,m)}$ at even earlier times. This is indication that the additional free component is fitting away some contamination that pollutes the recovery of A_{320} at earlier times.

Another noteworthy feature of Fig. 10 is that mismatches improve significantly when we keep the frequencies free.

For example, at $t_0=t_{\rm peak}$, $Q_1^{(f)}$ ($Q_2^{(f)}$) shows an improvement by one (two) orders of magnitude in the mismatch compared to Q_1 (Q_2). In fact, the $Q_1^{(f)}$ ($Q_2^{(f)}$) models even have better mismatches than the fixed-frequency models Q_2 (Q_4). As we discuss in Appendix C, $Q_3^{(f)}$ (a model with three free frequencies) yields smaller mismatches than the Q_7 . The inclusion of mode mixing has a dramatic effect on mismatches at late times. With the $Q_1^{(f)}$ model (two free modes), the mismatch saturates at $O(10^{-7})$; when we also include mode mixing in model $Q_1^{(f,m)}$, the mismatches are as low as $O(10^{-9})$. A similar improvement occurs going from $Q_1^{(f)}$ to $Q_2^{(f)}$, when the second free mode ($n_f=2$) saturates to the (3, 2, 0) mode frequency.

So far we have used two classes of fitting models $(Q_N^{(f)})$ or $Q_N^{(f,m)}$) in which all complex frequencies are free, except for the (3, 2, 0) mode-mixing component in $Q_N^{(f,m)}$. The results in Fig. 10 show that the free frequencies capture the modes (2, 2, 0), (2, 2, 1), and (3, 2, 0), but there is no evidence for higher overtones (n > 1).

Could the absence of "detectable" higher overtones be an artifact of the use of free-frequency models? To exclude this hypothesis, we now fit the (2, 2) mode by the $Q_{N,1}$ model. As a reminder, this means that one mode frequency is free $(N_f = 1)$, but N mode frequencies (n = 0, ..., N - 1) are fixed to the predicted overtone values. For N = 0, model $Q_{0,1}$ consists (trivially) of a single, free-frequency damped exponential.

In Fig. 11 (which can be compared to the toy-model investigations of Fig. 2) we plot the frequencies, amplitudes, and mismatches extracted in this way. For model $Q_{0.1}$, the free mode $(n_f = 0)$ converges to n = 0 at late

times (blue lines). This late-time convergence is lost when N > 0. For example, model $Q_{1,1}$ is reminiscent of Fig. 10: the free mode frequency "flies by" the (2, 2, 1) overtone at times $t_0 - t_{\text{peak}} \in [10 - 20]M$, but then it drifts to higher frequencies (orange lines).

Interestingly, when N > 1 the free frequency always tends to the (3, 2, 0) mode-mixing QNM, with $\delta\omega$ as low as $O(10^{-3})$ at some point after 20M. The free mode gets close to the expected overtone frequency at intermediate times only for N=2 (green line) and N=3 (red line). As N increases, the free mode latches onto ω_{320} at earlier and earlier times: the free frequency gets within $\delta\omega < 0.1$ of the (3, 2, 0) mode after 17M, 11M, and 7M for N=2, N=3 and N=4, respectively.

Finally, in Fig. 12 we repeat the same experiment, but this time we also fix the (3, 2, 0) mode, i.e., we fit the (2, 2) mode by the $Q_{N,1}^{(m)}$ model. The results are qualitatively similar to Fig. 11; only the fundamental mode and the first overtone saturate to the expected QNM frequencies. By including mode mixing, we manage to get the N=2 fitting model relatively closer to the expected n=2 overtone at $t_0-t_{\rm peak}\sim 13M$. However this mode does not asymptote to the n=2 overtone at late time, and instead it keeps drifting to large values of $M\omega_r$. In the next section, we give a physical interpretation to the observed low-frequency behavior of the free modes.

We have studied three different "agnostic" fitting models (Figs. 10–12). All three models support an important conclusion; the only identifiable physical modes in the $\ell=m=2$ mode of the radiation are (2, 2, 0), (2, 2, 1), and (3, 2, 0). Higher overtones (N>1) cannot be robustly identified by free-frequency fitting.

B. Are we fitting the low-frequency merger signal?

In Figs. 10–12, we observe a clear pattern (that becomes progressively clearer as we remove "undesired" physics, such as spherical-spheroidal mode mixing); at least one free mode starts from low $M\omega_r$ and evolves to higher $M\omega_r$, often (but not always) saturating to one of the expected QNM frequencies. For example, in Fig. 11 the free mode in $Q_{N.1}$ for N=1, 2, 3 always starts at $M\omega_r \simeq 0.47$ and evolves in time (at least initially) to higher $M\omega_r$. This behavior is reminiscent of the well-known "chirping" that characterizes the binary's inspiral. To make the comparison more quantitative, we can compute the instantaneous GW frequency [53,54],

$$\omega_{\rm GW} = \frac{\dot{h}_{22}(t)}{h_{22}(t)}. (20)$$

The black solid line in the right panels of Fig. 12 shows how $M\omega_{\rm GW}$ evolves in time. We also plot, for comparison, the values of $M\omega_r$ extracted using the $Q_{N,1}^{(m)}$ models. Note that the time axes are different; $M\omega_{\rm GW}$ is the instantaneous frequency at time t, while $M\omega_r$ is plotted as a function of the starting time t_0 of the fitting window (the ending point is $t_f = t_{\rm peak} + 90M$). This is why $M\omega_{\rm GW} < M\omega_r$; $M\omega_r$ contains information about physics in the range $[t_0, t_f]$, and as such it includes late-time contributions to $M\omega_{\rm GW}$.

At late times, the real frequencies $M\omega_r$ extracted using the $Q_{0,1}^{(m)}$ model (top right panel) and the $Q_{1,1}^{(m)}$ model (middle right panel) approach the real frequency of the n=0 and n=1 overtones, respectively. As we remarked earlier, this is not true for model $Q_{2,1}^{(m)}$; $M\omega_r$ keeps growing after $t_0-t_{\rm peak}\gtrsim 13M$, instead of saturating to the n=2

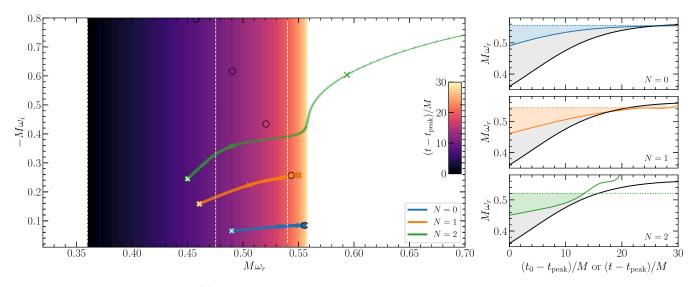


FIG. 12. Fits performed using the $Q_{N,1}^{(m)}$ model with N=0,1,2. Left panel: same as Fig. 11. We also show how $M\omega_{\rm GW}$ evolves (x-axis) as a function of time (see color bar). White dashed lines indicate $t=t_{\rm peak}+10M$ and $t=t_{\rm peak}+20M$ in the instantaneous frequency evolution. The panels on the right show $M\omega_r$ extracted with N=0,1,2 (top, middle and bottom, respectively) as a function of the starting time t_0 in solid colored lines). For comparison, we also plot $M\omega_{\rm GW}(t)$ (black line) and the real part of the expected QNM frequencies (colored horizontal dotted lines).

overtone. In Appendix C, we argue that this large- $M\omega_r$ evolution of the $Q_{2,1}^{(m)}$ model is not physical.

These findings are consistent with other results in the literature. By adding even lower frequency contributions (specifically, modes corresponding to counter-rotating perturbations, which should not be excited unless the binary has large spins antialigned with the orbital angular momentum [84,157,158]), Ref. [112] found that a linear QNM model with a large number of overtones yields low mismatches even before the waveform peak, where a linearized description should not be valid.

These results suggest that once the "physical" ringdown modes (typically the fundamental mode, the first overtone, and the mode-mixing contribution) have been fitted for, additional free modes have a tendency to simply track the GW frequency chirp at early times. Higher overtones have low frequencies, and this observation is consistent with the empirical fact that overtones do such a good job at effectively reducing the mismatch close to the peak.

It is important to emphasize again that lack of evidence is not evidence of absence; while we are unable to extract higher overtones even when including a large number of free modes, this does not imply the absence of overtones altogether. The extraction of overtones is challenging due to the presence of additional physics and nonlinear effects. It may still be possible to achieve it by incorporating a more comprehensive understanding of all postmerger nonlinearities, and allowing for the dynamical character of the timevarying background.

C. Extracting BH properties in the presence of nonlinearities

The comprehensive fitting experiments performed so far show that higher overtones are hard to extract from numerical waveforms, and they are often unlikely to be physical. In the infinite SNR limit, how do frequencies inferred by fitting free damped exponentials translate into mass and spin values? This is the important question to ask in the spirit of the original, agnostic BH spectroscopy proposal. The answer depends on the starting time of the fit, and it is affected by the inclusion of mode-mixing contributions from the (3, 2, 0) mode.

Consider first model $Q_1^{(f)}$, with two free complex frequencies. We will identify the $n_f = 0$ component with the (2, 2, 0) mode and the $n_f = 1$ component with the (2, 2, 1) mode. A posteriori, we now assume that these damped exponentials correspond to the GR QNMs. Then we can invert the (complex) functional relationships $\omega_{\ell mn}(M_f, \chi_f)$ provided in Ref. [40] separately for the (2, 2, 0) and (2, 2, 1) modes, as follows:

$$\chi_f^{\inf} = \mathcal{Q}_{\ell mn}^{-1} \left(-\frac{\omega_r}{2\omega_i} \right),$$

$$M_f^{\inf} = \frac{M\omega_r}{\text{Re}[\omega_{\ell mn}(\chi_f^{\inf}, M)]}.$$
(21)

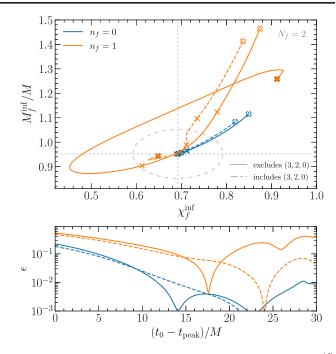


FIG. 13. Remnant mass and spin $(M_f^{\rm inf},\chi_f^{\rm inf})$ inferred using $Q_1^{(f)}$ [solid lines] and $Q_1^{(f,m)}$ [dashed lines] for t_0 ranging from $t_{\rm peak}$ (hollow crosses) to $t_{\rm peak}+30M$ (filled crosses). We also mark $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$ by crosses. Mass and spin are calculated directly from ω_{n_f} , without assuming any no-hair theorem enforced relation between the modes. The gray vertical (horizontal) lines mark the true χ_f (M_f) . The dashed ellipse corresponds to $\epsilon=0.1$.

Here $\mathcal{Q}_{\ell mn}^{-1}$ is the inverse of the monotonic function relating the quality factor and the spin, $\mathcal{Q}_{\ell mn}(\chi_f)$ [40]. If the modes can be robustly inferred from the data, the mass and spin inferred from each of these two inversions should coincide. In fact, performing this exercise with three or more "measured" parameters is a null test of GR.

We already know from Fig. 10 that in a fit with two free modes, neither of the extracted frequencies at $t_0 = t_{\rm peak}$ coincides with the first overtone: the difference between the "recovered first overtone" and the true (2, 2, 1) mode is $\delta \omega > 0.1$. These recovered frequencies will inevitably give wrong estimates for the mass and spin of the remnant.

Let us assume that the frequencies extracted at the peak with model $Q_1^{(f)}$ are physical QNMs. In Fig. 13 we compute the remnant mass and spin inferred from each overtone under this assumption. The (M_f,χ_f) values inferred by assuming that the free mode $n_f=0$

⁸This assumption is implicit in the discussion of Ref. [96]. However, the low SNR of GW150914 implies that the measurement is dominated by statistical errors, and therefore the inferred mass and spin can be compatible with GR. Our discussion here focuses on systematic errors, that will dominate at large SNR.

corresponds to the (2, 2, 0) mode converge monotonically towards the remnant BH value as t_0 increases, with $\epsilon < 0.01$ after $t_0 \gtrsim 11M$ (solid blue line in the bottom panel). The addition of mode mixing improves the recovered values only after $t_0 \gtrsim 11M$ (dashed blue line).

On the other hand, the values obtained through the $n_f=1$ mode assuming that it corresponds to the (2,2,1) mode are significantly biased $(\epsilon>0.1)$ until $t_0\sim 15M$. Around $t_0\sim 18M$, and for a short interval $\Delta t\sim 1M$ around it, $\epsilon<0.01$, but the accuracy in the inferred mass and spin decreases at later times. The late-time inference of mass and spin with this second mode can be marginally improved by including mode mixing (model $Q_1^{(f,m)}$); this keeps $\epsilon\lesssim 0.1$ even at late times. In any case, this experiment clarifies that even in the infinite SNR limit, it is not easy to get percent-level estimates of the remnant parameters with the first overtone in an agnostic search.

The hollow crosses in the top panel of Fig. 13 show that when the fit starts at the peak, both the mass and the spin are significantly overestimated. Note also that in this exercise we find mass and spin by inverting each of the two modes independently. This is different from Ref. [96], where both frequencies were assumed to correspond to their GR values for a given (M_f, χ_f) .

This exercise has an interesting implication for data analysis. At the relatively low SNR of GW150914, the errors on the recovered parameters are so large that the $Q_1^{(f)}$ model gives complex frequency estimates compatible with the true quasinormal frequencies [96,121]. However our slightly generalized version of the test, where we relax the assumption that all modes are related to mass and spin as in GR, reveals that the $Q_1^{(f)}$ model is effective at reducing the mismatch at the peak, but it does not reproduce the physical QNM content of signal.

Reference [96] claimed evidence for the first overtone, in part, because the $Q_1^{(f)}$ model yields better estimates of the remnant mass and spin than the $Q_0^{(f)}$ model at the peak. As we have shown in Fig. 10, the complex frequency extracted with the $Q_0^{(f)}$ model at the peak is indeed biased because of the low-frequency merger signal. It is true, as claimed in Ref. [96], that the $Q_1^{(f)}$ model yields better estimates of the remnant's parameters. However this is not because the additional exponential gives a good estimate of the first overtone frequency, but simply because it fits away the low-frequency merger. This allows the fundamental mode to be extracted more accurately, improving the remnant mass and spin estimate.

While the $Q_1^{(f)}$ model does not allow us to extract the first overtone at $t_0 = t_{\text{peak}}$, leading to large deviations in the inferred values of (M_f, χ_f) , one might argue that the addition of more free modes could improve the situation. However, this is not the case. The agnostic models we explored do not lead to an unbiased extraction of the first overtone at the peak. This is

illustrated, for example, by Fig. 25 in Appendix C, where we fit the (2, 2) multipole with as many as five damped sinusoids while also including mode mixing. Even with such a large number of modes, and despite the mismatch $\mathcal{M} \sim \mathcal{O}(10^{-9})$ being so small that one runs into the risk of fitting numerical noise, we cannot extract the first overtone at $t_0 = t_{\rm peak}$.

It is by now common practice to fit GW data with the fundamental mode plus one or two overtones, treating the remnant mass and spin (along with the mode amplitude and phase) as free parameters. Many studies claim evidence for an overtone when (i) fits with either one or two overtones lead to values of (M_f, χ_f) which are in better agreement with the values extracted from the full waveform than fits using only the fundamental mode; and (ii) the amplitude of the additional overtone (or overtones) is nonzero [96,105,115,120–123]. As we argued in Sec. IV, the fact that adding QNMs leads to a better extraction of mass and spin does not imply that these additional QNMs are physical. This is confirmed, e.g., by the plots in the central column (N = 1) of Fig. 10. By adding an extra mode (orange line) we can better fit the low-frequency merger signal at early times, which leads to a more accurate extraction of the fundamental mode and improves the estimate of (M_f, χ_f) —but not because of any physical contribution from the first overtone. It should also be clear that a nonzero amplitude of the overtone at the peak is not sufficient to claim detection of the overtone, as the additional QNM is simply (over)fitting the low-frequency merger signal. To make any credible detection claims, overtones should be extracted in an agnostic fashion.

There is another important caveat. In this whole analysis we focused on the dominant (2, 2) mode, but real data will also contain multipoles such as (4, 4), which are usually dominant over the (3, 2, 0) mode (with amplitude $A_{320} \sim 4 \times 10^{-3}$). This will inevitably complicate the problem in ways that should be addressed in future work.

The main goal of this section is to point out two important limitations of BH spectroscopy using only overtones: (1) at early times, we do not observe the excitation of the first overtone; instead, the second mode in a free-frequency fit tries to pick up the low-frequency merger signal. (2) while the fundamental mode gives reliable estimates of the remnant mass and spin at late times, the first overtone never allows for percent-level estimates of the mass and spin; it may allow for $\sim 10\%$ estimates of the mass and spin at late time for loud enough signals, but only if we carefully take into account mode mixing and higher multipoles.

D. Higher multipoles: Overtones or nonlinearities? The example of $(\mathcal{C}, m) = (4, 4)$

It is instructive to apply the techniques we deployed for the (2, 2) multipole to higher multipoles of the

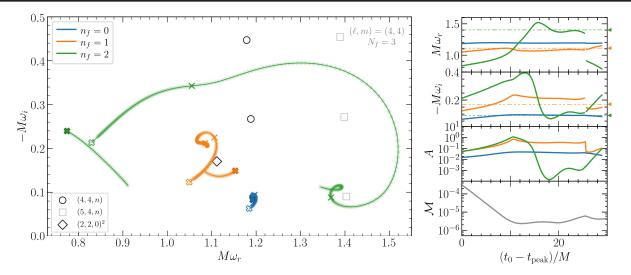


FIG. 14. Fits of the (4, 4) multipole using the $Q_2^{(f)}$ model. Left panel: the extracted frequencies $(M\omega_r, -M\omega_i)$ for $n_f = 0, 1, 2$, with t_0 ranging from t_{peak} (empty cross) to $t_{\text{peak}} + 30M$ (filled cross); other crosses correspond to times $t_{\text{peak}} + 10M$ and $t_{\text{peak}} + 20M$. We mark the QNMs (4,4,n) (black circles), (5,4,n) (gray squares), and the quadratic mode (2, 2, 0)(2, 2, 0) (black diamond). Right panels: $M\omega_r, -M\omega_i, A$ and M as functions of t_0 . The horizontal orange (green) dot-dot-dashed lines and triangle markers correspond to the expected frequencies for $M\omega_{220\times220}$ ($M\omega_{540}$).

radiation. To improve readability, we discuss the full waveform including higher multipoles in Appendix F. There we observe that: (1) BH spectroscopy with overtones at early times (assuming they are excited) requires good modeling of the low-frequency merger signal; and (2) BH spectroscopy with overtones at late times should be done after accounting for other multipoles, which can often be more robustly extracted from the waveform.

Here we single out a particularly interesting example: the (4, 4) multipole, which is known to contain a quadratic mode (2, 2, 0)(2, 2, 0) sourced by the square of the (2, 2, 0) mode [143-145]. We fit this multipole with the $Q_2^{(f)}$ model $(N_f=3$ free modes). In Fig. 14 we plot the three extracted frequencies, their amplitudes, and the corresponding mismatches.

The fundamental mode (4, 4, 0) can be extracted with $\delta\omega < 0.01$ after $t_0 - t_{\rm peak} > 5M$, and an average $\delta\omega \sim \mathcal{O}(10^{-3})$ in the time range shown. The quadratic mode (2, 2, 0)(2, 2, 0) can be identified with $\delta\omega \sim 10^{-2}-10^{-1}$ [144,145]. Its frequency is fairly well recovered, while the damping time tends to be overestimated compared to the predicted value. The third free mode starts off at lower $M\omega_r$, evolving towards the (5, 4, 0) mode in the time range $t_0 - t_{\rm peak} = 15-25M$, and moving to lower frequencies again at later times. We do not find any evidence of excitation of the first overtone. This conclusion holds even if we fix the lowest modes, (4, 4, 0) or (5, 4, 0), to their values as predicted in GR.

The main message of this analysis is clear—the nonlinear mode is easier to recover in a free-frequency fit than the first linear overtone.

VI. CAN THE FIRST OVERTONE BE EXTRACTED IN THE PRESENCE OF SUBDOMINANT MULTIPOLES?

In Sec. VA we discussed how spherical-spheroidal mixing of the (3, 2) multipole affects the extraction of the first overtone from the dominant (2, 2) multipole. We concluded that at least $N_f = 3$ free modes are required to extract the first overtone and that it is easier to extract (long-lived) fundamental modes than fast-decaying overtones, even if the latter have a much larger amplitude. In real-world data analysis problems, the waveform will in general be a superposition of several multipoles. We now ask how many free modes would be necessary to extract the first overtone (2, 2, 1) when the actual waveform is a superposition of several multipoles?

We fit the waveform

$$h(t,\theta,\phi) = \sum_{\ell m} h_{\ell m}(t)_{-2} Y_{lm}(\theta,\phi)$$
 (22)

using the $Q_N^{(f)}$ model. For simplicity we focus on $\ell \leq 4$. We quantify the relative contribution of subdominant multipoles through the ratio

$$R_{\ell m} = \frac{\max_{t} (h_{\ell m}(t)_{-2} Y_{lm}(\theta, \phi))}{\max_{t} (h_{22}(t)_{-2} Y_{22}(\theta, \phi))}.$$
 (23)

We consider two scenarios: $(\theta, \phi) = (0, \pi/4)$ ("GW150914-like scenario") and $(\theta, \phi) = (\pi/9, \pi/4)$ (higher-multipole scenario).

A. Face-on, GW150914-like scenario

If for simplicity we assume $(\theta, \phi) = (0, \pi/4)$, a face-on configuration compatible with the orientation measured for GW150914 [159], we find that the only observable spherical multipoles are (2, 2), (3, 2), and (4, 2), with $R_{32} = 0.08$ and $R_{42} = 0.006$.

In Fig. 15 we show the extracted frequencies, along with their amplitudes and the mismatch, obtained by fitting $h(t,0,\pi/4)$ with the $Q_2^{(f)}$ model (i.e., $N_f=3$ free modes). While we recover the (3, 2, 0) mode, the first overtone (2, 2, 1) cannot be reliably extracted. This is somewhat surprising, because the $Q_2^{(f)}$ model was sufficient to extract the first overtone from the $\ell=m=2$ multipole (compare the central row of Fig. 10).

In Fig. 16 we make the obvious next step, and we try again to fit $h(t, 0, \pi/4)$, but this time with the $Q_3^{(f)}$ model.

We now find that the first overtone (orange line) can be identified with sufficient accuracy at $\sim t_{\text{peak}} + 12M$. However, the (3, 2, 0) mode (green line) can be extracted more accurately and at earlier times, and therefore this may be a more natural "second mode" to choose when performing tests of GR. Besides (as expected) the recovery of the (3, 2, 0) mode improves as we include more free modes in the fit, moving from $\mathcal{Q}_2^{(f)}$ to $\mathcal{Q}_3^{(f)}$. As usual, this is because the additional free mode fits other early-time contributions which, if not accounted for, can harm the recovery of the (3, 2, 0) mode. In fact, the fourth free mode in the $Q_3^{(f)}$ model recovers the (4, 2, 0) spherical-spheroidal mixing mode frequency (red line) at about the same t_0 at which the (2, 2, 1) mode is found, and with comparable accuracy. The (4, 2) multipole—despite having a small relative amplitude, $R_{42} = \mathcal{O}(10^{-3})$ —can affect the extraction of the (2, 2, 1) overtone.

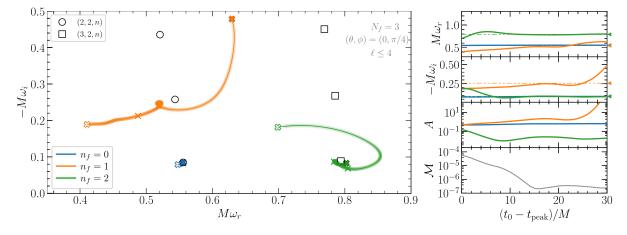


FIG. 15. Fitting $h(t,\theta,\phi)$ at $(\theta,\phi)=(0,\pi/4)$ using the $Q_2^{(f)}$ model $(N_f=3)$. Left panel: the extracted frequencies $(M\omega_r,-M\omega_i)$ for $n_f=0,1,2$, with t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross). We also mark $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$ by crosses. Right panels: $M\omega_r,-M\omega_i,A$, and M as a function of t_0 . We show $M\omega_{220},M\omega_{221},M\omega_{320}$ as blue, orange, and green dot-dot-dashed horizontal lines and triangle markers.

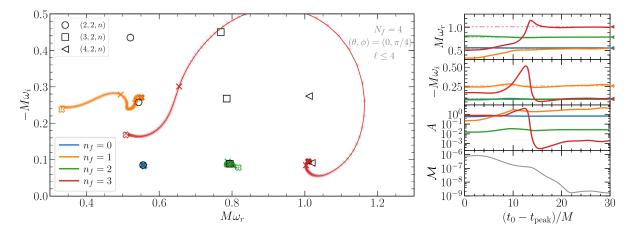


FIG. 16. Fitting $h(t,\theta,\phi)$ at $(\theta,\phi)=(0,\pi/4)$ using the $Q_3^{(f)}$ model $(N_f=4)$. Left panel: the extracted frequencies $(M\omega_r,-M\omega_i)$ for $n_f=0,1,2,3$, with t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross). We also mark $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$ by crosses. Right panels: $M\omega_r,-M\omega_i,A$, and M as a function of t_0 . We show $M\omega_{220},M\omega_{221},M\omega_{320},M\omega_{420}$ as blue, orange, green, and red dot-dot-dashed horizontal lines and triangle markers.

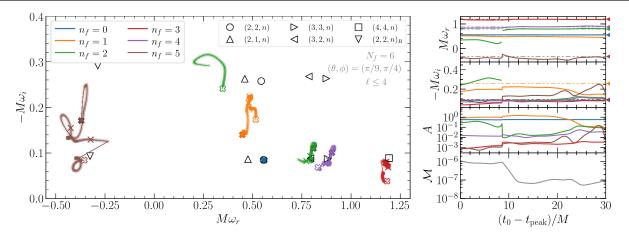


FIG. 17. Fitting $h(t,\theta,\phi)$ at $(\theta,\phi)=(\pi/9,\pi/4)$ using the $Q_5^{(f)}$ model $(N_f=6)$. Left panel: the extracted frequencies $(M\omega_r,-M\omega_i)$ for $n_f=[0,\ldots,5]$, with t_0 ranging from t_{peak} (empty cross) to $t_{\text{peak}}+30M$ (filled cross). We also mark $t_{\text{peak}}+10M$ and $t_{\text{peak}}+20M$ by crosses. Right panels: $M\omega_r,-M\omega_i,A$, and M as a function of t_0 . We also plot $M\omega_{220},M\omega_{221},M\omega_{320},M\omega_{330},M\omega_{440}$, and the retrograde mode $M\omega_{220,R}=-M\omega_{2-20}^*$ as blue, orange, green, purple, red, and brown dot-dot-dashed lines and triangle markers.

Even in the relatively optimistic case of a face-on binary, when only the (3, 2) and (4, 2) multipoles significantly contribute to the strain, the extraction of the first overtone (2, 2, 1) relies on the successful extraction of the (long-lived) fundamental modes of subdominant multipoles.

B. A nonface-on binary

The face-on scenario considered above is idealized and somewhat optimistic, since the assumption that $\theta=0$ artificially removes many multipolar components. For more generic binary orientations we would expect more multipoles to contribute to the strain, and therefore the recovery of the first overtone should be more difficult.

We now consider a more realistic case where the inclination is not exactly face-on; $(\theta,\phi)=(\pi/9,\pi/4)$. Now the dominant multipoles have $R_{32}=0.06$, $R_{33}=0.03$, $R_{44}=0.01$, $R_{21}=0.01$, and in addition there are five multipoles with $10^{-3} < R_{\ell m} < 10^{-2}$.

In Fig. 17 we show the extracted frequencies, amplitudes and mismatches when we fit $h(t,\pi/9,\pi/4)$ with the $Q_5^{(f)}$ model. Now the first overtone cannot be extracted even when we include $N_f=6$ free modes; the only modes that can be cleanly extracted are (2,2,0),(3,3,0),(3,2,0), and (4,4,0). The free mode labeled by $n_f=1$ (orange line) approaches the (2,1,0) mode, but the model fails to extract this mode at times $t_0-t_{\rm peak}<30M$. The $n_f=5$ free mode (brown line) tries to fit a "retrograde" mode, that we label by $(2,2,0)_R$. Here, following Ref. [85], we call "retrograde" those QNMs that have frequency $M\omega_{\ell mn,R}=-M\omega_{\ell-mn}^*$. In Appendixes C and F 3 we show that retrograde modes are present also in the (2,2) and (2,1) multipoles, respectively.

The case we studied in this section is still an optimistic scenario for overtone detection for two main reasons.

First of all, GW150914-like signals have nearly equal masses, so that the subdominant multipoles have relatively

small amplitudes. Even in this "optimal" case, the extraction of the (2, 2, 1) overtone is severely affected by higher multipoles. The extraction of the (2, 2, 1) overtone will be more difficult (while the extraction of subdominant multipoles will become easier) for binary configurations with asymmetric masses and large spins, where higher multipoles are more excited. For this reason, asymmetric systems are better targets for BH spectroscopy.

Secondly, if we assume random orientations, more than 94% of the sources have $\theta > \pi/9$. As we increase θ , the contribution of subdominant multipoles increases, and isolating the (2, 2, 1) overtone gets harder. In fact, we have experimented and found that the (2, 2, 1) overtone cannot be extracted with good accuracy at even smaller inclinations, such as $\theta = 5^{\circ}(=\pi/36)$.

We conclude that in practical scenarios, it would be hard to extract the first overtone from GW150914-like signals. This is because at the times $t_0 \gtrsim t_{\rm peak} + 12M$ when the (2, 2, 1) mode is recovered by the fits, the fundamental modes of subdominant multipoles are usually dominant. Even when the overtone can be confidently identified, the fundamental modes of subdominant multipoles are probably more reliable for performing BH spectroscopy tests.

These considerations ignore detector noise. Depending on the mass and spin of the remnant, noise could further reduce the "effective" amplitude of subdominant mutipoles, impeding their extraction from the data. Our inability to extract these subdominant multipoles can significantly affect the extraction of the (2, 2, 1) overtone.

VII. CONCLUSIONS AND DISCUSSION

The examples discussed so far show that the inclusion of a large number of overtones in BH spectroscopy has no physical justification. Claims that one or more overtones can be observed at the peak of the waveform are theoretically inconsistent with the analysis of NR waveforms. In fact, the analysis of the previous section leads to an unavoidable conclusion; BH spectroscopy using overtones is either unfeasible or suboptimal compared to spectroscopy using additional multipoles.

A. Observational implications

Independently of the current debate on the statistical significance of high-overtone fits to the GW150914 data [96,105,119–123], our results demonstrate that fits with multiple modes starting at $t_{\rm peak}$ do not correspond to the physical excitation of the QNMs of the remnant, but are simply effective at overfitting the low-frequency components of the waveform near merger. As shown in Appendix E, these models would give rise to false evidence in favor of overtones even when the postpeak signal is unphysical and overtones should not be present. Any interpretation of overtone measurements close to the peak (i.e., outside the regime of validity of linear perturbation theory) is devoid of physical meaning and leads to unreliable BH spectroscopy tests.

Tests based on early-time overtone models are conceptually equivalent to parametrized tests where low-frequency effective components are used to model the near-merger signal. Such tests have already been performed by the LIGO-Virgo Collaboration at the time of the discovery of GW150914 [160], and they have been extended to the GW catalogs in subsequent analyses [68–70]. More recently, modifications to merger quantities were considered within an effective-one-body model [161]. These are certainly valuable investigations, but they cannot be regarded as BH spectroscopy tests as they do not probe directly the excitation of the QNMs of the remnant BH, but rather the nonlinear portion of the signal, whose structure is not analytically understood.

These arguments also apply to tests of the Hawking area law based on overtone measurements [162]. For GW150914-like binaries, these tests are a time-domain reformulation of the inspiral-merger-ringdown consistency tests and they are equivalent to frequency-domain tests, since frequency leakage between pre and postmerger is negligible (as shown in Ref. [163]).

B. Pseudospectral instabilities and nonlinearities

Our analysis raises some interesting questions. Why are we unable to extract higher overtones even in linear perturbation theory? Why are the overtone frequencies that we infer by fitting time-domain signals different from the expected values?

Some answers to these questions could come from recent work on the pseudospectrum. Nollert demonstrated a long time ago that small changes in the potential can destabilize the whole QNM spectrum [164] (see also [165]). Recent calculations of the so-called "pseudospectrum" in the

context of BH physics put these ideas on a more solid mathematical foundation, showing that the overtones are more easily destabilized than the fundamental mode (while claiming that the fundamental mode is, in fact, stable under certain classes of perturbations) [138]. Later work showed that even the fundamental mode is not immune from spectral instabilities [140]. In practice, this is unlikely to affect our ability to do BH spectroscopy with the fundamental mode in observationally relevant scenarios [141].

What is still unclear is whether spectral instabilities affect overtones, and to what extent. The intrinsic nonlinearity of GR implies that the linear wave equations are always affected by second-order terms that can be treated as perturbations [166–173]. There has been some encouraging progress in our understanding of nonlinearities in BH physics [47,142,174,175]. However several open questions remain. Are the perturbations induced by nonlinear effects sufficient to destabilize most overtones in the context of comparable-mass BH mergers? Can the spectral instability of the overtones explain our failure to find them in timedomain waveforms? Do these spectral instabilities explain why overtones are so hard to identify?

C. Amplitude parametrizations and Occam penalties

Besides being physically inconsistent for reasons we have already discussed, multiple overtone models with free amplitudes are problematic from the point of view of data analysis. Depending on the SNR of the signal, it may be necessary to include several overtones to obtain unbiased inference. However, each QNM comes with two additional free parameters—an amplitude and a phase. As pointed out in Ref. [109], a large number of parameters generally decreases the model evidence due to the associated Occam penalty. Reference [109] concluded that even for extremely loud signals, the Occam penalty would not allow us to use overtones to test the "no-hair" hypothesis. It also reported a significant increase in the Bayesian evidence if one uses phenomenological templates rather than an agnostic superposition of damped sinusoids.

In principle, this problem can be alleviated by parametrizing the overtone amplitudes as function of the parameters of the progenitors. For example, Ref. [75] presented a fit of the first overtone amplitude as a function of binary parameters for nonspinning progenitors valid at times $t_0 > 10 M_f$ (see also Ref. [78]). However, later work found that when we try to push such "global" fits to earlier times, the result is not robust under variations of t_0 (see e.g., Figs. 3 and 4 of Ref. [82]). Even if we can model the overtone amplitudes as functions of the properties of the remnant progenitors, measuring several overtone frequencies in real data may be impractical. Fisher matrix estimates suggest that it is easier to obtain evidence for multiple modes using higher-angular multipoles rather than overtones [82]. These conclusions are in disagreement with

Ref. [105], which however assumed the overtone model to be valid at the peak and employed a different detection criterion.

In future work we will systematically rank the most important QNMs in the spectrum as functions of the progenitor parameters, investigate strategies for a robust modeling and measurement of multiple modes, and revisit previous BH spectroscopy horizon estimates [83,92].

D. Weak and strong "no-hair" tests

Recent observational implementations of BH spectroscopy tests are quite different in nature [69,70,95,96,105,119–121,176]. Since the work of Ref. [86], it has been customary to perform "no-hair" tests not by directly measuring multiple QNM frequencies—as envisioned in the original BH spectroscopy proposals [21,30,40]—but rather by measuring the mass and the spin of the BH (the impact of charge is currently negligible [100]) together with a single parametric deviation, e.g., $\delta\omega_{221}$.

This parametrization is sensible. The use of physical degrees of freedom reduces degeneracies, decreases the number of parameters in the test to three (as opposed to the four real parameters that correspond to measuring two complex frequencies), and consequently decreases the SNR required for the test. However, the parametrization only allows for a "weak" BH spectroscopy test, because we must assume *a priori* which modes are present in the data. As such, it must be applied with great care. This "weak" version of the test could be consistent with GR just because it enforces the presence of (spurious) low-frequency modes, which are effective at fitting nonlinearities or other physical artifacts near the peak.

Conversely, the original BH spectroscopy proposal relies on a "strong," more conservative version of the test: we should perform an agnostic search for multiple damped exponentials, and only *a posteriori* identify them with a given set of modes, provided that the results of the agnostic frequency search are robust enough. This strong version was used in most of the fitting experiments reported in this manuscript, it does not impose the presence of a given mode in the data, and (while harder to implement) it leads to more robust conclusions.

The key take-away message for the purpose of this discussion is that if the mode is actually present in the data, in the large-SNR limit it must be possible to replace the weak form of the test by its strong version. As we have shown in Sec. V C above, if we carry out BH spectroscopy tests with one overtone close to $t_{\rm peak}$, this does not happen even in the infinite SNR limit, i.e., when we fit state-of-theart NR simulations. The opposite is true for tests that use the fundamental modes corresponding to different multipoles at later times. In either form of the test, using the fundamental modes has an additional observational advantage: at times t_0 significantly larger than $t_{\rm peak}$, the statistical uncertainty on the starting time has a much smaller impact.

For sufficiently large SNR, fundamental modes with different values of (ℓ, m) can be robustly identified over multiple cycles, and errors $\Delta t_0 \sim \mathcal{O}(1) M_f$ have negligible impact on their recovery. On the contrary, the starting time uncertainty has a large impact on overtones-based tests near the peak [119,121].

E. Beyond-Kerr parametrizations: Effective detection templates versus physical parametrizations

Any BH spectroscopy tests relying on templates with a single parametric deviation, such as those in Ref. [69], are unphysical. This is because virtually all modified-gravity theories are expected to change all of the QNM frequencies, not just one [177–183] (this discussion applies also to parametrized tests in the inspiral: see e.g., Sec. II of Ref. [104]). However, it is important to make a distinction between two different goals; the detection of beyond-GR effects, and the identification of their origin.

We can safely use unphysical models as effective detection templates with the minimal number of parameters required to detect hypothetical deviations from GR. While using multiple parameters may become beneficial in the context of next-generation detectors [184,185], adding too many parameters can introduce correlations, increase the SNR required for a detection of beyond-GR effects, and decrease the sensitivity of the search. In addition, single-parameter templates have been shown to be effective at recovering deviations which affect multiple parameters [186]. If future observations using these simple parametric tests show any hints of deviations from GR, they should be followed up with better-motivated templates to identify their origin.

One proposal to bypass this problem in ringdown tests is to use parametrized frameworks such as ParSpec [103,104]. The idea is to perform a perturbative expansion around GR, in which corrections to the individual QNM frequencies are proportional to the (small) couplings that appear in the action of beyond-GR theories. This is a sensible parametrization that introduces additional information, and increases the sensitivity of searches for new physics at fixed SNR [104]. The sensitivity can be further increased for specific beyond-GR theories [187], which can be easily embedded within a ParSpec-like framework in the small-coupling limit. A valid objection to this approach is that ParSpec-like frameworks, being intrinsically perturbative, could miss nonperturbative deviations from the GR spectrum. This is true but perhaps unlikely, given mounting evidence that the observed GW signals must be quite close to the predictions of GR. It may be advisable to run two parallel searches: a more sensitive, ParSpec-like search for perturbative deviations, and a less sensitive (but more general) search for nonperturbative signatures.

Note added.—Soon after this manuscript appeared on the arXiv, a preprint by Nee et al. [188] studied the excitation of overtones by comparing Regge-Wheeler and Pöschl-Teller potentials. Their findings are similar to those presented in Sec. III, i.e., "... large overtone numbers may [...] remove nonquasinormal mode contributions that are relevant at early times of a ringdown, but do not necessarily correspond to the physical excitation of modes of the system" [188].

ACKNOWLEDGMENTS

We are grateful to Swetha Bhagwat, Yanbei Chen, Kyriakos Destounis, Will Farr, Pedro Ferreira, Xisco Jiménez Forteza, Matt Giesler, Lam Hui, Max Isi, Danny Laghi, Macarena Lagos, Paul Lasky, Lionel London, Sizheng Ma, Keefe Mitman, Paolo Pani, Rodrigo Panosso-Macedo, Lorenzo Pierini, Leo Stein, Saul Teukolsky, Eric Thrane, and Sam Wong for helpful discussions. M. H. Y. C. is a Croucher Scholar supported by the Croucher Foundation. E. B., R. C., and M. H. Y. C. are supported by NSF Grants No. AST-2006538, No. PHY-2207502, No. PHY-090003, and No. PHY20043, and NASA Grants No. 19-ATP19-0051, No. 20-LPS20-0011, and No. 21-ATP21-0010. This research project was conducted using computational resources at the Maryland Advanced Research Computing Center (MARCC). V.C. is a Villum Investigator and a DNRF Chair, supported by VILLUM Foundation (Grant No. VIL37766) and the DNRF Chair program (Grant No. DNRF162) by the Danish National Research Foundation. V. C. acknowledges financial support provided under the European Union's H2020 ERC Advanced Grant "Black holes: gravitational engines of discovery" Grant Agreement No. Gravitas-101052587. Views and opinions expressed are, however, those of the author only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. G. C. acknowledges support by the Della Riccia Foundation under an Early Career Scientist Fellowship. G. C. acknowledges funding from the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie Grant Agreement No. 847523 "INTERACTIONS," from the Villum Investigator program supported by the VILLUM Foundation (Grant No. VIL37766) and the DNRF Chair program (Grant No. DNRF162) by the Danish National Research Foundation. F. D. acknowledges financial support provided by FCT/Portugal through Grant No. SFRH/BD/ 143657/2019. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie Grant Agreement No. 101007855. We thank FCT for financial support through Project No. UIDB/00099/2020. We acknowledge financial support provided by FCT/Portugal through Grants No. PTDC/MAT-APL/30043/2017 and No. PTDC/FIS-AST/7002/2020. Part of E. B.'s work was

performed at the Aspen Center for Physics, which is supported by National Science Foundation Grant No. PHY-1607611. The authors would like to acknowledge networking support by the GWverse COST Action No. CA16104, "Black holes, gravitational waves and fundamental physics."

APPENDIX A: AMBIGUITIES IN THE DEFINITION OF RINGDOWN STARTING TIME AND RINGDOWN EXCITATION

A ringdown waveform is formally defined as the superposition of damped exponential given in Eq. (2). One of the main issues we address in this paper is the following: Is this definition *ever* a good approximation of a given waveform, and to what level of approximation?

In this appendix we use a simple toy model and some examples from linear BH perturbation theory to clear up possible misunderstandings about (i) the very possibility to define a ringdown "starting time," and (ii) the meaning of "ringdown excitation."

1. A Green's function toy model

The first observation is that conservative systems described by linear wave equations can always be described in terms of their eigenmodes. However, it is well known that a simple description in terms of constant-amplitude eigenmodes may not be adequate even in such systems. Consider for illustration the following wave equation in one spatial dimension, explored already in Refs. [47,189]:

$$\frac{\partial^2 \Psi}{\partial x^2} - \frac{\partial^2 \Psi}{\partial t^2} - 2V_0 \delta(x) \Psi = 0. \tag{A1}$$

Through a Laplace transform, $\psi(s,x) = \int_0^\infty e^{-st} \Psi(t,x) dt$, we can transform the partial differential equation into the ordinary nonhomogeneous equation

$$\frac{d^2\psi}{dx^2} + \left[\omega^2 - 2V_0\delta(x)\right]\psi = -\frac{\partial\Psi}{\partial t}\Big|_{t=0} + i\omega\Psi\Big|_{t=0} \equiv S,$$
(A2)

where we have defined the frequency $\omega = is$ and a source term S. Define the "left" homogeneous solution as that satisfying outgoing boundary conditions at large negative values of x,

$$\psi_L = \begin{cases} e^{-i\omega x}, & x \to -\infty \\ A_{\text{in}} e^{-i\omega x} + A_{\text{out}} e^{i\omega x}, & x \to +\infty \end{cases}$$
 (A3)

and the "right" homogeneous solution as

$$\psi_R = e^{i\omega x}, \qquad x \to +\infty.$$
 (A4)

The (constant) Wronskian W of these two solutions is $W \equiv \psi_L \psi_R' - \psi_L' \psi_R = 2i\omega A_{\rm in}$. By demanding continuity of ψ_L and a discontinuity in its derivative consistent with a δ function, we find

$$A_{\rm in} + A_{\rm out} = 1,\tag{A5}$$

$$A_{\text{out}} - A_{\text{in}} + 1 = \frac{2V_0}{i\omega},$$
 (A6)

which is equivalent to

$$A_{\text{out}} = -i\frac{V_0}{\omega}, \qquad A_{\text{in}} = 1 + i\frac{V_0}{\omega}. \tag{A7}$$

The Wronskian vanishes at the single zero $\omega = -iV_0$, which is therefore a QNM frequency, and the only one supported by this potential.

The solution of Eq. (A2) satisfying outgoing boundary conditions as $x \to \infty$ is

$$\psi = e^{i\omega x} \int dx \frac{\psi_L S}{W}. \tag{A8}$$

Consider for simplicity an initial pulse $S = \delta(x - x_0)$. The integral then yields

$$\psi = \frac{e^{i\omega(x-x_0)}}{2} + e^{i\omega(x+x_0)} \frac{A_{\text{out}}}{2A_{\text{in}}}.$$
 (A9)

This is the full solution in the Fourier plane. To get the time-domain solution, we calculate the inverse,

$$2\pi\Psi = \int d\omega e^{-i\omega t} \psi = \int d\omega \frac{e^{-i\omega(x-x_0-t)}}{2} + \int d\omega e^{-i\omega(x+x_0-t)} \frac{A_{\text{out}}}{2A_{\text{in}}}, \tag{A10}$$

but $2\pi\delta(x-a) = \int d\omega e^{i\omega(x-a)}$, thus we have

$$\Psi = \frac{1}{2}\delta(x - x_0 - t) + \frac{1}{4\pi} \int d\omega e^{-i\omega(x + x_0 - t)} \frac{-iV_0}{\omega + iV_0}.$$

The remaining integral is performed in the complex plane, by closing the contour of integration as shown in Fig. 18. There are no branch cuts. For $t < x_0 + x$, we can close the contour in the upper half of the complex plane. Since there are no poles in the upper-half plane, the integral vanishes and we get

$$\Psi = \frac{1}{2}\delta(x - x_0 - t), \qquad t < x_0 + x.$$
 (A11)

This is the prompt response: in flat space, signals travel at the speed of light (here c=1) on the light cone. Thus, the

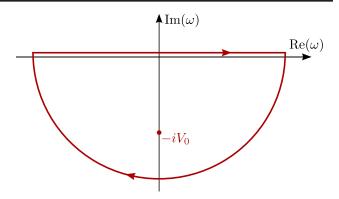


FIG. 18. Integration contour to get the time response in the presence of a δ -function potential and localized initial data.

contribution above corresponds to a signal emitted from the source at x_0 and reaching the observer at $t = x - x_0$.

For $t > x_0 + x$, we must close the contour in the lower half of the complex plane. Taking the residue at the pole, we find

$$\Psi = \frac{1}{2}\delta(x - x_0 - t) + \frac{V_0}{2}e^{-V_0[t - (x + x_0)]}, \qquad t > x_0 + x.$$
(A12)

Note that the second term peaks at $t = x_0 + x$. It corresponds to the initial burst traveling to the potential barrier (which takes a time x_0) and then back to the observer at x, emerging after interaction with the potential barrier. The signal is exponentially suppressed afterwards at a rate determined by the single, purely damped QNM frequency.

In the simple scenario above we assumed sharp, localized initial data. The effect of a broadened pulse can be understood by adding an extra Dirac delta, so that the source $S = \delta(x - x_0) + \delta(x - x_1)$ with $x_1 > x_0$. It is now easy to see that

$$\Psi = \frac{1}{2}\delta(x - x_0 - t) + \frac{1}{2}\delta(x - x_1 - t), \qquad t < x_0 + x$$
(A13)

$$\Psi = \frac{1}{2}\delta(x - x_0 - t) + \frac{1}{2}\delta(x - x_1 - t) + \frac{V_0}{2}e^{-V_0(t - (x + x_0))},$$

$$x_1 + x > t > x_0 + x$$
(A14)

$$\begin{split} \Psi &= \frac{1}{2} \delta(x - x_0 - t) + \frac{1}{2} \delta(x - x_1 - t) \\ &+ \frac{V_0}{2} e^{-V_0(t - x)} (e^{V_0 x_0} + e^{V_0 x_1}), \quad t > x_1 + x \end{split} \tag{A15}$$

We see then two "ringdown" stages of different amplitudes, driven of course by the "extended" initial data. The implication is clear—a constant-in-time QNM amplitude

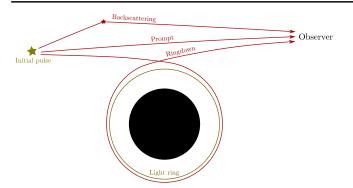


FIG. 19. Schematic representation of the radiation emitted by sources moving around BHs and their path towards faraway observers. The yellow circle represents the light ring.

description is not valid for generic initial data, even when we consider conservative systems.

In curved spacetimes, this picture is further complicated by the absence of propagation on the light cone; backscattering off the spacetime curvature produces tails, beyond the effects already described. A depiction of what happens when a source is emitting close to a BH is shown in Fig. 19. There is direct, prompt emission from the source to the observer, the analog of propagation on the light cone in Minkowski. Concurrent with this process, waves are continuously backscattered off the spacetime curvature, a low-frequency effect that is also present for massive fields or for spacetimes of odd dimensionality (see e.g. [190,191]). In addition, light is trapped close to the potential barrier (i.e., close to the BH

light ring) and continuously reemitted. The sum of all these effects produces the signal seen by an observer. Once the source crosses the BH event horizon, direct emission ceases to exist, but reemission by the light ring and backscattering can still occur.

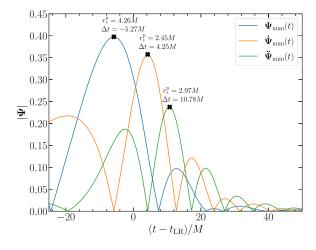
There are two important lessons to be learned from this simple toy model: (1) the signal is never expected to be described by a pure ringdown waveform because other components (including prompt emission and backscattering) will always be present; and (2) even in a Minkowski background and in the linear regime, the "ringdown starting time" is not a well-defined quantity.

We further clarify this point below by considering a curved background. The relative importance of the different emission components, and their effect in searching for damped sinusoids, is explored in more detail in Sec. III of the main text with a toy model.

2. When is ringdown excited?

In a BH spacetime the situation becomes even murkier. To stress how nontrivial it is to define the ringdown starting time (or the "ringdown stage," more in general), we repeat the classic Gedankenexperiment first performed by Vishveshwara [4].

We let a Gaussian wave packet scatter off a nonrotating BH, and we compare the location of the peaks of the signal with the respective position of the center of the Gaussian at the emission of the direct signal, i.e., at time $t_{\text{peak}} = t + r_{\text{ext}}^{\text{ext}} - r_{\text{p}}^{\text{p}}$. Here $r_{\text{ext}}^{\text{ext}}$ is the extraction location



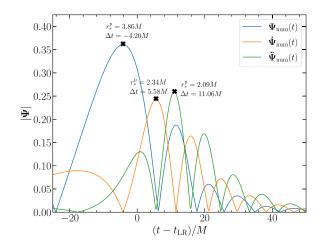


FIG. 20. Left panel: Time evolution of the $\ell=2$ mode for the Zerilli function Ψ (and its first and second time derivatives) extracted at null infinity, for initial data given by Eqs. (A16) and (A17) with $r_*^{(0)}=50M$ and $\sigma=5.0M$. Here, t_{LR} corresponds to the instant when the Gaussian crosses the light ring, located at $r_{LR}=3M$. The first burst occurs slightly before t_{LR} , when part of the front of the wave packet reaches the peak of the gravitational potential (which is close to the light ring), leading to the direct emission of radiation. The following peaks of Ψ occur at time intervals close to half of the light-ring period $T_{LR}/2 \sim 16M$ and correspond to the ringdown phase, where trapped waves at the light ring leak out both to infinity and into the BH. Note that the peaks in the time derivatives of the Zerilli function occur later. Right panel: same, but for $r_*^{(0)}=50M$ and $\sigma=0.5M$. Now the first burst of radiation appears later, since the Gaussian is narrower. The ringdown is more excited, i.e., the relative magnitude of the second peak of Ψ with respect to the first is larger. Thus, wide Gaussians suppress the ringdown.

on the numerical grid (in terms of the Schwarzschild tortoise coordinate), $r_*^{(0)}$ is the initial (tortoise) radius of the Gaussian center, and $r_*^p = r_*^{(0)} - t$ is the radius of the center of the Gaussian at time t. This can be achieved by evolving the Zerilli equation [2] with initial data

$$\Psi(t=0,r) = 0, (A16)$$

$$\partial_t \Psi(t=0,r) = \exp\left[-(r-r_*^{(0)})^2/(2\sigma^2)\right].$$
 (A17)

The resulting time evolution of the Zerilli function and of its first two time derivatives are shown in Fig. 20 for a quadrupolar ($\ell = 2$) perturbation. The two panels refer to Gaussian wave packets with different widths ($\sigma = 5.0M$ on the left, and $\sigma = 0.5M$ on the right). It is apparent that different physical quantities have different peak times. The first peak in Ψ (as extracted at null infinity) occurs slightly before the Gaussian crosses the light ring, and it corresponds to direct radiation excited from the peak of the gravitational potential. The broader Gaussian excites the burst first, because the "leading edge" of the wave packet reaches the peak first when compared to the narrower Gaussian. However, the zero of $\ddot{\Psi}$, which roughly corresponds to the peak of energy emission because $\dot{E} \propto |\dot{\Psi}|^2$. occurs shortly after the Gaussian has crossed the light ring. This can be interpreted in terms of the leakage of waves that are quasitrapped at the light ring on unstable circular orbits, i.e., as "the beginning of the ringdown.". This interpretation is supported by the time difference between the peaks of Ψ, which is comparable to half the light-ring period $T_{\rm LR}/2 \sim 16M$. Note also that for the larger Gaussian, the ringdown phase is suppressed and the second peak of Ψ has smaller relative magnitude with respect to the first peak of direct radiation, due to the interaction of the infalling wave packet with radiation escaping to infinity.

In Fig. 21 we show that the results are similar when we consider the radial infall of a point particle from rest at infinity, rather than the scattering of a Gaussian wave packet.

However, it would be wrong to read too much into these data. The light ring is constantly being hit with forward-beamed radiation, thus GWs from the light ring are in principle present as soon as the radiation has time to interact with it. The intensity of these waves becomes larger when the particle gets closer, but there is some inherent delay time, since the waves are trapped in quasicircular motion. Therefore there is a large uncertainty in trying to associate the peak of the waveform, or of the energy flux, to a particular particle position. For example, in the point-particle case, it is far fetched to claim that the flux peak occurs when the particle is close to the horizon $(r_*^p = 2.07M)$ because of the particle being close to the horizon: the peak is at that location because the particle has traveled some distance before the GWs it emitted close to

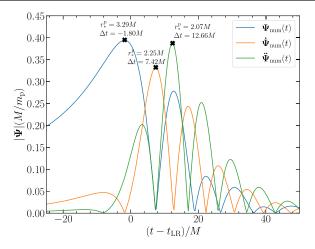


FIG. 21. Same as Fig. 20, but instead of a Gaussian wave packet we let a point particle fall radially onto the BH from rest at infinity. The waveforms are scaled by the ratio between the particle mass m_p and the BH mass M. Results are qualitatively similar to the Gaussian case.

the light ring start leaking out. For much more detailed studies of ringdown excitation in the point-particle limit, see e.g., Refs. [192–203].

The punchline of this discussion is that even for linear perturbations of a Schwarzschild spacetime, the ringdown starting time is not a well-defined quantity. It depends on which quantity we monitor (differing e.g., for the wave amplitude and energy flux), and there is no mathematical or physical basis to claim a well-defined instant as "the" ringdown starting time.

APPENDIX B: IMPACT OF NUMERICAL RELATIVITY ERRORS ON RINGDOWN AMPLITUDE FITS

Two important sources of error in the SXS NR waveforms analyzed here are the finite numerical resolution of the simulations (henceforth resolution error) and the GW extraction procedure, which involves extrapolation to infinity of the waveforms computed at finite radii (henceforth extrapolation error). Our goal in this appendix is to understand to what extent these errors affect our ringdown amplitude fits (see Ref. [151] for a more detailed discussion of numerical errors).

The resolution in the SpEC code depends on the number of points and subdomains in the discrete spacetime grid. These are controlled by a tolerance parameter in the adaptive mesh refinement (AMR) algorithm, which modifies the number of points and subdomains when the local measurement of the error is larger than the tolerance. In converged NR simulations the discretization error decreases when lowering the AMR tolerance parameter, which corresponds to increasing the resolution level. Hence, a common way to gauge the impact of the

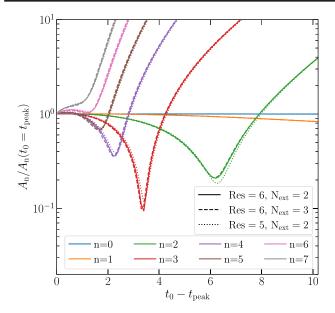


FIG. 22. Same as the bottom-right panel with N=7 in Fig. 6, but using the NR waveform with (i) the second-highest resolution level Lev = 5 and $N_{\rm ext}=2$ (dotted lines), and (ii) the highest resolution level Lev = 6 but with $N_{\rm ext}=3$ (dash-dotted lines). Solid lines correspond to the fits in Fig. 6, which use the reference NR simulation with resolution level Lev = 6 and $N_{\rm ext}=2$.

resolution error in NR simulations is to compare the two highest-resolution levels of the same simulation.

The extrapolation error is caused by the fact that the NR waveforms are not computed directly at future null infinity \mathcal{I}^+ . Hence they have to be extrapolated to \mathcal{I}^+ using a polynomial fit of the strain, produced using its value computed at different extrapolation radii. Since polynomials of different order $N_{\rm ext}$ will give slightly different extrapolations of the waveform at \mathcal{I}^+ , it is common to estimate the impact of the extrapolation error by comparing the waveform extrapolated using polynomials with different values of $N_{\rm ext}$. The reference extrapolation order for the merger and ringdown part of the waveform is $N_{\rm ext}=2$ (see Sec. 2.4.1 of Ref. [151]). This waveform is typically compared with the same waveform extrapolated using an $N_{\rm ext}=3$ polynomial to estimate the extrapolation error.

We want to assess the effect of these two sources of error on the fits of the postmerger NR waveform performed in Secs. IV and V. In particular, we are interested in understanding whether the variation of the amplitude of the modes when fitting the waveform starting at different initial times t_0 could be due to this error.

In Fig. 6 of the main text we use the reference NR waveform at the highest-resolution level (Lev = 6) with $N_{\rm ext}=2$. In Fig. 22 we perform again the same fit of the

postmerger NR waveform using a model with N=7 overtones, but this time we compare with (i) the NR waveform with the second highest-resolution level (Lev = 5) and $N_{\rm ext}=2$, as well as (ii) the NR waveform with Lev = 6, but $N_{\rm ext}=3$. We find that using NR waveforms with different resolution levels and different extrapolation orders makes almost no visible difference in the fits. These sources of error in the NR simulations cannot explain the variation of the fitted amplitudes of the modes when varying t_0 .

APPENDIX C: EFFECT OF A SPURIOUS CONSTANT

In Sec. V we pointed out that in the $Q_{3,1}^{(m)}$ model, the free mode is merely fitting the late-merger GW frequency at early times $t_0 - t_{\text{peak}} < 15M$, while at late times $(t_0 - t_{\text{peak}} > 15M)$ it is unphysical, and it is likely fitting noise. Here we justify the second statement.

The waveform SXS:BBH:0305 used in this study contains a spurious constant (see also Ref. [207]). To determine this constant we subtract the fundamental mode $A_{220}e^{-i(\omega_{220}t+\phi_{220})}$, which can be very accurately determined, from the late-time ringdown, with the result

$$c_n = -(0.8 + 1.8 i) \times 10^{-5}.$$
 (C1)

In this section we gauge the impact of this term on our fits, using waveforms where the constant has not been perfectly subtracted,

$$h_{\lambda} = h_{22}^{\text{SXS}} - \lambda c_n. \tag{C2}$$

Here, $h_{22}^{\rm SXS}$ is the original extrapolated waveform provided by the SXS collaboration, and we have introduced a continuous parameter λ such that $\lambda = 0$ gives the waveform from which the spurious constant c_n has not been subtracted, while $\lambda = 1$ corresponds to perfect subtraction.

In Fig. 23, we plot the waveform amplitude for different values of λ . The original waveform ($\lambda = 0$,

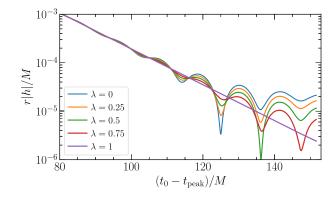


FIG. 23. $r|h_{\lambda}|/M$: Amplitudes of SXS:BBH:0305 when the spurious constant is not subtracted perfectly [see Eq. (C2)].

 $^{^9}$ NR waveforms can be extracted at \mathcal{I}^+ using Cauchy characteristic extraction [204–206]. Unfortunately, at the moment the availability of these NR simulations is limited.

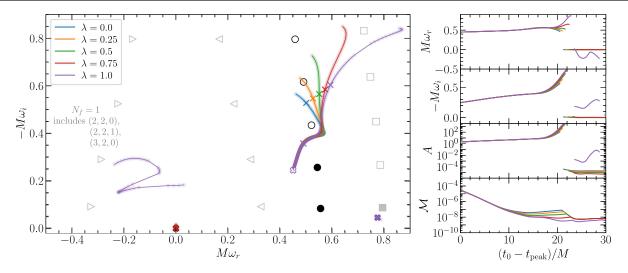


FIG. 24. Fitting h_{λ} with the $Q_{2,1}^{(m)}$ model: one free mode in addition to (2, 2, 0), (2, 2, 1), (3, 2, 0) (which are all marked by filled symbols in the left panel). Left panel: extracted frequencies $(M\omega_r, -M\omega_i)$ for $n_f = 0$, with t_0 ranging from t_{peak} (empty cross) to $t_{\text{peak}} + 30M$ (filled cross). We also mark $t_{\text{peak}} + 10M$ and $t_{\text{peak}} + 20M$ by crosses. We denote known QNMs as follows: (2, 2, n) (black circles), (2, -2, n) (gray left-pointing triangles), $(2, 2, n)_R$ (gray right-pointing triangles), (3, 2, n) (gray squares). Right panels: evolution of $M\omega_r$, $-M\omega_i$, A, and A as functions t_0 .

blue line) does not show an exponential decay at late times, when c_n dominates over the fundamental mode. Once we extract c_n perfectly ($\lambda = 1$, purple line), we recover the "pure" exponential decay due to the (2, 2, 0) mode.

Next, we fit h_{λ} with the $Q_{2,1}^{(m)}$ model. Recall this means one free mode in addition to the (2, 2, 0), (2, 2, 1), (3, 2, 0) fixed modes. In Fig. 24 we plot the extracted complex frequency, the amplitude of the free mode and the

mismatch, as usual. The early-time behavior is the same for all waveforms, being controlled by the late-merger GW frequency. However the late-time behavior is very sensitive to λ . For $t_0 - t_{\rm peak} = 15M-20M$, the imaginary part of the free mode increases for all λ , but the trajectory of the real part of the frequency in the complex plane depends on λ . Whenever $\lambda \neq 1$ we find that $M\omega_r$ drops to zero for $t_0 - t_{\rm peak} \gtrsim 20M$: the free mode is fitting a spurious constant with value $(1 - \lambda) c_n$. This is a consistency check

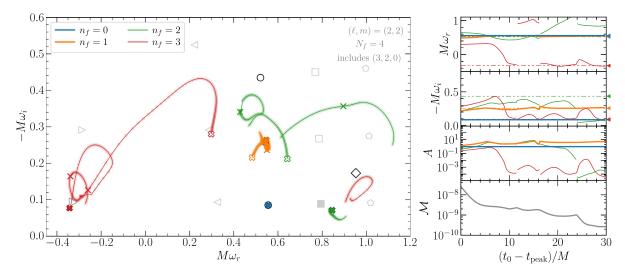


FIG. 25. Fits with the $Q_3^{(f,m)}$ model, i.e., four free modes $(N_f=4)$ in addition to (3,2,0) (marked by a filled symbol in the left panel). Left panel: extracted frequencies $(M\omega_r, -M\omega_i)$ for $n_f=0,1,2,3$, with t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross). We also mark $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$ by crosses. We denote known QNMs as follows: (2,2,n) (black circles), (2,-2,n) (gray left-pointing triangles), $(2,2,n)_R$ (gray right-pointing triangles), (3,2,n) (gray squares), (4,2,n) (gray pentagons). Right panels: evolution of $M\omega_r, -M\omega_i, A$, and M as functions of t_0 . We also plot $M\omega_{220}, M\omega_{221}, M\omega_{222}, M\omega_{22n,R} = -M\omega_{2-22}^*$ using blue, orange, green and red dot-dashed horizontal lines, and triangle markers in the same colors.

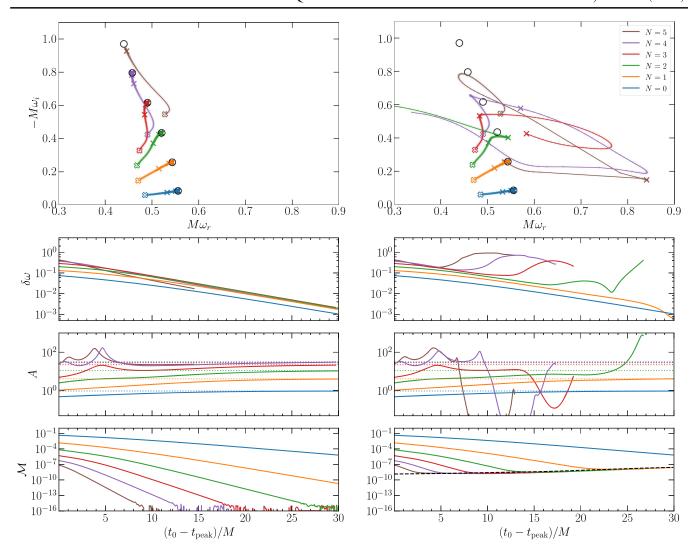


FIG. 26. Fits of the $Q_7(t)$ (left column) and $Q_7^e(t)$ (right column) toy waveforms with model $Q_{N,1}$. First row: evolution of the best-fit frequency of the highest overtone with t_0 . Second row: difference between the fitted free frequency and the predicted value for the corresponding overtone, as defined in Eq. (16). Third row: amplitude of the free overtone in each fit. The injected amplitudes of the toy waveforms are indicated by horizontal dotted lines. Bottom row: mismatch between the best-fit waveform and the actual waveform for each fit. In the bottom-right panel we also show with a black dashed line the mismatch between $Q_7(t)$ and $Q_7^e(t)$, i.e., the approximate mismatch induced by the injected noise floor. As in Fig. 2 we remove spurious data points. This is why some of the curves end prematurely.

of our fitting procedure, and shows that we can extract even very small spurious noise components ($\sim \mathcal{O}(10^{-6})$).

The situation is different for $\lambda=1$, when the spurious noise has been subtracted and the free mode now finds evidence for a retrograde mode with $M\omega_r<0$.

The impact of the spurious constant and the presence of retrograde modes can be better understood by adding multiple free modes. In Fig. 25 we fit the $\lambda=1$ waveform with model $Q_3^{(f,m)}$; $N_f=4$ free modes in addition to the fixed (3,2,0) mode. We observe that the fundamental mode can be extracted nearly perfectly $\delta\omega\sim\mathcal{O}(10^{-5})$ and the $n_f=1$ mode converges to the first overtone, with $\delta\omega\sim\mathcal{O}(10^{-2})$. However the behavior of the $n_f=2,3$

modes is very sensitive to numerical noise. The mismatch when fitting with $Q_3^{(f,m)}$ is $\mathcal{M} \sim \mathcal{O}(10^{-9})$, which is much smaller than differences between waveforms computed at different resolutions, further indicating the possibility than one or two free modes could be simply fitting numerical noise. However, the real frequency $M\omega_r$ of the $n_f=3$ mode is negative and indicates the presence of a retrograde mode (shown by a down-pointing triangle) around $t_0-t_{\rm peak}=10M$, with an amplitude of $\mathcal{O}(10^{-4})$. We also see that none of the free modes coincide with higher overtones (n>1). Once again, it is not possible to extract the first overtone at $t_0=t_{\rm peak}$. We conclude that the presence of the spurious constant has only a small impact

at very late times, and it does not affect the extraction of overtones near the peak.

APPENDIX D: A TOY MODEL FOR SXS:BBH:0305

In Sec. III we examined toy model waveforms to argue that the overtone model may not be a good description of the ringdown as early as the peak even in linear theory. While we expect the same conclusions to hold for NR waveforms, the results in Sec. III should not be carelessly used as a benchmark, because there are some important (and not obvious) differences between the linear and NR waveforms. Firstly, the linear model we explored refers to the Zerilli and Regge-Wheeler functions. A linear combination of these functions yields the Newman-Penrose scalar Ψ , while the NR waveforms refer to the strain h. The two quantities are connected by a second derivative in time, $\Psi \sim \ddot{h}$, meaning that the peak of Ψ and h may differ by a few M (see e.g., the discussion in Appendix A). Secondly, the linear models we considered are strictly real, because the real and imaginary parts of the Regge-Wheeler equation [Eq. (7)] decouple and we specified real initial conditions, while the complex NR waveforms must be found as a linear superposition of complex damped exponentials as in Eq. (10). This technical difference could affect the quality of the fits. In fact, in this section we will see that fits for the NR waveforms often give smoother results than the linear waveforms across different starting times. This is because the full phase evolution of the waveform is encoded in the complex waveform, and this typically improves the quality of fits in terms of complex exponentials which are essentially spinning phasors (with an exponentially decaying amplitude) in the complex plane. Thirdly, while the Price power-law tail is apparent in the linear waveform, to our knowledge there is no conclusive evidence of its presence in the NR waveforms. While we could demonstrate that the exponential blow up of the QNM amplitudes at late times for the linear waveforms is not driven by numerical noise, it is uncertain whether the same conclusion applies to NR waveforms.

To explore the effect of these differences, in the spirit of Sec. III we construct two toy waveforms,

$$Q_7(t) \equiv \sum_{n=0}^7 A_n e^{-i[\omega_n(t-t_{\text{peak}}+\phi_n]}, \qquad (D1)$$

$$Q_7^e(t) \equiv Q_7(t) + e(t), \tag{D2}$$

where $Q_7(t)$ is a clean combination of seven overtones, and $Q_7^e(t)$ contains in addition the estimated numerical

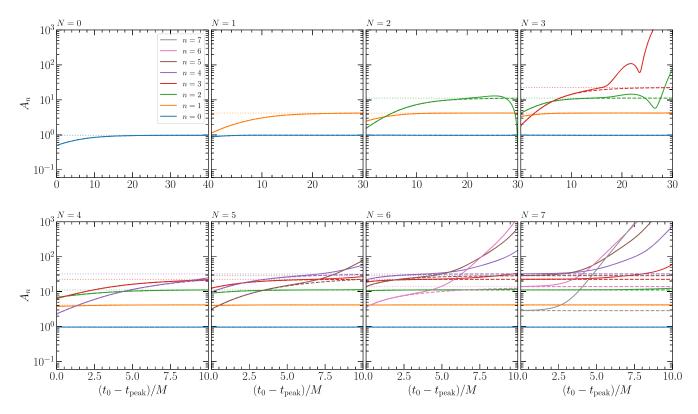


FIG. 27. The amplitude of the overtones for the $Q_7(t)$ (dashed) and $Q_7^e(t)$ (solid) toy waveforms, similar to Fig. 4. Each panel correspond to a fit with a different number of overtones in the fitting model. All mode frequencies are fixed to their predicted values. Faint, dotted horizontal lines are the injected values of the amplitude for each mode. Note that the time range of the plots is different across different panels.

error e(t). These toy waveforms are constructed to mimic the $\ell=m=2$ multipole ringdown waveform of the SXS: BBH:0305 simulation (resolution Lev = 6, extrapolation order N=2: see Appendix B), so we use the first seven overtone frequencies of its remnant BH, and the injected values of A_n and ϕ_n are chosen to be equal to those obtained by fitting the SXS:BBH:0305 waveform starting at the peak with 7 overtones. The curly symbol $Q_7(t)$ is a reminder that, contrary to the real waveforms in Sec. III, these waveforms are complex. The error is approximated as the difference between the Lev = 6 and Lev = 5 waveforms minimized over a relative time and phase difference.

We fit these toy waveforms by a linear combination of complex QNMs. We first repeat the free-frequency search of Sec. III by adding more overtones to our fitting model while leaving only the highest overtone frequency free. As shown in Fig. 26, the free mode converges to the expected overtone frequencies of the clean toy waveform $Q_7(t)$ at least for $N \leq 5$. Once we add noise in $\mathcal{Q}_{7}^{e}(t)$, the fit fails to converge unless $N \leq 1$. This proves that even if there were no nonlinearities near the peak and the amplitudes of the overtones were constant starting at the peak, a subdominant contamination to the waveform could affect our ability to extract the correct overtone frequencies. However, we should be able to recover the first overtone at late enough starting time, contrary to the linear waveform case, where we cannot extract any overtone unless we use a very specific starting time. In fact, this is just what we observe in Fig. 11 for the SXS:BBH:0305 waveform.

Next, we repeat the fits with fixed frequencies, also as in Sec. III. Figure 27 shows the amplitude of the overtones as a function of t_0 for both $Q_7(t)$ (dashed lines) and $Q_7^e(t)$ (solid lines). Adding more overtones stabilizes the lower overtone amplitudes. For N = 7 there is a brief period of time where all overtone amplitudes are constant even near the peak, even for the $Q_7^e(t)$ toy waveform. While the agnostic (free-frequency) mode search would fail for $N \ge 3$, if the waveform could be modeled as QNMs modulo small contamination, we should expect the amplitudes of the overtones to be flat for a brief moment. For example, with an N = 7 fit model (bottom right panel of Fig. 27), the amplitudes of the n = 2 and 3 overtones are approximately constant over a time range $\sim 10M$ and $\sim 5M$, respectively. The corresponding range of flatness (if any) in Fig. 6 is significantly shorter. Therefore the numerical error alone cannot explain why we fail to consistently model SXS:BBH:0305 as a combination of overtones starting at the peak.

APPENDIX E: AN UNPHYSICAL HYBRID WAVEFORM

In Sec. IV we show that the overtones found when fitting the early postpeak BBH waveforms are unphysical. In this appendix, we confirm this finding by producing an unphysical postpeak BBH waveform, and showing that the overtones can still fit it with similar accuracy.

We produce the unphysical postpeak BBH waveform by hybridizing the $\ell=m=2$ mode of our prototypical NR waveform SXS:BBH:0305 (henceforth h_2) with another waveform in the catalog, SXS:BBH:0220 (henceforth h_1). The latter represents an equal-mass BH binary with dimensionless spins $\chi_1=-0.4$ and $\chi_2=-0.8$ antialigned with the orbital angular momentum. We first align the two waveforms such that t=0M correspond to the peak of their $\ell=m=2$ mode. Then we construct the hybrid waveform by smoothly blending their amplitudes A_1 and A_2 and their frequencies ω_1 and ω_2 as follows:

$$A_{\text{hyb}}(t) = A_1(t)W(t) + A_2(t)(1 - W(t)),$$

$$\omega_{\text{hyb}}(t) = \omega_1(t)W(t) + \omega_2(t)(1 - W(t)),$$
 (E1)

where W(t) is a window function defined as

$$W(t) = \begin{cases} 1 & t \le t_1 \\ S\left(\frac{t_2 - t_1}{t - t_1} + \frac{t_2 - t_1}{t - t_2}\right) & t_1 < t < t_2 \\ 0 & t \ge t_2, \end{cases}$$
 (E2)

where S(t) is the logistic function, $t_1 = t_{\rm peak}$ and $t_2 = t_{\rm peak} + 20M$. With this definition, the hybrid waveform exactly matches $h_1(h_2)$ for $t \le t_{\rm peak}$ ($t \ge t_{\rm peak} + 20M$), and it is an unphysical combination of the two in the region $t_1 < t < t_2$.

We fit the hybrid waveform by the Q_N model, with QNM frequencies set to those of the BH remnant of the SXS: BBH:0305 waveform. In the upper panel of Fig. 28, we show the mismatch \mathcal{M}_h between the fits and the hybrid waveform as a function of $t_0 - t_{\text{peak}}$, where t_{peak} is the peak of the hybrid waveform. For comparison, we also show the mismatch \mathcal{M} between SXS:BBH:0305 and the Q_N model fits. Using more and more overtones improves the fit of the hybrid waveform in the unphysical region ($t_{\text{peak}} < t_0 < t_{\text{peak}} + 20M$) as much as it does it for the SXS:BBH:0305 waveform. In fact, in some ranges of t_0 the fit of the hybrid waveform in the unphysical region has an even smaller mismatch than the fit of the SXS:BBH:0305 waveform. This is clear from the second panel of the figure, where we show the ratio $\mathcal{M}/\mathcal{M}_h$.

We find that the overtones typically add in counterphase in the fit of the hybrid waveform, just as they do when fitting SXS:BBH:0305. The extension of the overtone model up to the peak requires very large, presumably unphysical amplitudes, which however can be fine-tuned to cancel out. Modes in counterphase are very effective at (pathologically) canceling each other. This is precisely what is required to overfit the early postmerger signal.

Interestingly, including overtones in the fit of the postpeak hybrid waveform also improves the measurement

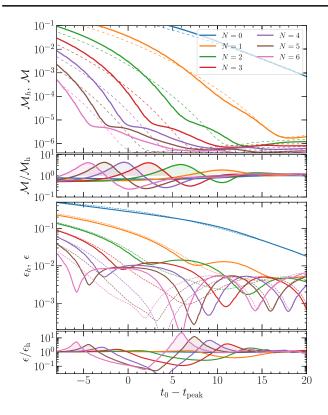


FIG. 28. Mismatch and mass/spin errors (first and third panel from the top, respectively) for the hybrid waveform (solid line) and SXS:BBH:0305 (dashed line). We also plot the ratio of the corresponding quantities for the hybrid waveform (solid line) and SXS:BBH:0305 (dashed line), shading the region where the overtones fit the hybrid waveform better.

of the remnant spin χ_f and mass M_f . We prove this by repeating the fit of this waveform with models Q_N with free χ_f and M_f (in addition to the free amplitudes and phases of the modes). We determine the deviation of the fitted χ_f and M_f from their asymptotic values using the quantity ϵ_h , defined in analogy with Eq. (19). Smaller values of ϵ_h indicate a better measurement of χ_f and M_f .

In the third panel of Fig. 28, we show ϵ_h as a function of $t_0 - t_{\rm peak}$. For comparison we also plot the value of ϵ obtained when fitting the Q_N models against the SXS: BBH:0305 waveform. The results confirm what we found by computing the mismatch i.e., ϵ_h gets smaller when adding a larger number of overtones, even when we start the fit at values of t_0 for which the hybrid waveform is unphysical. The bottom panel of Fig. 28 shows that ϵ_h is even smaller than ϵ in some ranges of t_0 . The fact that using the overtones allows us to improve the measurement of χ_f and M_f even for the unphysical hybrid waveform is clear evidence that their role is to match whichever early postpeak waveform, allowing the dominant mode to correctly fit the late postpeak waveform, which really carries information about the remnant BH properties.

APPENDIX F: HIGHER MULTIPOLES

In Sec. V D we focused on the (4, 4) subdominant multipole, noting that (quite remarkably) the nonlinear mode is easier to extract than the first linear overtone.

In this section we further analyze the (3, 3), (3, 2), and (2, 1) subdominant multipoles to understand which modes

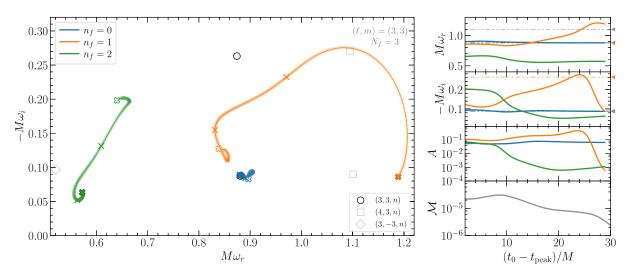


FIG. 29. Fits of the (3, 3) multipole with the $Q_2^{(f)}$ model. Left panel: the extracted frequencies $(M\omega_r, -M\omega_i)$ for $n_f=0, 1, 2$, with t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross); other crosses correspond to times $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$. We also mark the QNMs (3,3,n) (black circles), (3,-3,n) (gray diamonds) and (4,3,n) (gray squares), for reference. Right panels: $M\omega_r, M\omega_i$, A and M as functions of t_0 . Horizontal gray (orange) dot-dot-dashed lines and triangle markers correspond to the expected frequencies for $M\omega_{430}$ ($M\omega_{331}$).

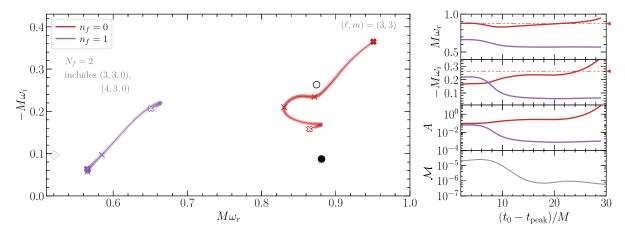


FIG. 30. Fits of the (3, 3) multipole using the $Q_{2,2}^{(m)}$ model, i.e., two free modes $(N_f=2)$ in addition to two additional modes with frequencies fixed to the (3, 3, 0) and (4, 3, 0) values. Left panel: the extracted frequencies $(M\omega_r, -M\omega_i)$ for $n_f=0$, 1 and $N_f=2$, for t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross); other crosses correspond to times $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$. We mark the QNMs (3,3,n) (black circles) and (3,-3,n) (gray diamonds). Right panels: $M\omega_r, M\omega_i, A$ and M as functions of t_0 . The horizontal red dot-dot-dashed line and triangle marker corresponds to the expected frequency for $M\omega_{331}$.

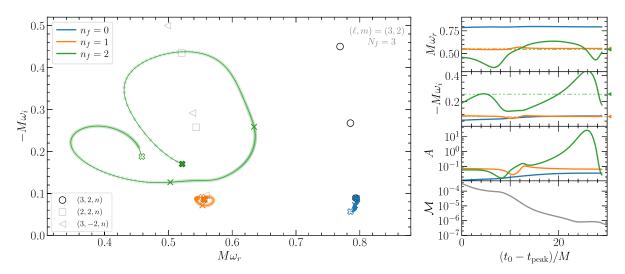


FIG. 31. Fits of the (3, 2) multipole with the $Q_2^{(f)}$ model. Left panel: the extracted frequencies $(M\omega_r, -M\omega_i)$ for $n_f = 0, 1, 2$ for t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak} + 30M$ (filled cross); other crosses correspond to times $t_{\rm peak} + 10M$ and $t_{\rm peak} + 20M$. We mark the QNMs (3, 2, n) (black circles), (3, -2, n) (gray triangle), and (2, 2, n) (gray squares). Note that the marker for the (2, 2, 0) mode almost exactly overlaps with that of the (3, -2, 0) mode. Right panels: $M\omega_r, -M\omega_i, A$ and M as functions of t_0 . The horizontal orange (green) dot-dot-dashed lines and triangle markers correspond to the expected frequencies for $M\omega_{220}$ ($M\omega_{221}$).

can be extracted from the numerical waveforms, and whether observations that apply to the (2, 2) mode generalize to other multipoles of the radiation.

1.
$$(\mathcal{C}, m) = (3, 3)$$

We fit the (3, 3) multipole with the $Q_2^{(f)}$ model. In Fig. 29 we plot the extracted frequencies, amplitudes, and mismatches. The fundamental mode (3, 3, 0) (blue lines) can be extracted with $\delta\omega < 0.01$ as long as $t_0 - t_{\rm peak} \gtrsim 10M$. The second free mode $(n_f = 1, \text{ orange line})$ shows a tendency to approach the (4, 3, 0) mode-mixing QNM, and it still

seems to be changing in frequency and amplitude after $t_0 - t_{\rm peak} > 30M$. The third free mode ($n_f = 2$, green line) moves between two QNM frequencies: one (at early times) with large $|M\omega_i|$ and large amplitude, and one at late times with small $|M\omega_i|$ and small amplitude.

As shown in Fig. 30, this mode is present even when we fit with the $Q_2^{(f,m)}$ model, i.e., we add $N_f=2$ free modes along with two fixed frequencies corresponding to the modes (3,3,0) and (4,3,0). This model also finds a mode moving between two QNM frequencies; one with $A \simeq 0.053$ and $M\omega \simeq 0.65-0.20i$ at early times $(t_0-t_{\rm peak}<5M)$, and one

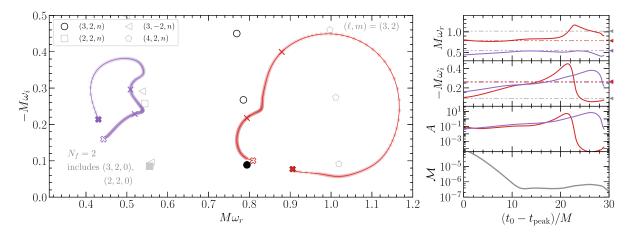


FIG. 32. Fits of the (3,2) multipole using the $Q_{2,2}^{(m)}$ model: two free modes $(N_f=2)$ in addition to (2,2,0) and (3,2,0). Left panel: the extracted frequencies $(M\omega_r, -M\omega_i)$ for $n_f=0$, 1 and $N_f=2$, for t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross); other crosses correspond to times $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$. We mark the QNMs (3,2,n) (black circles), (3,-2,n) (gray triangles), (2,2,n) (gray squares), and (4,2,n) (gray pentagons). Right panels: $M\omega_r, -M\omega_i, A$, and M as functions of t_0 . The horizontal gray, purple, and red dot-dot-dashed lines and gray triangle markers correspond to the expected frequencies for $M\omega_{420}, M\omega_{221}$, and $M\omega_{321}$, respectively.

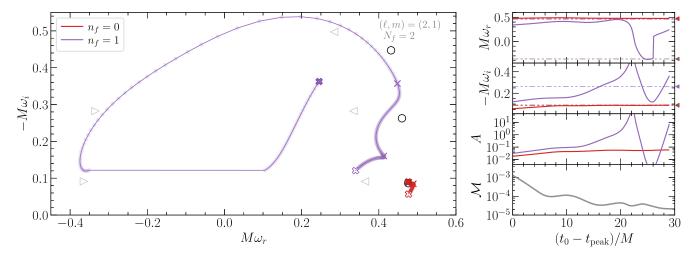


FIG. 33. Fits of the (2, 1) multipole using the $Q_1^{(f)}$ model. Left panel: the extracted frequencies $(M\omega_r, -M\omega_i)$ for $n_f=0, 1$ with t_0 ranging from $t_{\rm peak}$ (empty cross) to $t_{\rm peak}+30M$ (filled cross); other crosses are at $t_{\rm peak}+10M$ and $t_{\rm peak}+20M$. We mark the known QNM frequencies as follows: (2,1,n) (black circles), (2,-1,n) (gray left-pointing triangles), $(2,1,n)_R$ (gray right-pointing triangles). Right panels: $M\omega_r, M\omega_i, A$, and M as functions of t_0 . Horizontal red, purple and brown dot-dot-dashed lines and triangle markers are the expected frequencies for the modes $M\omega_{210}, M\omega_{211}$, and $-M\omega_{2-11}^*$, respectively.

with $A \simeq 0.0012$ and $M\omega \simeq 0.56 - 0.063i$ at late times $(t_0 - t_{\rm peak} > 15M)$. The free mode with $n_f = 0$ in Fig. 30 finds a frequency close to the (3, 3, 1) mode, with "closest approach" happening at $t_0 - t_{\rm peak} \approx 25M$.

2.
$$(\mathcal{C}, m) = (3, 2)$$

We start by fitting the (3, 2) multipole with the free-frequency $Q_2^{(f)}$ model. We plot the three extracted frequencies, amplitudes and mismatches in Fig. 31. The fundamental mode (3, 2, 0) can be extracted with

 $\delta\omega$ < 0.01 when $t_0-t_{\rm peak}\gtrsim 12M$ (blue line). The spherical-spheroidal mixing contamination from the (2, 2, 0) mode can also be accurately identified, with $\delta\omega\lesssim 10^{-2}$ in the time range we consider [note that in Fig. 31 the (2, 2, 0) mode almost exactly overlaps with the (3, -2, 0) mode, although no evidence is found for this additional mode]. The third free mode circles around the (2, 2, 1) overtone, but never converges to any known QNM frequency value.

In Fig. 32 we try better extract some of the modes by repeating the fit with the $Q_{2,2}^{(m)}$ template, i.e., two

fixed modes [(3, 2, 0) and (2, 2, 0)] and two free modes. The $n_f = 0$ mode (red line) moves towards the (3, 2, 1) overtone, but then drifts towards the (4, 2, 0) QNM due to mode-mixing contamination after $t_0 - t_{\text{peak}} > 25M$. The $n_f = 1$ mode hovers around to the (2, 2, 1) overtone, but never gets really close to it.

In conclusion, mode mixing contamination seems to prevent an accurate recovery of the first overtone with any of the free-frequency fits that we attempted.

- 3. $(\mathcal{C}, m) = (2, 1)$
- In Fig. 33 we plot the two extracted frequencies, amplitudes and mismatches found by fitting the (2, 1) multipole with the $Q_1^{(f)}$ model. As usual, the fundamental mode (2, 1, 0) can be extracted with $\delta\omega < 0.01$ after $t_0 t_{\rm peak} > 5M$. We also find some evidence of the retrograde mode $(2, 1, 0)_R$. If we repeat the fit with the $Q_2^{(f)}$ model the fits are generally very noisy, and there is no robust evidence for the first overtone or for spherical-spheroidal mixing modes.
- [1] T. Regge and J. A. Wheeler, Phys. Rev. 108, 1063 (1957).
- [2] F. J. Zerilli, Phys. Rev. Lett. 24, 737 (1970).
- [3] F. J. Zerilli, Phys. Rev. D 2, 2141 (1970).
- [4] C. V. Vishveshwara, Nature (London) 227, 936 (1970).
- [5] H.-P. Nollert, Classical Quantum Gravity 16, R159 (1999).
- [6] K. D. Kokkotas and B. G. Schmidt, Living Rev. Relativity 2, 2 (1999).
- [7] E. Berti, V. Cardoso, and A.O. Starinets, Classical Quantum Gravity 26, 163001 (2009).
- [8] R. A. Konoplya and A. Zhidenko, Rev. Mod. Phys. 83, 793 (2011).
- [9] W. H. Press, Astrophys. J. Lett. 170, L105 (1971).
- [10] M. Davis, R. Ruffini, W. H. Press, and R. H. Price, Phys. Rev. Lett. 27, 1466 (1971).
- [11] S. Chandrasekhar and S. L. Detweiler, Proc. R. Soc. A **344**, 441 (1975).
- [12] V. Ferrari and B. Mashhoon, Phys. Rev. D 30, 295 (1984).
- [13] B. Mashhoon, Phys. Rev. D 31, 290 (1985).
- [14] H.-J. Blome and B. Mashhoon, Phys. Lett. 100A, 231 (1984).
- [15] B. F. Schutz and C. M. Will, Astrophys. J. Lett. 291, L33 (1985).
- [16] S. A. Teukolsky, Phys. Rev. Lett. 29, 1114 (1972).
- [17] S. A. Teukolsky, Astrophys. J. 185, 635 (1973).
- [18] W. H. Press and S. A. Teukolsky, Astrophys. J. 185, 649 (1973).
- [19] S. A. Teukolsky and W. H. Press, Astrophys. J. **193**, 443 (1974).
- [20] S. L. Detweiler, Proc. R. Soc. A 352, 381 (1977).
- [21] S. L. Detweiler, Astrophys. J. 239, 292 (1980).
- [22] G. Jaffé, Z. Phys. 87, 535 (1934).
- [23] W. G. Baber and H. R. Hassé, Proc. Cambridge Philos. Soc. 31, 564 (1935).
- [24] E. W. Leaver, Proc. R. Soc. A 402, 285 (1985).
- [25] F. Echeverria, Phys. Rev. D 40, 3194 (1989).
- [26] L. S. Finn, Phys. Rev. D 46, 5236 (1992).
- [27] K. S. Thorne, in *Three Hundred Years of Gravitation* (Cambridge University Press, New York, 1987), pp. 330–458.
- [28] S. Hod, Phys. Rev. Lett. 81, 4293 (1998).
- [29] O. Dreyer, Phys. Rev. Lett. **90**, 081301 (2003).

- [30] O. Dreyer, B. J. Kelly, B. Krishnan, L. S. Finn, D. Garrison, and R. Lopez-Aleman, Classical Quantum Gravity 21, 787 (2004).
- [31] K. D. Kokkotas, T. A. Apostolatos, and N. Andersson, Mon. Not. R. Astron. Soc. 320, 307 (2001).
- [32] C. T. Cunningham, R. H. Price, and V. Moncrief, Astrophys. J. 224, 643 (1978).
- [33] V. Ferrari and R. Ruffini, Phys. Lett. 98B, 381 (1981).
- [34] R. F. Stark and T. Piran, Phys. Rev. Lett. 55, 891 (1985); 56, 97(E) (1986).
- [35] E. Berti, K. Yagi, H. Yang, and N. Yunes, Gen. Relativ. Gravit. 50, 49 (2018).
- [36] V. Cardoso and P. Pani, Living Rev. Relativity 22, 4 (2019).
- [37] E. Berti, V. Cardoso, and M. Casals, Phys. Rev. D **73**, 024013 (2006); **73**, 109902(E) (2006).
- [38] E. Berti and A. Klein, Phys. Rev. D 90, 064012 (2014).
- [39] E. W. Leaver, Phys. Rev. D 34, 384 (1986).
- [40] E. Berti, V. Cardoso, and C. M. Will, Phys. Rev. D 73, 064030 (2006).
- [41] Webpages with *Mathematica* notebooks and numerical quasinormal mode tables: https://pages.jh.edu/eberti2/ringdown/; https://centra.tecnico.ulisboa.pt/network/grit/files/; https://paolopani.weebly.com/notebooks.html.
- [42] N. Andersson, Phys. Rev. D 51, 353 (1995).
- [43] N. Andersson, Phys. Rev. D 55, 468 (1997).
- [44] E. Berti and V. Cardoso, Phys. Rev. D 74, 104020 (2006).
- [45] Z. Zhang, E. Berti, and V. Cardoso, Phys. Rev. D 88, 044018 (2013).
- [46] N. Oshita, Phys. Rev. D 104, 124032 (2021).
- [47] M. Lagos and L. Hui, Phys. Rev. D 107, 044040 (2023).
- [48] E. E. Flanagan and S. A. Hughes, Phys. Rev. D **57**, 4535 (1998).
- [49] E. Berti, V. Cardoso, and C. M. Will, AIP Conf. Proc. 848, 687 (2006).
- [50] F. Pretorius, Phys. Rev. Lett. **95**, 121101 (2005).
- [51] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, Phys. Rev. Lett. 96, 111101 (2006).
- [52] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, Phys. Rev. Lett. 96, 111102 (2006).
- [53] A. Buonanno, G. B. Cook, and F. Pretorius, Phys. Rev. D 75, 124018 (2007).

- [54] E. Berti, V. Cardoso, J. A. Gonzalez, U. Sperhake, M. Hannam, S. Husa, and B. Bruegmann, Phys. Rev. D 76, 064034 (2007).
- [55] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **116**, 061102 (2016).
- [56] J. Aasi *et al.* (LIGO Scientific Collaboration), Classical Quantum Gravity **32**, 074001 (2015).
- [57] F. Acernese *et al.* (Virgo Collaboration), Classical Quantum Gravity **32**, 024001 (2015).
- [58] T. Akutsu *et al.* (KAGRA Collaboration), Prog. Theor. Exp. Phys. **2021**, 05A101 (2021).
- [59] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. X **9**, 031040 (2019).
- [60] R. Abbott et al. (LIGO Scientific and Virgo Collaborations), Phys. Rev. X 11, 021053 (2021).
- [61] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), arXiv:2108.01045.
- [62] R. Abbott et al. (LIGO Scientific, Virgo, and KAGRA Collaborations), arXiv:2111.03606 [Phys. Rev. X (to be published)].
- [63] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, Astrophys. J. 872, 195 (2019).
- [64] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, Astrophys. J. 891, 123 (2020).
- [65] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, Astrophys. J. 922, 76 (2021).
- [66] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, Phys. Rev. D 101, 083030 (2020).
- [67] B. Zackay, L. Dai, T. Venumadhav, J. Roulet, and M. Zaldarriaga, Phys. Rev. D 104, 063030 (2021).
- [68] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. D 100, 104036 (2019).
- [69] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. D **103**, 122002 (2021).
- [70] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), arXiv:2112.06861 [Phys. Rev. D (to be published)].
- [71] E. Berti, J. Cardoso, V. Cardoso, and M. Cavaglia, Phys. Rev. D 76, 104044 (2007).
- [72] V. Baibhav, E. Berti, V. Cardoso, and G. Khanna, Phys. Rev. D 97, 044048 (2018).
- [73] I. Kamaretsos, M. Hannam, S. Husa, and B. S. Sathyaprakash, Phys. Rev. D 85, 024018 (2012).
- [74] I. Kamaretsos, M. Hannam, and B. Sathyaprakash, Phys. Rev. Lett. 109, 141102 (2012).
- [75] L. London, D. Shoemaker, and J. Healy, Phys. Rev. D **90**, 124032 (2014); **94**, 069902(E) (2016).
- [76] S. Bhagwat, D. A. Brown, and S. W. Ballmer, Phys. Rev. D **94**, 084024 (2016); **95**, 069906(E) (2017).
- [77] E. Thrane, P.D. Lasky, and Y. Levin, Phys. Rev. D **96**, 102004 (2017).
- [78] L. T. London, Phys. Rev. D 102, 084052 (2020).
- [79] V. Baibhav and E. Berti, Phys. Rev. D 99, 024005 (2019).
- [80] V. Baibhav, E. Berti, and V. Cardoso, Phys. Rev. D 101, 084053 (2020).
- [81] G. B. Cook, Phys. Rev. D 102, 024027 (2020).

- [82] X. Jiménez Forteza, S. Bhagwat, P. Pani, and V. Ferrari, Phys. Rev. D 102, 044053 (2020).
- [83] I. Ota and C. Chirenti, Phys. Rev. D 105, 044015 (2022).
- [84] X. Li, L. Sun, R. K. L. Lo, E. Payne, and Y. Chen, Phys. Rev. D 105, 024016 (2022).
- [85] L. Magaña Zertuche et al., Phys. Rev. D 105, 104015 (2022).
- [86] S. Gossan, J. Veitch, and B. S. Sathyaprakash, Phys. Rev. D 85, 124056 (2012).
- [87] J. Meidam, M. Agathos, C. Van Den Broeck, J. Veitch, and B. S. Sathyaprakash, Phys. Rev. D 90, 064009 (2014).
- [88] M. Punturo et al., Classical Quantum Gravity 27, 194002 (2010).
- [89] S. Bhagwat, M. Cabero, C. D. Capano, B. Krishnan, and D. A. Brown, Phys. Rev. D 102, 024023 (2020).
- [90] S. Borhanian, K. G. Arun, H. P. Pfeiffer, and B. S. Sathyaprakash, Classical Quantum Gravity 37, 065006 (2020).
- [91] M. Cabero, J. Westerweck, C. D. Capano, S. Kumar, A. B. Nielsen, and B. Krishnan, Phys. Rev. D 101, 064044 (2020).
- [92] E. Berti, A. Sesana, E. Barausse, V. Cardoso, and K. Belczynski, Phys. Rev. Lett. 117, 101102 (2016).
- [93] G. Carullo et al., Phys. Rev. D 98, 104020 (2018).
- [94] W. Del Pozzo and A. Nagar, Phys. Rev. D 95, 124034 (2017).
- [95] G. Carullo, W. Del Pozzo, and J. Veitch, Phys. Rev. D 99, 123029 (2019); 100, 089903(E) (2019).
- [96] M. Isi, M. Giesler, W. M. Farr, M. A. Scheel, and S. A. Teukolsky, Phys. Rev. Lett. 123, 111102 (2019).
- [97] G. Carullo, W. D. Pozzo, D. Laghi, M. Isi, and J. Veitch, pyRing: A time-domain ringdown analysis Python package, https://git.ligo.org/lscsoft/pyring.
- [98] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **125**, 101102 (2020).
- [99] R. Abbott et al. (LIGO Scientific and Virgo Collaborations), Astrophys. J. Lett. 900, L13 (2020).
- [100] G. Carullo, D. Laghi, N. K. Johnson-McDaniel, W. Del Pozzo, O. J. C. Dias, M. Godazgar, and J. E. Santos, Phys. Rev. D 105, 062009 (2022).
- [101] G. Carullo, D. Laghi, J. Veitch, and W. Del Pozzo, Phys. Rev. Lett. 126, 161102 (2021).
- [102] D. Laghi, G. Carullo, J. Veitch, and W. Del Pozzo, Classical Quantum Gravity 38, 095005 (2021).
- [103] A. Maselli, P. Pani, L. Gualtieri, and E. Berti, Phys. Rev. D 101, 024043 (2020).
- [104] G. Carullo, Phys. Rev. D 103, 124043 (2021).
- [105] M. Isi and W. M. Farr, arXiv:2107.05609.
- [106] M. Giesler, M. Isi, M. A. Scheel, and S. Teukolsky, Phys. Rev. X 9, 041060 (2019).
- [107] I. Ota and C. Chirenti, Phys. Rev. D 101, 104005 (2020).
- [108] S. Bhagwat, X. J. Forteza, P. Pani, and V. Ferrari, Phys. Rev. D 101, 044033 (2020).
- [109] J. C. Bustillo, P. D. Lasky, and E. Thrane, Phys. Rev. D 103, 024041 (2021).
- [110] M. Okounkova, arXiv:2004.00671.
- [111] P. Mourier, X. Jiménez Forteza, D. Pook-Kolb, B. Krishnan, and E. Schnetter, Phys. Rev. D **103**, 044054 (2021).
- [112] A. Dhani, Phys. Rev. D **103**, 104048 (2021).

- [113] A. Dhani and B. S. Sathyaprakash, arXiv:2107.14195.
- [114] E. Finch and C. J. Moore, Phys. Rev. D **103**, 084048 (2021).
- [115] E. Finch and C. J. Moore, Phys. Rev. D **104**, 123034 (2021).
- [116] X. J. Forteza, S. Bhagwat, S. Kumar, and P. Pani, Phys. Rev. Lett. 130, 021001 (2023).
- [117] X. J. Forteza and P. Mourier, Phys. Rev. D 104, 124072 (2021).
- [118] J. L. Jaramillo and B. Krishnan, arXiv:2211.03405.
- [119] R. Cotesta, G. Carullo, E. Berti, and V. Cardoso, Phys. Rev. Lett. 129, 111102 (2022).
- [120] M. Isi and W. M. Farr, arXiv:2202.02941.
- [121] E. Finch and C. J. Moore, Phys. Rev. D 106, 043005 (2022).
- [122] S. Ma, L. Sun, and Y. Chen, Phys. Rev. D **107**, 084010 (2023).
- [123] S. Ma, L. Sun, and Y. Chen, Phys. Rev. Lett. 130, 141401 (2023).
- [124] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, Phys. Rev. D 84, 124052 (2011).
- [125] T. Damour and A. Nagar, Phys. Rev. D 90, 024054 (2014).
- [126] R. Brito, A. Buonanno, and V. Raymond, Phys. Rev. D 98, 084038 (2018).
- [127] K. Glampedakis, A. D. Johnson, and D. Kennefick, Phys. Rev. D 96, 024036 (2017).
- [128] A. Zenginoglu and G. Khanna, Phys. Rev. X 1, 021017 (2011).
- [129] V. Cardoso, F. Duque, and G. Khanna, Phys. Rev. D 103, L081501 (2021).
- [130] W. Krivan, P. Laguna, P. Papadopoulos, and N. Andersson, Phys. Rev. D 56, 3395 (1997).
- [131] E. Pazos-Avalos and C.O. Lousto, Phys. Rev. D 72, 084022 (2005).
- [132] A. Zenginoğlu, G. Khanna, and L. M. Burko, Gen. Relativ. Gravit. 46, 1672 (2014).
- [133] L. M. Burko and G. Khanna, Phys. Rev. D **89**, 044037 (2014)
- [134] P. A. Sundararajan, G. Khanna, and S. A. Hughes, Phys. Rev. D 76, 104005 (2007).
- [135] V. Cardoso, K. Destounis, F. Duque, R. Panosso Macedo, and A. Maselli, Phys. Rev. Lett. **129**, 241103 (2022).
- [136] P. Virtanen et al., Nat. Methods 17, 261 (2020).
- [137] R. H. Price, Phys. Rev. D 5, 2419 (1972).
- [138] J. L. Jaramillo, R. Panosso Macedo, and L. Al Sheikh, Phys. Rev. X 11, 031003 (2021).
- [139] J. L. Jaramillo, R. Panosso Macedo, and L. A. Sheikh, Phys. Rev. Lett. 128, 211102 (2022).
- [140] M. H.-Y. Cheung, K. Destounis, R. P. Macedo, E. Berti, and V. Cardoso, Phys. Rev. Lett. 128, 111103 (2022).
- [141] E. Berti, V. Cardoso, M. H.-Y. Cheung, F. Di Filippo, F. Duque, P. Martens, and S. Mukohyama, Phys. Rev. D 106, 084011 (2022).
- [142] L. Sberna, P. Bosch, W. E. East, S. R. Green, and L. Lehner, Phys. Rev. D 105, 064046 (2022).
- [143] S. Ma, K. Mitman, L. Sun, N. Deppe, F. Hébert, L. E. Kidder, J. Moxon, W. Throwe, N. L. Vu, and Y. Chen, Phys. Rev. D 106, 084036 (2022).
- [144] K. Mitman et al., Phys. Rev. Lett. 130, 081402 (2023).

- [145] M. H.-Y. Cheung et al., Phys. Rev. Lett. 130, 081401 (2023).
- [146] A. Kehagias, D. Perrone, A. Riotto, and F. Riva, Phys. Rev. D 108, L021501 (2023).
- [147] A. Kehagias and A. Riotto, Phys. Rev. D 108, 021501 (2023).
- [148] E. N. Dorband, E. Berti, P. Diener, E. Schnetter, and M. Tiglio, Phys. Rev. D 74, 084028 (2006).
- [149] E. Berti, V. Cardoso, J. A. Gonzalez, and U. Sperhake, Phys. Rev. D 75, 124017 (2007).
- [150] A. H. Mroue et al., Phys. Rev. Lett. 111, 241104 (2013).
- [151] M. Boyle et al., Classical Quantum Gravity 36, 195006 (2019).
- [152] M. Ruiz, R. Takahashi, M. Alcubierre, and D. Nunez, Gen. Relativ. Gravit. **40**, 2467 (2008).
- [153] S. Bhagwat, M. Okounkova, S. W. Ballmer, D. A. Brown, M. Giesler, M. A. Scheel, and S. A. Teukolsky, Phys. Rev. D 97, 104065 (2018).
- [154] J. Bamber, O. J. Tattersall, K. Clough, and P. G. Ferreira, Phys. Rev. D 103, 124013 (2021).
- [155] F. Dyson, Nature (London) **427**, 297 (2004).
- [156] B. J. Kelly and J. G. Baker, Phys. Rev. D 87, 084004 (2013).
- [157] S. Bernuzzi and A. Nagar, Phys. Rev. D 81, 084056 (2010).
- [158] E. Barausse, A. Buonanno, S. A. Hughes, G. Khanna, S. O'Sullivan, and Y. Pan, Phys. Rev. D 85, 024046 (2012).
- [159] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **116**, 241102 (2016).
- [160] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **116**, 221101 (2016); **121**, 129902(E) (2018).
- [161] E. Maggio, H. O. Silva, A. Buonanno, and A. Ghosh, Phys. Rev. D 108, 024043 (2023).
- [162] M. Isi, W. M. Farr, M. Giesler, M. A. Scheel, and S. A. Teukolsky, Phys. Rev. Lett. 127, 011103 (2021).
- [163] A. Ghosh, N. K. Johnson-Mcdaniel, A. Ghosh, C. K. Mishra, P. Ajith, W. Del Pozzo, C. P. L. Berry, A. B. Nielsen, and L. London, Classical Quantum Gravity 35, 014002 (2018).
- [164] H.-P. Nollert, Phys. Rev. D 53, 4397 (1996).
- [165] E. Barausse, V. Cardoso, and P. Pani, Phys. Rev. D 89, 104059 (2014).
- [166] R. J. Gleiser, C. O. Nicasio, R. H. Price, and J. Pullin, Classical Quantum Gravity 13, L117 (1996).
- [167] M. Campanelli and C. O. Lousto, Phys. Rev. D 59, 124022 (1999).
- [168] D. Brizuela, J. M. Martin-Garcia, and M. Tiglio, Phys. Rev. D 80, 024021 (2009).
- [169] K. Ioka and H. Nakano, Phys. Rev. D **76**, 061503 (2007).
- [170] H. Nakano and K. Ioka, Phys. Rev. D 76, 084007 (2007).
- [171] E. Pazos, D. Brizuela, J. M. Martin-Garcia, and M. Tiglio, Phys. Rev. D 82, 104028 (2010).
- [172] N. Loutrel, J. L. Ripley, E. Giorgi, and F. Pretorius, Phys. Rev. D 103, 104017 (2021).
- [173] J. L. Ripley, N. Loutrel, E. Giorgi, and F. Pretorius, Phys. Rev. D 103, 104018 (2021).
- [174] P. Bosch, S. R. Green, and L. Lehner, Phys. Rev. Lett. 116, 141102 (2016).

- [175] P. Bosch, S. R. Green, L. Lehner, and H. Roussille, Phys. Rev. D 102, 044014 (2020).
- [176] C. D. Capano, M. Cabero, J. Westerweck, J. Abedi, S. Kastha, A. H. Nitz, Y.-F. Wang, A. B. Nielsen, and B. Krishnan, arXiv:2105.05238.
- [177] O. J. Tattersall, P. G. Ferreira, and M. Lagos, Phys. Rev. D 97, 044021 (2018).
- [178] O. J. Tattersall, P. G. Ferreira, and M. Lagos, Phys. Rev. D 97, 084005 (2018).
- [179] V. Cardoso, M. Kimura, A. Maselli, E. Berti, C. F. B. Macedo, and R. McManus, Phys. Rev. D 99, 104077 (2019).
- [180] R. McManus, E. Berti, C. F. B. Macedo, M. Kimura, A. Maselli, and V. Cardoso, Phys. Rev. D 100, 044061 (2019).
- [181] S. H. Völkel, N. Franchini, and E. Barausse, Phys. Rev. D 105, 084046 (2022).
- [182] D. Li, P. Wagle, Y. Chen, and N. Yunes, Phys. Rev. X 13, 021029 (2023).
- [183] A. Hussain and A. Zimmerman, Phys. Rev. D 106, 104018 (2022).
- [184] A. Gupta, S. Datta, S. Kastha, S. Borhanian, K. G. Arun, and B. S. Sathyaprakash, Phys. Rev. Lett. 125, 201101 (2020).
- [185] S. Datta, M. Saleem, K. G. Arun, and B. S. Sathyaprakash, arXiv:2208.07757.
- [186] N. K. Johnson-McDaniel, A. Ghosh, S. Ghonge, M. Saleem, N. V. Krishnendu, and J. A. Clark, Phys. Rev. D 105, 044020 (2022).
- [187] H. O. Silva, A. Ghosh, and A. Buonanno, Phys. Rev. D 107, 044030 (2023).
- [188] P. J. Nee, S. H. Völkel, and H. P. Pfeiffer, Phys. Rev. D 108, 044032 (2023).

- [189] Vitor Cardoso, in ICTP Summer School Lectures 2018, https://indico.ictp.it/event/8317/.
- [190] J. Alsing, E. Berti, C. M. Will, and H. Zaglauer, Phys. Rev. D 85, 064041 (2012).
- [191] V. Cardoso, O. J. C. Dias, and J. P. S. Lemos, Phys. Rev. D 67, 064026 (2003).
- [192] S. Hadar and B. Kol, Phys. Rev. D 84, 044019 (2011).
- [193] S. Hadar, B. Kol, E. Berti, and V. Cardoso, Phys. Rev. D 84, 047501 (2011).
- [194] S. Hadar, A. P. Porfyriadis, and A. Strominger, Phys. Rev. D 90, 064045 (2014).
- [195] S. Hadar, A. P. Porfyriadis, and A. Strominger, J. High Energy Phys. 07 (2015) 078.
- [196] A. Folacci and M. Ould El Hadj, Phys. Rev. D 98, 084008 (2018).
- [197] A. Apte and S. A. Hughes, Phys. Rev. D **100**, 084031 (2019).
- [198] H. Lim, G. Khanna, A. Apte, and S. A. Hughes, Phys. Rev. D 100, 084032 (2019).
- [199] S. A. Hughes, A. Apte, G. Khanna, and H. Lim, Phys. Rev. Lett. 123, 161101 (2019).
- [200] H. Lim, S. A. Hughes, and G. Khanna, Phys. Rev. D 105, 124030 (2022).
- [201] N. Oshita, J. Cosmol. Astropart. Phys. 04 (2023) 013.
- [202] N. Oshita and D. Tsuna, arXiv:2210.14049.
- [203] B. Rom and R. Sari, Phys. Rev. D 106, 104040 (2022).
- [204] N. T. Bishop, R. Gomez, L. Lehner, and J. Winicour, Phys. Rev. D 54, 6153 (1996).
- [205] M. Babiuc, B. Szilagyi, I. Hawke, and Y. Zlochower, Classical Quantum Gravity 22, 5089 (2005).
- [206] J. Winicour, Living Rev. Relativity 12, 3 (2009).
- [207] T. Pereira and R. Sturani, arXiv:2210.07299.