Low evidence for ringdown overtone in GW150914 when marginalizing over time and sky location uncertainty

Alex Correia,[1] Yi-Fan Wang,[2] Julian Westerweck,[3,4,5] and Collin D. Capano[6,1,*]

[1]*Physics Department, University of Massachusetts Dartmouth, North Dartmouth, Massachusetts 02747, USA*
[2]*Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Am Mühlenberg 1, D-14476 Potsdam, Germany*
[3]*Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Callinstraße 38, D-30167 Hannover, Germany*
[4]*Leibniz Universität Hannover, D-30167 Hannover, Germany*
[5]*Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, England, United Kingdom*
[6]*Department of Physics, Syracuse University, Syracuse, New York 13244, USA*

Tests of the no-hair theorem using astrophysical black holes involve the detection of at least two quasinormal modes (QNMs) in the gravitational wave emitted by a perturbed black hole. A detection of two modes—the dominant, $(\ell, m, n) = (2, 2, 0)$, mode and its first overtone, the $(2, 2, 1)$ mode—in the postmerger signal of the binary black hole merger GW150914 was claimed by Isi *et al.* [Testing the no-hair theorem with GW150914, Phys. Rev. Lett. **123**, 111102 (2019)], with further evidence provided by Isi and Farr [Revisiting the ringdown of GW150914 (2022)]. However, Cotesta *et al.* [Analysis of ringdown overtones in GW150914, Phys. Rev. Lett. **129**, 111102 (2022)] disputed this claim, finding that evidence for the overtone only appeared if the signal was analyzed before merger, when a QNM description of the signal is not valid. Because of technical challenges, both of these analyses fixed the merger time and sky location of GW150914 when estimating the evidence for the overtone. At least some of the contention can be attributed to fixing these parameters. Here, we surmount this difficulty and fully marginalize over merger time and sky-location uncertainty while doing a postmerger QNM analysis of GW150914. We find that marginalizing over all parameters yields low evidence for the presence of the overtone, with a Bayes factor of 1.10 ± 0.03 in favor of a QNM model with the overtone versus one without. The arrival time uncertainty of GW150914 is too large to definitively claim detection of the $(2, 2, 1)$ mode.

Introduction. The direct detection of gravitational waves has made it possible to perform "black hole spectroscopy" on real astrophysical black holes. Black holes that are formed from a binary black hole merger are initially perturbed. The perturbation is radiated away in a gravitational-wave "ringdown" as the black hole settles down to a Kerr spacetime. The ringdown is a superposition of damped sinusoids, or quasinormal modes (QNMs), that are uniquely determined by three indices: two angular indices $\ell \geq 2$ and $|m| \leq \ell$ and an overtone number $n \geq 0$ [1,2]. As a consequence of the no-hair theorem, we expect that the frequencies and damping times of all the QNMs are fully characterized by the final black hole's mass and spin. The Kerr nature of the final object can be tested if more than one QNM can be observed in a postmerger gravitational wave (GW) [3]. Should the QNMs be found to need more

than two parameters to characterize their frequency and damping times, it would indicate that the final object is not a Kerr black hole as expected from General Relativity (GR).

Black hole spectroscopy is challenging to accomplish in practice. The QNMs are exponentially damped, and so the signal-to-noise ratio (SNR) of the signal falls off quickly after merger. This would favor searching for QNMs immediately after merger, but it was thought that one had to wait $\sim 10M$ (where M is the mass of the black hole, in seconds) or longer (depending on the sensitivity of the detectors [4,5]) after the time at which the GW amplitude peaks (t_c) for nonlinearities to radiate away and a QNM description of the signal to be valid [6–10]. The $10M$ rule was arrived at by studying the stability of the "fundamental" $n = 0$ modes in numerical signals. However, Giesler *et al.* [11] show that including overtones of the dominant, $(\ell, m, n) = (2, 2, 0)$ mode—i.e., modes with $(2, 2, n \geq 1)$—yielded surprisingly good fits to a numerical relativity simulation at t_c. This led Giesler *et al.*

---

*Contact author: cdcapano@syr.edu

to claim that the spacetime could be described by a linearly perturbed black hole *at the merger*. Isi *et al.* [12] then performed an overtone analysis of the first detected binary black hole merger, GW150914 [13], and claimed detection of the (2, 2, 1) mode at 3.6σ. They performed a no-hair test using the overtone by allowing its frequency and damping time to vary in a Bayesian analysis, finding that it was broadly consistent with GR.

The validity of using a QNM description of the signal at t_c is controversial. Although it was known prior to the work of Giesler *et al.* that overtones yielded a better fit to the GW at merger [6], it was not clear whether this indicated that the spacetime was truly described by a linearly perturbed black hole (as claimed in Giesler *et al.*) or if it was simply due to the overtones overfitting the signal (meaning that any sufficiently compatible signal could be fit by them). While Giesler *et al.* provided some evidence for the former, subsequent analyses have challenged their interpretation [14–20], leaving the question unresolved.

Even the claim of a detection of an overtone in GW150914 by Isi *et al.* [12] has been challenged. Cotesta *et al.* [21] performed an overtone analysis at multiple times around the purported t_c of GW150914. At each time, they calculated the Bayes factor for a signal model consisting of the $(2, 2, 0) + (2, 2, 1)$ QNMs versus one consisting of just the (2, 2, 0) mode. Assuming equal prior weight for the two models, the Bayes factor $\mathcal{B} = Z_{220+221}/Z_{220}$ (where Z_x is the Bayesian evidence for model x) gives the odds ratio that the overtone model is a better description of the signal. Cotesta *et al.* found that the Bayes factor was ~ 1 at t_c, indicating that the data were uninformative as to the presence of the (2, 2, 1) mode. They claimed that positive evidence for the overtone only appeared if the QNM model was used *prior* to the merger. Since a QNM model is not valid there, this would indicate that the overtone model was simply overfitting the data. The low Bayes factor was also consistent with what had been reported by Bustillo *et al.* [22] and in publications from the LIGO and Virgo Collaborations [23,24]. However, in a rebuttal, Isi and Farr [27] found that the Bayes factor for the overtone was nearly 100 at t_c, reinforcing the initial finding by Isi *et al.* [12].

The disagreement between Isi *et al.* and Cotesta *et al.* may be due to several factors, including data conditioning choices and the manner in which the likelihood function is evaluated. The effects of these choices were investigated by Wang *et al.* [28]. It is evident, however, that the choice of start time for the ringdown is one of the most important factors—if not the defining factor [29]—in determining if the overtone is present. While they may disagree on the magnitude of the Bayes factor at a given time, both Cotesta *et al.* and Isi and Farr found that its value changes rapidly as the ringdown start time is varied around t_c, increasing by several orders of magnitude in less than $5M$. For GW150914, $5M$ corresponds to ~ 1.7 ms [13]. This is strikingly small; for comparison, the statistical uncertainty on t_c when using an inspiral-merger-ringdown (IMR) template to model the full signal is ~ 1.2 ms [30,31].

The statistical uncertainty in t_c arises from an overall uncertainty in the arrival time of the signal in a single detector and the uncertainty in the sky location of GW150914. The latter affects the relative offset in arrival times between the Hanford and Livingston detectors, which observed the event [13]. A full accounting of the evidence for or against the overtone should marginalize over the time and sky location uncertainty of GW150914. However, the analysis pipelines used by Isi *et al.* and Cotesta *et al.* fix the arrival time and sky location of the signal to their maximum likelihood values from IMR studies.

Even when changing the start time to recalculate the Bayes factor, the sky location was still fixed in the previous analyses. This was done to keep the ringdown start time unchanged during these pipelines' Bayesian inference analysis. Fixing the start time is necessary due to technical hurdles, discussed below, when calculating the likelihood. For this reason, almost every ringdown analysis that has been performed to date—both on GW150914 and other events—has fixed the sky location and merger time [12,21,23,27,28,32–36].

Finch and Moore were able to marginalize over the sky location and arrival time in their postmerger analysis of GW150914 [37]. They accomplished this by attaching an arbitrary signal model to the start of the ringdown. This had the effect of smoothly tapering the signal model to zero, thereby allowing them to use the traditional frequency-domain likelihood function that is used in standard parameter estimation analyses. Doing so, they recovered Bayes factors B_{220}^{221} between 0.1 and 10, depending on the prior chosen on t_c. Their prior choices consisted of Gaussian distributions of width $1M$ centered at various t_c values.

By using the traditional frequency-domain likelihood function, Finch and Moore coupled information from the premerger part of the signal with the postmerger QNM model. The frequency-domain likelihood is equivalent to the convolution of the "whitened" data with the whitened signal model. The whitening filter couples adjacent points in time over several milliseconds. This means, for example, that the beginning of the whitened ringdown will be a weighted sum over a time span starting before the merger that is on order with the damping time of the overtone (~ 1.4 ms). Consequently, the recovered QNM parameters will be influenced by both the choice of signal model and the data just prior to the start of the ringdown.

Here, we overcome the technical hurdles that arise when marginalizing over time and sky location uncertainty while doing a purely QNM analysis of the postmerger portion of GW150914. We simultaneously analyze the premerger data using the inspiral-merger portion of an IMR template waveform. The pre- and postmerger portions of the analysis only share a common sky location and geocentric t_c.

This allows us to use the full signal to isolate the binary, giving the most accurate assessment of the timing uncertainty, while still maintaining independence between the pre- and postmerger analyses. In a companion paper [38], we show that it is necessary to model the entire observable signal in order to marginalize over timing uncertainties. Consequently, our analysis here is the most complete accounting of all uncertainties that can be preformed when using a QNM-only model of the postmerger signal.

Methods. An in-depth description of our analysis is presented in a following paper [38]. Here, we give a brief overview.

The Bayesian evidence for a signal model is given by marginalizing the likelihood function times the prior over all the model parameters ϑ. In GW astronomy, it is common to assume that each interferometer outputs independent, wide-sense stationary Gaussian noise with zero mean in the absence of a GW signal. Under this assumption, the likelihood function is given by

$$p(\mathbf{s}|\vartheta, h) = \prod_{D=H,L,\dots} \frac{\exp\left[-\frac{1}{2}\mathbf{r}_D^\mathsf{T}\mathbf{\Sigma}_D^{-1}\mathbf{r}_D\right]}{\sqrt{(2\pi)^N \det \mathbf{\Sigma}_D}}. \qquad (1)$$

Here, the product is taken over the number of detectors, $\mathbf{r}_D = \mathbf{s}_D - \mathbf{h}_D(\vartheta)$ is the residual after subtracting the signal model $\mathbf{h}(\vartheta)$ from the time-domain data $\mathbf{s}_D$, which contains N samples. The elements of the covariance matrix $\mathbf{\Sigma}_D$ are given by the autocorrelation function of the data.

In standard GW parameter estimation (in which $\mathbf{h}$ models the entire observable signal), $\mathbf{\Sigma}_D$ can be approximated as a circulant matrix. That yields an analytic solution for $\mathbf{\Sigma}_D^{-1}$, which in turn yields a weighted inner product for the log likelihood that is most efficiently calculated in the frequency domain. However, since a QNM template only models the ringdown portion of the signal, data prior to the start of the ringdown must be excised from the analysis, or else biases will arise. Excising data breaks the circulant approximation of $\mathbf{\Sigma}_D$, and so the standard frequency-domain likelihood function cannot be used.

There are two approaches around this. One is to numerically invert $\mathbf{\Sigma}_D$ and explicitly calculate the matrix multiplication in Eq. (1). This is what is done by the "RINGDOWN" [39] and "pyRing" [25,26] analyses used by Isi *et al.* and Cotesta *et al.*, respectively. Numerically inverting $\mathbf{\Sigma}_D$ is computationally expensive but needs only to be done once per analysis if the start time of the ringdown does not change. If the start time of the ringdown does change (meaning different times need to be excised from the data), then $\mathbf{\Sigma}_D^{-1}$ and $\det \mathbf{\Sigma}_D$ need to be recalculated, making the analysis substantially more expensive.

The other approach is to use "gating and in-painting" to solve the likelihood [40]. This is mathematically equivalent to the brute force numerical method but is more efficient for longer stretches of data [28,41]. In this approach, a gate is applied to a stretch of data, within which the data are zeroed out. The gated portion is then in-painted so that it contributes nothing to the likelihood. For a ringdown-only analysis, the gate ends at the start of the ringdown; we use a 1s gate to ensure that the rest of the observable signal is removed. The gate shifts around in time as the start time of the ringdown is varied, thereby removing more or less of the signal from the analysis.

The cost of computing the numerator in Eq. (1) with the gating and in-painting method is the same regardless of the start time of the ringdown. However, $\det \mathbf{\Sigma}_D$ still needs to be calculated, as its value will change for different ringdown start times. For this reason, studies that have used gating and in-painting have still fixed the start time and sky location when doing QNM analyses [28,32,34–36,42].

In Ref. [38], we show that to high precision $\det \mathbf{\Sigma}_D$ can be approximated by a simple linear interpolation. This allows us to quickly evaluate all terms in Eq. (1) using gating and in-painting, regardless of the ringdown start time, thereby allowing for time and sky-location marginalization.

We also show in Ref. [38] that it is not possible to only model part of a signal when marginalizing over t_c and sky location. Marginalizing over these parameters while only modeling part of the signal means that the portion of data that is excised from the analysis is allowed to shift around in time. If the shift is such that the entire observable signal is excised from the analysis, then a relatively large likelihood can be obtained for any parameter value by reducing the overall amplitude of the template to noise level. Consequently, the posterior probability will favor excising as much of the signal as possible, even the part that the template is meant to model. This means that for a ringdown-only analysis the start time of the template will be pushed to as *late* as the prior will allow. Likewise, if an inspiral-only analysis is performed (in which the postmerger is excised and the premerger is modeled with a template), then the end time of the template will be pushed to as *early* as the prior will allow. It is necessary to have a signal model for all times that a signal is observable in the data when doing any sort of marginalization over arrival time.

We overcome this problem by simultaneously yet independently modeling the pre- and postmerger data. We treat the pre- and postmerger signal as separate analyses that share a common set of hyperparameters—the sky location and geocentric t_c. We use gating and in-painting for both. For the premerger analysis, we excise the data *after* t_c, while for the postmerger analysis, we excise the data *before* t_c. The premerger portion of the signal is modeled using the IMRPhenomXPHM approximant [43], which we zero out for all times after t_c. The postmerger signal is modeled using a QNM model consisting of the $(2, 2, 0) + (2, 2, 1)$ modes, which is zeroed out for all times before t_c. The final mass and final spin in the QNM model are varied over uniform priors such that $M_f \in [10, 200]M_\odot$ and $\chi_f \in [-0.99, 0.99]$. For the coalescence time, we use a

FIG. 1. Whitened strain in the Hanford (top) and Livingston (bottom) detectors around GW150914. Plotted is the original data (black line) along with the maximum likelihood waveform, which consists of an inspiral portion (blue lines) and ringdown portion (red lines). The inspiral and ringdown models are independent of each other and only share a common sky location and geocentric t_c. Ringdown (inspiral) waveforms are gated and in-painted before (after) the t_c in each detector; the first (last) nonzero point in the waveform is indicated by the solid dot. The data are also gated and in-painted before (after) the detector t_c for the ringdown (inspiral) model so that neither model simultaneously includes the same data point from the signal in the likelihood. Individual waveforms drawn from the upper 90th percentile of the posterior are also plotted to indicate the uncertainty region. Vertical purple lines indicate the gate time for each of the waveforms plotted. The time axis is zeroed with respect to the maximum likelihood t_c in each detector $t_{\text{ref}}^{H/L}$ from the 4-OGC IMR analysis [31].

uniform prior spanning $t_{\text{ref}} \pm 0.1$ s, where t_{ref} is the maximum likelihood geocentric coalescence time from Ref. [31]. This is the same prior width on t_c as used in standard IMR analyses.

For the (2, 2, 0) amplitude A_{220}, we sample over the range $[10^{-25}, 8 \times 10^{-17}]$ uniformly in log base 10. For the (2, 2, 1) amplitude A_{221}, we use a uniform prior on $A_{221}/A_{220} \in [0, 5]$. We additionally impose the constraint that the (2, 2, 0) mode's contribution to the postmerger SNR be >2. This is to ensure that at least the fundamental mode is present in the postmerger signal. Otherwise, we obtain samples where the (2, 2, 1) mode appears to have large amplitude but is in fact matching the (2, 2, 0) mode in the signal (the "label switching" problem).

Gating the pre- and postmerger means that no coupling occurs over the boundary due to whitening. Simultaneously modeling both portions has the added benefit that we use the entire signal to constrain the sky location. We use PyCBC Inference [44] with the DYNESTY sampler [45] to perform the Bayesian analysis.

Results and discussion. The results of our analysis are summarized in Figs. 1 and 2. Figure 1 shows the recovered

inspiral-only and ringdown-only templates. In Fig. 2, we plot the marginal posterior on the (2, 2, 1) amplitude A_{221} and t_c in the Hanford detector. The median and 90% credible interval for the (2, 2, 1) amplitude is $A_{221}/A_{220} = 0.9^{+1.8}_{-0.7}$; for the coalescence time, it is $t_c^H - t_{\text{ref}}^H = 0.0^{+1.0}_{-1.1}$ ms in Hanford and $t_c^L - t_{\text{ref}}^L = 0.2^{+1.1}_{-1.0}$ ms in Livingston [46].

Figure 3 compares the sky location posteriors of our analysis with a full IMR analysis from the fourth Open Gravitational-wave Catalog (4-OGC) [31]. Our analysis recovers most of the 90th percentile of the 4-OGC posterior. Our marginal posterior on coalescence time is also consistent with the full IMR analysis (see Fig. 2). This indicates that our sky and time marginalization implementation is able to recover nominal sky parameters for GW150914.

We estimate the Bayes factor $\mathcal{B}$ between $Z_{220+221}$ and Z_{220} by calculating the Savage-Dickey ratio [47] at $A_{221}/A_{220} = 0$. Doing so, we obtain $\mathcal{B} = 1.10 \pm 0.03$, indicating low support in favor of the (2, 2, 1) mode. Similarly, the median of the posterior distribution for A_{221} deviates from 0 by 1.2σ [48]. The reason for the small Bayes factor for the (2, 2, 1) mode is evident from the marginal posterior plot in Fig. 2. While the maximum likelihood point is at $A_{221}/A_{220} = 1.03$ and $t_c^H - t_{\text{ref}}^H = 0.03$ ms,

FIG. 2. Marginal posteriors on A_{221}/A_{220} and the coalescence time as measured in the Hanford detector t_c^H; similar results are obtained for the Livingston detector (not shown). The horizontal axis shows the ratio between the (2, 2, 1) and (2, 2, 0) mode amplitudes. The vertical axis shows the difference between the measured t_c^H and the reference time t_{ref}^H, which is the maximum likelihood t_c^H from Ref. [31]. The filled marginal histograms represent the one-dimensional distributions for each parameter. The prior on A_{221}/A_{220} is overlaid on the corresponding marginal posterior with a dotted black line. The 4-OGC t_c^H posterior is overlaid on the t_c marginal posterior in gray. Each point in the two-dimensional scatter plot is colored according to its log likelihood value relative to the maximum log likelihood. The maximum likelihood point is denoted with a black cross. The two-dimensional credible regions are indicated by the contour lines; shown are the 50th (dashed contour) and 90th (solid contours) percentiles. The Bayes factor for the (2, 2, 1) mode is estimated from the A_{221}/A_{220} marginal posterior using the Savagey-Dickey ratio. This is the ratio between the posterior density (solid black line) and the prior density (dotted black line) at $A_{221}/A_{220} = 0$.

the A_{221}/A_{220} posterior has support at zero over much of the t_c^H posterior (similarly for t_c^L). It is only at earlier times that the posterior on A_{221} is markedly shifted away from zero. This is consistent with previous studies which found the evidence for the (2, 2, 1) mode increased at earlier times [21,27,28]. However, due to the relatively large timing uncertainty—our credible interval on $t_c^{H/L}$ spans roughly $5M$—we cannot confidently claim detection of the overtone, as reflected by the Bayes factor. A signal with more SNR in the ringdown is needed.

Including the overtone does yield a mass and spin estimate of the final black hole that is more consistent with the full IMR analysis. This can be seen in Fig. 4, which compares the redshifted final mass and final spin parameters of the full IMR analysis to the QNM model that includes the (2, 2, 1) overtone and to the QNM model that includes only the (2, 2, 0) mode [49]. This is consistent with previous analyses of GW150914 [50], which also found


FIG. 3. Sky location posteriors for our analysis depicted in Fig. 2 and a full IMR analysis from 4-OGC [31]. Contours indicate the 90th percentile of points for each run. The red contour denotes the analysis conducted with sky and time marginalization described in Methods. The black contour denotes the 4-OGC analysis.


FIG. 4. Final mass and spin posteriors for our analysis depicted in Fig. 2, a (2, 2, 0)-only ringdown analysis, and a full IMR analysis from 4-OGC [31]. Contours indicate the 90th percentile of points for each run. The blue contour denotes our ringdown analysis with both the (2, 2, 0) and (2, 2, 1) modes. The red contour denotes an analysis with only the dominant mode present in the ringdown. The black contour denotes the 4-OGC analysis. The maximum likelihood points of the models are indicated by the crosses. The median and 90% credible intervals from the overtone model are $71^{+30}_{-16}M_\odot$ and $0.70^{+0.25}_{-0.47}$ for the redshifted mass and spin, respectively, while for the (2, 2, 0)-only model, they are $84^{+13}_{-27}M_\odot$ and $0.84^{+0.09}_{-0.55}$.

that using a (2, 2, 0)-only QNM model at merger yielded biased results. Indeed, a major motivation for including the (2, 2, 1) mode is that it yields better mass and spin estimates at merger [11].

However, the overtone posterior is also *wider*: the 90% credible intervals from the overtone model span $46M_\odot$ and 0.72 for the red-shifted mass and spin, respectively, while for the (2, 2, 0)-only model, they span $40M_\odot$ and 0.64. It is therefore unclear whether the better fit is simply due to the overtone model's additional freedom. This is borne out by the Bayes factor. A key feature of Bayes factors is they account for model complexity. That it is ∼1 between the overtone and (2, 2, 0)-only models indicates that the better fit obtained by the former is offset by the Occam penalty from its larger number of parameters.

A detection of the (3, 3, 0) subdominant QNM in the ringdown of GW190521 was claimed by Capano *et al.* [32,51]. The sky location was also fixed in that analysis, and the Bayes factor was maximized over a grid of possible merger times. Our results here raise the question of whether the evidence for the subdominant mode in GW190521 would also decrease if the sky location and t_c were marginalized over.

A reanalysis of GW190521 is beyond the scope of this paper. However, there are a few differences between the overtone claim in GW150914 and the (3, 3, 0) claim in GW190521 that indicate that sky/time marginalization will not negatively affect the Bayes factor for the latter, and in fact may increase it. First, the SNR in the ringdown of GW190521 is about twice that of GW150914 [13,32]. Second, the difference in arrival times between Hanford and Livingston for GW190521 is $\sim 0.1\tau_{330}$ [31], where τ_{330} is the damping time of the putative (3, 3, 0) mode. For comparison, the difference in arrival times in units of the overtone damping time is $\sim 5\tau_{221}$ for GW150914. Marginalizing over sky location—which roughly corresponds to marginalizing over uncertainty in the difference in arrival times—is likely to have less of an effect on the evidence for the GW190521 (3, 3, 0) mode.

Assuming the sky location used by Capano *et al.*, the evidence for the (3, 3, 0) mode stays elevated across a large range of times after merger [52]. In contrast, the evidence for the (2,2,1) mode for GW150914 drops quickly after merger. A naive marginalization over the discrete times published by Capano *et al.*—accomplished by summing the log evidence over those times—yields a Bayes factor for the (3, 3, 0) mode of ∼5000 [53]. A reanalysis that uses a model for the pre- and postmerger signal and marginalizes over time and sky-location uncertainty is necessary to fully evaluate the evidence for the (3, 3, 0) mode, and any other subdominant QNM, in GW190521. Nevertheless, our cursory analysis indicates that the evidence for the (3, 3, 0)

mode in GW190521 may in fact increase with full sky and time marginalization. We plan to revisit this in a future study.

Returning to overtones, our results show that GW150914 is too weak to definitively claim detection of the (2, 2, 1) mode. Gravitational-wave detectors are continuing to improve in sensitivity [31,54,55]. A signal with sufficient SNR to resolve questions about the presence of one or more overtones is likely to eventually be observed [56]. When such an event occurs, it will be necessary to fully marginalize over both time and sky-location uncertainties when performing a QNM analysis of the ringdown, as we have done here.

All computations were performed on Unity, a collaborative, multi-institutional high-performance computing cluster managed by UMass Amherst Research Computing and Data.

This research was conducted using PyCBC [57]. Data are available in Ref. [58].

[1] C. V. Vishveshwara, Scattering of gravitational radiation by a Schwarzschild black-hole, Nature (London) **227,** 936 (1970).

[2] S. Chandrasekhar and S. Detweiler, The quasi-normal modes of the Schwarzschild black hole, Proc. R. Soc. A **344,** 441 (1975).

[3] O. Dreyer, B. Kelly, B. Krishnan, L. S. Finn, D. Garrison, and R. Lopez-Aleman, Black-hole spectroscopy: Testing general relativity through gravitational-wave observations, Classical Quantum Gravity **21,** 787 (2004).

[4] L. T. London, Modeling ringdown. II. Aligned-spin binary black holes, implications for data analysis and fundamental theory, Phys. Rev. D **102,** 084052 (2020).

[5] G. Carullo et al., Empirical tests of the black hole no-hair conjecture using gravitational-wave observations, Phys. Rev. D **98,** 104020 (2018).

[6] A. Buonanno, G. B. Cook, and F. Pretorius, Inspiral, merger and ring-down of equal-mass black-hole binaries, Phys. Rev. D **75,** 124018 (2007).

[7] E. Berti, V. Cardoso, J. A. Gonzalez, U. Sperhake, M. Hannam, S. Husa, and B. Bruegmann, Inspiral, merger and ringdown of unequal mass black hole binaries: A multipolar analysis, Phys. Rev. D **76,** 064034 (2007).

[8] I. Kamaretsos, M. Hannam, S. Husa, and B. S. Sathyaprakash, Black-hole hair loss: Learning about binary progenitors from ringdown signals, Phys. Rev. D **85,** 024018 (2012).

[9] L. London, D. Shoemaker, and J. Healy, Modeling ring-down: Beyond the fundamental quasinormal modes, Phys. Rev. D **90,** 124032 (2014); **94,** 069902(E) (2016).

[10] A. Gupta, B. Krishnan, A. Nielsen, and E. Schnetter, Dynamics of marginally trapped surfaces in a binary black hole merger: Growth and approach to equilibrium, Phys. Rev. D **97,** 084028 (2018).

[11] M. Giesler, M. Isi, M. A. Scheel, and S. A. Teukolsky, Black hole ringdown: The importance of overtones, Phys. Rev. X **9,** 041060 (2019).

[12] M. Isi, M. Giesler, W. M. Farr, M. A. Scheel, and S. A. Teukolsky, Testing the no-hair theorem with GW150914, Phys. Rev. Lett. **123,** 111102 (2019).

[13] B. P. Abbott et al. (LIGO Scientific and Virgo Collaborations), Observation of gravitational waves from a binary black hole merger, Phys. Rev. Lett. **116,** 061102 (2016).

[14] S. Bhagwat, X. J. Forteza, P. Pani, and V. Ferrari, Ringdown overtones, black hole spectroscopy, and no-hair theorem tests, Phys. Rev. D **101,** 044033 (2020).

[15] J. L. Jaramillo, R. Panosso Macedo, and L. Al Sheikh, Pseudospectrum and black hole quasinormal mode instability, Phys. Rev. X **11,** 031003 (2021).

[16] S. Ma, K. Mitman, L. Sun, N. Deppe, F. Hébert, L. E. Kidder, J. Moxon, W. Throwe, N. L. Vu, and Y. Chen, Quasinormal-mode filters: A new approach to analyze the gravitational-wave ringdown of binary black-hole mergers, Phys. Rev. D **106,** 084036 (2022).

[17] M. H.-Y. Cheung et al., Nonlinear effects in black hole ringdown, Phys. Rev. Lett. **130,** 081401 (2023).

[18] K. Mitman et al., Nonlinearities in black hole ringdowns, Phys. Rev. Lett. **130,** 081402 (2023).

[19] P. J. Nee, S. H. Völkel, and H. P. Pfeiffer, Role of black hole quasinormal mode overtones for ringdown analysis, Phys. Rev. D **108,** 044032 (2023).

[20] V. Baibhav, M. H.-Y. Cheung, E. Berti, V. Cardoso, G. Carullo, R. Cotesta, W. Del Pozzo, and F. Duque, Agnostic black hole spectroscopy: Quasinormal mode content of numerical relativity waveforms and limits of validity of linear perturbation theory, Phys. Rev. D **108,** 104020 (2023).

[21] R. Cotesta, G. Carullo, E. Berti, and V. Cardoso, Analysis of ringdown overtones in GW150914, Phys. Rev. Lett. **129,** 111102 (2022).

[22] J. Calderón Bustillo, P. D. Lasky, and E. Thrane, Black-hole spectroscopy, the no-hair theorem, and GW150914: Kerr versus Occam, Phys. Rev. D **103,** 024041 (2021).

[23] R. Abbott et al. (LIGO Scientific and Virgo Collaborations), Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog, Phys. Rev. D **103,** 122002 (2021).

[24] It should be noted that the pipeline [25,26] used to report overtone Bayes factors by the LVC is the same as used by Cotesta et al., so the agreement between the two is not entirely surprising.

[25] G. Carullo, W. D. Pozzo, and J. Veitch, Observational black hole spectroscopy: A time-domain multimode analysis of GW150914, Phys. Rev. Lett. D **99,** 123029 (2019).

[26] G. Carullo, W. Del Pozzo, and J. Veitch, pyRing: A time-domain ringdown analysis Python package, https://git.ligo.org/lscsoft/pyring (2023).

[27] M. Isi and W. M. Farr, Revisiting the ringdown of GW150914, arXiv:2202.02941.

[28] Y.-F. Wang, C. D. Capano, J. Abedi, S. Kastha, B. Krishnan, A. B. Nielsen, A. H. Nitz, and J. Westerweck, A frequency-domain perspective on GW150914 ringdown overtone, arXiv:2310.19645.

[29] G. Carullo, R. Cotesta, E. Berti, and V. Cardoso, Reply to comment on "Analysis of ringdown overtones in GW150914", Phys. Rev. Lett. **131,** 169002 (2023).

[30] B. P. Abbott et al. (LIGO Scientific and Virgo Collaborations), Properties of the binary black hole merger GW150914, Phys. Rev. Lett. **116,** 241102 (2016).

[31] A. H. Nitz, S. Kumar, Y. Wang, S. Kastha, S. Wu, M. Schäfer, R. Dhurkunde, and C. D. Capano, 4-OGC: Catalog of gravitational waves from compact-binary mergers, Astrophys. J. **946,** 59 (2023).

[32] C. D. Capano, M. Cabero, J. Westerweck, J. Abedi, S. Kastha, A. H. Nitz, Y.-F. Wang, A. B. Nielsen, and B. Krishnan, Multimode quasinormal spectrum from a perturbed black hole, Phys. Rev. Lett. **131,** 221402 (2023).

[33] M. Isi, W. M. Farr, M. Giesler, M. A. Scheel, and S. A. Teukolsky, Testing the black-hole area law with GW150914, Phys. Rev. Lett. **127,** 011103 (2021).

[34] S. Kastha, C. D. Capano, J. Westerweck, M. Cabero, B. Krishnan, and A. B. Nielsen, Model systematics in time domain tests of binary black hole evolution, Phys. Rev. D **105,** 064042 (2022).

[35] C. D. Capano, J. Abedi, S. Kastha, A. H. Nitz, J. Westerweck, Y.-F. Wang, M. Cabero, A. B. Nielsen, and B. Krishnan, Statistical validation of the detection of a sub-dominant quasi-normal mode in GW190521, arXiv:2209.00640.

[36] J. Abedi, C. D. Capano, S. Kastha, A. H. Nitz, Y.-F. Wang, J. Westerweck, A. B. Nielsen, and B. Krishnan, Spectroscopy

for asymmetric binary black hole mergers, Phys. Rev. D **108**, 104009 (2023).

[37] E. Finch and C. J. Moore, Searching for a ringdown overtone in GW150914, Phys. Rev. D **106**, 043005 (2022).

[38] A. Correia and C. D. Capano, companion paper, Sky marginalization in black hole spectroscopy and tests of the area theorem, Phys. Rev. D **110**, 044018 (2024).

[39] M. Isi and W. M. Farr, maxisi/ringdown: Initial ringdown release, 10.5281/zenodo.5094068 (2021).

[40] B. Zackay, T. Venumadhav, J. Roulet, L. Dai, and M. Zaldarriaga, Detecting gravitational waves in data with non-stationary and non-Gaussian noise, Phys. Rev. Lett. D **104**, 063034 (2021).

[41] M. Isi and W. M. Farr, Analyzing black-hole ringdowns, arXiv:2107.05609.

[42] H. Siegel, M. Isi, and W. M. Farr, Ringdown of GW190521: Hints of multiple quasinormal modes with a precessional interpretation, Phys. Rev. D **108**, 064008 (2023).

[43] G. Pratten, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, M. Mateu-Lucena, R. Jaume, M. Haney, D. Keitel, J. E. Thompson, and S. Husa, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, Phys. Rev. D **103**, 104056 (2021).

[44] C. M. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, PyCBC Inference: A Python-based parameter estimation toolkit for compact binary coalescence signals, Publ. Astron. Soc. Pac. **131**, 024503 (2019).

[45] J. S. Speagle, DYNESTY: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, Mon. Not. R. Astron. Soc. **493**, 3132 (2020).

[46] We use the maximum likelihood sky location and coalesence time from the 4-OGC analysis [31] to calculate the reference GPS time in each detector. Specifically, we use $t_{\mathrm{ref}}^H = 1126259462.423701$ and $t_{\mathrm{ref}}^L = 1126259462.416674$.

[47] I. Verdinelli and L. Wasserman, Computing Bayes factors using a generalization of the Savage-Dickey density ratio, J. Am. Stat. Assoc. **90**, 614 (1993).

[48] The 1.2σ quoted here was computed using the relative amplitude posterior A_{221}/A_{220}. Recalculating using the absolute A_{221} posterior yields the same result.

[49] To obtain the (2, 2, 0)-only posterior on mass and spin, we preform a full Bayesian analysis in which the (2, 2, 0)-only QNM starts at t_c.

[50] B. P. Abbott et al. (LIGO Scientific and Virgo Collaborations), Tests of general relativity with GW150914, Phys. Rev. Lett. **116**, 221101 (2016); **121**, 129902(E) (2018).

[51] See also the work by Siegel et al. [42], which found evidence for the (2, 1, 0) and (3, 2, 0) modes in the same event.

[52] C. Capano, S. Kastha, M. Cabero, J. Westerweck, and Y. Wang, gwastro/BH-Spectroscopy-GW190521: Published data release (v4.0). Zenodo, 10.5281/zenodo.10056546.

[53] We report the Bayes factor for the (3, 3, 0) mode as being $Z_{220+330}/\max(Z_{220}, Z_{220+221})$, as was done by Capano et al. [32].

[54] R. Abbott et al. (LIGO Scientific, Virgo, and KAGRA Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, Phys. Rev. X **13**, 041039 (2023).

[55] LIGO, Virgo, and KAGRA observering run plans, https://observing.docs.ligo.org/plan/.

[56] Although larger SNR may result in an ambiguous number of overtones needed to fully model the merger, raising questions of overfitting [22].

[57] A. Nitz et al., gwastro/pycbc: v2.3.3 release of PyCBC (v2.3.3), Zenodo, 10.5281/zenodo.10473621.

[58] A. Correia and C. Capano, Sky marginalization in black hole spectroscopy, 10.5281/zenodo.10493247 (2024).

[59] https://gwosc.org/.