

Sky marginalization in black hole spectroscopy and tests of the area theorem

Alex Correia 

*University of Massachusetts Dartmouth, 285 Old Westport Rd,
North Dartmouth, Massachusetts 02747, USA*

Collin D. Capano *

*Department of Physics, Syracuse University, Syracuse, New York 13244, USA
and University of Massachusetts Dartmouth, 285 Old Westport Rd,
North Dartmouth, Massachusetts 02747, USA*



(Received 18 January 2024; accepted 28 June 2024; published 7 August 2024)

Direct observation of gravitational waves from binary black hole (BBH) mergers has made it possible to test the laws of black hole thermodynamics using real astrophysical sources. These tests rely on accurate and unbiased parameter estimates from the pre and postmerger portions of a signal. Due to numerical complications, previous analyses have fixed the sky location and coalescence time when independently estimating the parameters of the pre and postmerger signal. Here we overcome the numerical complications and present a novel method of marginalizing over sky location and coalescence time. Doing so, we find that it is not possible to model only the pre or postmerger portions of the signal while marginalizing over timing uncertainty. We surmount this problem by simultaneously yet independently modeling the pre and postmerger signal, with only the sky location and coalescence time being shared between the models. This allows us to marginalize over all parameters. We use our method to measure the change in area $\Delta A_{\text{measured}} = A_f - A_i$ between the final and initial black holes in the BBH merger GW150914. To measure the final black hole's area A_f we do an analysis using quasinormal modes (QNMs) to model the postmerger signal, and another analysis using the postmerger portion of an inspiral-merger-ringdown (IMR) template. We find excellent agreement with expectations from general relativity. The Hawking area theorem (which states that $A_f \geq A_i$) is confirmed to 95.4% and 99.5% confidence using the QNM and IMR postmerger models, respectively. Both models yield $\Delta A_{\text{measured}}/\Delta A_{\text{expected}} \sim 1$, where $\Delta A_{\text{expected}}$ is the expected change in area derived from fits to numerical relativity simulations.

DOI: [10.1103/PhysRevD.110.044018](https://doi.org/10.1103/PhysRevD.110.044018)

I. INTRODUCTION

The direct detection of gravitational waves (GWs) throughout the past decade has been one of the most significant advancements in observational relativity. These observations not only confirmed the existence of GWs, but also opened the door to direct tests of general relativity (GR) in the strong-field regime. One such prediction that can be verified is Hawking's area theorem, which states that the remnant from a binary black hole (BBH) merger must have an event horizon surface area greater than the sum of the progenitor horizon areas [1].

Tests of the area theorem using GWs have been carried out previously [2–4]. Generally, they involve measuring the initial two black holes' mass and spin (from which the total initial area A_i is derived) from the premerger part of the signal, during which the two black holes inspiral into each other. The area of the final black hole is independently

measured using the GW that is emitted during the postmerger, or “ringdown” phase. In both cases Bayesian inference is used to produce a “posterior” probability density on the black holes' parameters. Ideally, all parameters that describe the BBH should be allowed to vary in the analysis, in order to fully account for all statistical uncertainties. However, previous tests of the area theorem have fixed the sky position and coalescence time t_c of the events to nominal values when doing their analysis [2,4]. Fixing the values in this way may lead to biases in the resultant parameter estimates, obfuscating the true nature of the system [3,5]. At the very least, it may cause an underestimate of the statistical uncertainty of measured parameters, yielding constraints on deviations from GR that are misleadingly strong.

The sky location and t_c have been fixed in earlier studies due to technical hurdles in calculating the likelihood function. In order to independently analyze the pre and postmerger portion of the signal it is necessary to excise the post and premerger data, respectively, from the analysis.

*Contact author: cdcapano@syr.edu

Excising the premerger data is also necessary when just analyzing the postmerger signal for tests of the no-hair theorem using quasinormal modes (known as black hole spectroscopy). In either case, the excision (or “gate”) results in a modified likelihood function that cannot be solved using conventional, frequency-domain means. Existing pipelines for doing such analyses in the time domain, such as `pyring` [6,7] and `ringdown` [8], instead calculate the likelihood by numerically inverting the noise covariance matrix for the data. This calculation can be numerically costly [5,9]. If the sky location and t_c are not fixed, the gate will vary in time, and the full likelihood will need to be recalculated for every unique gate time, significantly increasing computational costs. However, if the sky location and t_c are fixed, the gate is static, and the most computationally demanding part of the calculation need only be done once.

Finch and Moore [10] devised a method to overcome this issue and vary the sky location and t_c . In their method the inspiral and merger are modeled with a wavelet series stitched to the beginning of the ringdown. This allows them to use the traditional frequency-domain likelihood, and vary the start time of the ringdown. However, this does not allow for an area theorem test, since the parameters of the initial black holes cannot be estimated from the wavelet signal model. Furthermore, the traditional frequency-domain likelihood causes the start of the ringdown to be coupled to information from just before the merger due to convolution with the whitening filter. This means that recovered pre and postmerger parameters will not be independent measurements.

This paper presents a novel method of calculating the likelihood function using the parameter estimation code `PyCBC Inference` [11]. This code already utilizes “gating and in-painting” [12] to excise data from the analysis, which allows for cheaper likelihood evaluation under certain conditions as compared to the methods used by `pyring` and `ringdown` [5,9]. However, an additional normalization factor has traditionally been omitted from this calculation, as this also required computationally expensive numerical methods to evaluate. This paper presents a method of linear interpolation that calculates this normalization factor with good approximation. Together, these methods allow for fast likelihood calculations regardless of gate position, allowing for marginalization over t_c and sky location and a full accounting of parameter uncertainties.

This paper is structured as follows. Section II describes the modifications to the likelihood calculation in `PyCBC` in detail. Full waveform analyses can be conducted with these modifications to marginalize over sky location and t_c . However, the method cannot be used in partial waveform analyses. Section III documents the issues that arise in these models. Using this method, a test of the Hawking area theorem is conducted using data from GW150914 [13]. Section IV describes the configuration of the analyses used

in this test, and Sec. V describes and analyzes the results. Section VI summarizes the findings and the implications of the marginalization method for future GW analyses.

II. DETERMINANT APPROXIMATION

`PyCBC Inference` [11] utilizes Bayesian inference to conduct parameter estimation on GW events. Bayes’ theorem is used to extract information about the parameter space $\boldsymbol{\theta}$ from a dataset s [14]. Assuming the dataset is composed of a signal h and noise n such that $s = n + h$, the probability of observing a specific $\boldsymbol{\theta}$ is given by

$$p(\boldsymbol{\theta}|s, h) = \frac{p(s|\boldsymbol{\theta}, h)p(\boldsymbol{\theta}|h)}{p(s|h)}. \quad (1)$$

The term $p(s|\boldsymbol{\theta}, h)$ is the likelihood, which describes the probability of observing a signal s assuming the event has a parameter space $\boldsymbol{\theta}$. The term $p(\boldsymbol{\theta}|h)$ is the prior, a distribution that describes the *a priori* probability of observing $\boldsymbol{\theta}$ given a signal model h . The denominator $p(s|h)$ is the evidence, a normalization factor used to compare analyses using different models for h . The resultant probability distribution on the left-hand side of the equation is known as the posterior.

The prior is chosen at the discretion of the analyst based on assumed plausible values for $\boldsymbol{\theta}$, whereas the likelihood is calculated directly from the data. For a system of K detectors each taking N time series samples of an event with a stochastic Gaussian noise background, the likelihood can be written as [15]

$$p(s_{\text{net}}|n) = \frac{\exp[-\frac{1}{2}\sum_{d=1}^K s_d^T \boldsymbol{\Sigma}_d^{-1} s_d]}{[(2\pi)^{NK} \prod_{d=1}^K \det \boldsymbol{\Sigma}_d]^{1/2}}, \quad (2)$$

where $\boldsymbol{\Sigma}_d$ is the covariance matrix associated with n in detector d .

To evaluate the likelihood, further assumptions must be made to calculate $\boldsymbol{\Sigma}_d^{-1}$. The result of these assumptions is

$$p(s_{\text{net}}|n) \propto \exp\left[-\frac{1}{2}\sum_{d=1}^K \langle s_d, s_d \rangle\right], \quad (3)$$

where the inner product is defined by Eq. (A14). (The full derivation of this likelihood function is given in Appendix A). This gives the likelihood function for a signal that is assumed to be entirely noise. To get the likelihood function with respect to the signal model h evaluated for a parameter space $\boldsymbol{\theta}$, Eq. (3) can be rewritten by substituting $s_d \rightarrow s_d - h_d(\boldsymbol{\theta})$ on the right-hand side, yielding

$$\log p(s_{\text{net}}|\boldsymbol{\theta}, h) = -\frac{1}{2} \sum_{d=1}^K \langle s_d - \mathbf{h}_d(\boldsymbol{\theta}), s_d - \mathbf{h}_d(\boldsymbol{\theta}) \rangle - \frac{1}{2} \left[NK \log 2\pi + \sum_{d=1}^K \log \det \boldsymbol{\Sigma}_d \right]. \quad (4)$$

(Here and throughout, we use \log to refer to the natural logarithm.) For a GW analysis that examines the entirety of s , this expression is easily evaluated using the approximated eigenvalues of $\boldsymbol{\Sigma}_d$. However, many analyses only examine a portion of the time series (generally either the pre or postmerger signal). Therefore, a region of s is omitted, or “gated”, to reduce numerical biases due to the merger. Doing this also requires excising rows and columns from $\boldsymbol{\Sigma}_d$, thereby breaking the Toeplitz form of the matrix and the corresponding approximations. The inverse of the truncated covariance matrix $\boldsymbol{\Sigma}_{d,tr}$ is evaluated in PyCBC using “gating and in-painting” (explained in detail in [12] and Appendix B), allowing for easy calculation of the first term of Eq. (4).

The second term, however, requires the calculation of $\log \det \boldsymbol{\Sigma}_{d,tr}$, which is a notoriously expensive numerical problem. The most widely used matrix decomposition methods are $\sim O(N^3)$ [16], which become impractical for GW time series such as GW150914, where $N \sim 10^4$. This problem is amplified when varying sky position and t_c in a gated analysis. The start and end times of the gate must be converted between the geocentric and individual detector frames. The time conversions directly depend on the time of merger t_c and the sky location of the event relative to the detectors. Therefore, when varying sky location and t_c , the gate times in the detector frames will also vary. Subsequently, the elements of $\boldsymbol{\Sigma}_{d,tr}$ will vary with each unique sky location and t_c value. Analyses that marginalize over these parameters require recalculating $\log \det \boldsymbol{\Sigma}_{d,tr}$ on multiple steps of the sampler, which would be impractical using numerical decomposition methods. This is not an issue if the sky location and t_c are fixed, since in that case $\det \boldsymbol{\Sigma}_{d,tr}$ will not change throughout the analysis (and in fact can be ignored, as it amounts to a constant normalization term).

In order to marginalize over sky location and t_c it is necessary to have a fast method for calculating $\det \boldsymbol{\Sigma}_{d,tr}$. We find that $\log \det \boldsymbol{\Sigma}_{d,tr}$ is strongly linearly correlated to the length of the power spectral density (PSD) of the data, which is equivalent to the number of rows and columns in $\boldsymbol{\Sigma}_{d,tr}$. Figure 1 shows this relationship using $\log \det \boldsymbol{\Sigma}_{d,tr}$ values calculated for various gate sizes applied to GW150914 data. Using a least squares fit to the points, the correlation coefficient was calculated as $0.999 < R^2 < 1$, which for a sample of eight points implies a highly significant correlation [17].

We also find that the position of the gate in the time series has no significant effect on the value of $\log \det \boldsymbol{\Sigma}_{d,tr}$.




FIG. 1. Plot of $\log \det \boldsymbol{\Sigma}_{d,tr}$ versus PSD length for GW150914 in the Hanford detector. Blue points represent determinants calculated exactly using SciPy functions [18], while the black dashed line is a least squares linear fit. The correlation coefficient of $0.999 < R^2 < 1$ indicates a highly significant linear correlation between $\log \det \boldsymbol{\Sigma}_{d,tr}$ and the size of $\boldsymbol{\Sigma}_{d,tr}$.

The maximum range over which the determinant values varied was $\sim 10^{-6}$. However, since these values were generally of order $\sim 10^6$, this represents a maximum fractional change of $\sim 10^{-12}$ over the entire time domain. While gate position does have a minor effect on determinant value, the relative effect is so small that the value is well-approximated as a constant for a static gate length.

Using these two facts, we calculate the normalization term in Eq. (4) using a linear interpolation based solely on the size of $\boldsymbol{\Sigma}_{d,tr}$. Specifically, a least squares linear fit is generated using the determinant values for the full matrix and three differently-sized truncated matrices. The full matrix determinant is calculated using its approximated eigenvalues, while the truncated determinants are calculated numerically using SciPy [18]. The determinant can then be calculated on each step through linear interpolation based on the size of the truncated covariance matrix.

III. NEED FOR JOINT ANALYSES

The strategy outlined in Sec. II fixes the technical hurdle of calculating the likelihood when the gate length and position are varied. Even with that, however, we find that it is not possible to only model a portion of a signal—whether it be the postmerger or the premerger—if t_c or sky location are uncertain. This is because a larger likelihood can generally be obtained if the gate is shifted such that it removes as much of a signal as possible.

From Eq. (4) it is evident that the likelihood is maximized as $s_d - \mathbf{h}_d(\boldsymbol{\theta}) \rightarrow 0$. When analyzing the entire time series this can only occur if the template $\mathbf{h}_d(\boldsymbol{\theta})$ is similar to the signal that exists in the data (as desired). However, if a variable gate is involved, then a large likelihood can also be obtained if the gate is shifted such that it excises most

(or all) of the signal in the data. The template then only needs to match the remaining noise. Since the noise manifold is larger than the signal manifold, a larger volume of prior space will generally be able to match the noise than the volume that matches the signal. For example, when using a quasinormal mode (QNM) model, if the gate is shifted such that it excises the entire signal then one only needs to reduce the amplitude of the QNM template below noise level to get a relatively large likelihood. The same likelihood would then be obtained regardless of what the other parameters describing the template are set to. The end result is that the posterior probability will favor excising the entire signal even if the template that is being used can match the signal.

Specifically, in *premerger-only* analyses, the gate will be shifted to as *early* a time that is possible while in *postmerger-only* analyses, the gate will be shifted to as *late* as possible. These phenomena are illustrated in the middle row of Fig. 2. Shown are the results from separate premerger and postmerger analyses of a simulated signal in which the gate time is varied due to time and sky-location marginalization. The maximum likelihood template in the premerger-only analysis has a t_c roughly 25 ms earlier

than what was injected. This significantly shortens the premerger signal, leading to inaccurate measurements in parameters such as total spin. Meanwhile, the maximum likelihood template for the postmerger-only analysis had a ringdown start time $\sim 10M_\odot$ later than the injected value, well into the regime where the signal is expected to be noise dominated. This leads to a “prior-in prior-out” posterior—a roughly uniform distribution across parameter space—as the signal model can fit any arbitrary model to late-time noise.

Fundamentally, this problem is due to the gated signal model not being the appropriate model for the observed data. Excising data from the analysis is mathematically equivalent to assuming that the excised data is Gaussian noise and marginalizing over all possible realizations of it. A property of multivariate Gaussian distributions is that marginalizing over a subset of dimensions yields another multivariate Gaussian distribution with the marginalized dimensions excised from the covariance matrix. This is exactly the same form as Eq. (3); in our case, each time sample is a dimension in the multivariate Gaussian. The problem is the excised times are *not* Gaussian noise. They contain a signal, albeit a portion of the signal that we want

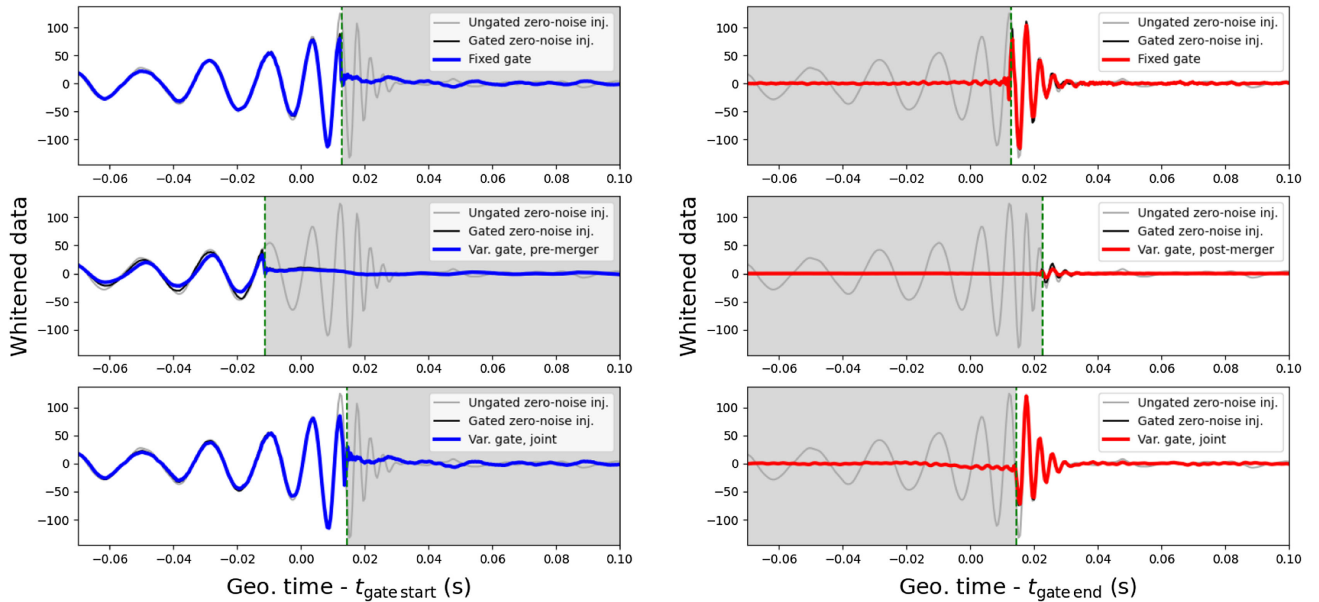


FIG. 2. Comparison of waveform templates fitting to a simulated signal similar to GW150914 in the Hanford detector. The simulated signal (“injection”) is generated using the maximum likelihood values from a 4-OGC analysis of GW150914 [19] and added to zero noise. The ungated zero-noise injection is plotted with a gray line on each set of axes. The plots on the left (right) contain maximum likelihood premerger (postmerger) waveform templates plotted with blue (red) lines. In each analysis, a 1-second gate is applied to the data, represented by a shaded gray region and a dashed green line representing the start (end) of the gate. The data with this gate applied is shown with a black line. The horizontal axes represent the time in seconds relative to the geocentric gate start (end) time. The top plots show the waveform templates for analyses where the sky location was fixed to the same values as the injection. The center plots show the templates for analyses where the sky location was allowed to vary and the pre and postmerger signals were modeled separately. Notice that the premerger (postmerger) gate is significantly earlier (later) than the injected values shown in the corresponding fixed sky location templates. The bottom plots show the templates from an analysis with variable sky location where the pre and postmerger signals were modeled simultaneously. Besides sky location and t_c , the parameter distributions were independent of each other. The resultant gate positions are similar to those observed in the fixed sky location analyses.

to ignore. Consequently, a gated signal model is not representative of the data.

This issue could be mitigated by modifying the prior to exclude the portion of parameter space that matches the noise while keeping the portion that matches the signal. With a QNM template this would mean setting a lower bound on the amplitude such that it is above noise level. However, modifying the prior in this manner is challenging to do in an unbiased way, as it would require *a priori* knowledge of both the noise and signal properties. It also does not solve the fundamental problem that the gated signal model is not representative of the data.

We resolve this issue by simultaneously yet independently modeling both the pre and postmerger signals. Each domain is treated as a separate gated analysis. The signal parameters used in each domain are independent of each other, except for a set of common parameters $\boldsymbol{\vartheta}_{\text{com}}$. In our case the common parameters are the right ascension α , declination δ , and geocentric coalescence time t_c . These parameters determine the coalescence time t_c^d in each detector d , which the gate time depends on. The likelihood is calculated in each domain, then combined to get the overall likelihood.

Explicitly, our algorithm to calculate the likelihood for a given set of parameter values $\boldsymbol{\vartheta} = \{\boldsymbol{\vartheta}_{\text{insp}}, \boldsymbol{\vartheta}_{\text{rd}}, \boldsymbol{\vartheta}_{\text{com}}\}$ is as follows:

- (1) Generate the premerger (“inspiral”) template h^{insp} using parameters $\boldsymbol{\vartheta}_{\text{insp}}$;
- (2) Project the template into each detector’s frame using $\boldsymbol{\vartheta}_{\text{com}} = \{\alpha, \delta, t_c\}$;
- (3) Excise times after t_c^{det} by gating and in-painting the residual $s_d - h_d^{\text{insp}}(\boldsymbol{\vartheta}_{\text{insp}}, \boldsymbol{\vartheta}_{\text{com}})$ using a 1s gate that spans $[t_c^d, t_c^d + 1)$;
- (4) Calculate the premerger likelihood $p(s_{\text{net}} | \boldsymbol{\vartheta}_{\text{insp}}, \boldsymbol{\vartheta}_{\text{com}}, h^{\text{insp}})$ via Eq. (4);
- (5) Repeat steps 1–4 for the postmerger (“ringdown”) template h^{rd} to get the postmerger likelihood $p(s_{\text{net}} | \boldsymbol{\vartheta}_{\text{rd}}, \boldsymbol{\vartheta}_{\text{com}}, h^{\text{rd}})$. However, for the ringdown the gate spans $[t_c^d - 1, t_c^d)$; i.e., it ends at t_c^d whereas it starts at t_c^d for the premerger. The data used in each domain is thereby (nearly) mutually exclusive (see Appendix B for more details);
- (6) The total likelihood is then

$$\begin{aligned}
 & p(s_{\text{net}} | \boldsymbol{\vartheta}, h) \\
 &= p(s_{\text{net}} | \boldsymbol{\vartheta}_{\text{insp}}, \boldsymbol{\vartheta}_{\text{com}}, h^{\text{insp}}) p(s_{\text{net}} | \boldsymbol{\vartheta}_{\text{rd}}, \boldsymbol{\vartheta}_{\text{com}}, h^{\text{rd}}). \quad (5)
 \end{aligned}$$

We fix the issue of the gate trying to excise the signal by using this hierarchical likelihood. The two domains offset each other; in order for the postmerger gate to shift to later times the premerger template must match more of the signal, and vice versa. This also addresses the fundamental issue with the gated signal model highlighted above; our

global model for the entire dataset now contains a non-Gaussian element (the other domain’s signal model) in each domain’s excised region. Note also that no coupling occurs across the domain boundaries due to the whitening filter.

The bottom plots in Fig. 2 show the maximum likelihood waveform templates from an analysis with this configuration. Besides sky location and t_c , both models generated independent parameter measurements. The resultant waveform templates closely match those obtained by fixing sky location and t_c to the injected values (top row of Fig. 2); the erroneous gate motion observed in the pre and postmerger-only analyses is no longer present.

IV. METHODS

We performed eight analyses in total. Each analysis was differentiated by the waveform used, postmerger approximant, and whether or not sky location and t_c were marginalized over. All other aspects of the analyses were kept constant. Both the pre and postmerger signals in all models utilized the PyCBC model GatedGaussianMargPol, which inherits the normalization protocols described in Sec. II. The DYNESTY sampler [20] was used to generate posterior distributions. The samplers in these analyses used 4000 live points to ensure the convergence of each model (see Appendix C).

Half of the analyses used the original data of GW150914 with a sample rate of 2048 Hz obtained from the Gravitational Wave Open Science Center [21]. To validate these results, we repeat each run on a simulated signal in zero noise. The simulated signal was generated using the IMRPhenomXPHM waveform approximant [22] with the maximum likelihood parameters for GW150914 from Ref. [19].

All template models utilized IMRPhenomXPHM (abbreviated here on as IMR) to model the premerger signal of the waveform. Half of the models used this IMR approximant to model the postmerger signal. The other half utilized a QNM approximant to model the postmerger section of the waveform. Table I lists the sampled parameters and priors used in the premerger models, and Table II lists the same for the postmerger models. The QNM approximant was configured such that the ringdown was composed of a dominant (2, 2, 0) mode and a subdominant (2, 2, 1) mode as proposed by [23]. All models apply the ringdown model starting at merger time t_c . The priors of the QNM postmerger models were restricted such that the (2, 2, 0) mode contribution to the postmerger signal-to-noise ratio (SNR) was at least 2. This condition was imposed to ensure that the dominant mode was present in the model, preventing possible “label switching” that may occur due to the (2, 2, 1) mode erroneously matching to the (2, 2, 0) mode in the signal.

Only half of the models allowed for sky location and t_c to vary. The other half fixed these parameters to nominal values for GW150914 to replicate previous works. Specifically, the fixed parameter analyses set $\alpha = 1.95$, $\delta = -1.27$, and

TABLE I. Varied parameters in IMR models and their associated prior distributions. The subscript (1/2) indicates that the same prior was used for the primary and secondary masses. The third column indicates the sampling method for each prior. Parameters listed in this column indicate uniform sampling over those parameters rather than what is listed in column 1. (For example, M_{chirp} is sampled using uniform priors for M_1 and M_2 .)

Parameter	Description	Prior distribution	Prior range
t_c	Coalescence time	Uniform	$1126259462.43 + [-0.05, 0.05]$ s
α	Right ascension	Uniform	$[0, 2\pi]$
δ	Declination	Sine angle	$[-\pi/2, \pi/2]$
i	Inclination	Sine angle	$[0, \pi]$
M_{chirp}	Source frame chirp mass	M_1, M_2	$[23, 42]M_\odot$
q	Mass ratio M_1/M_2	M_1, M_2	$[1, 4]$
$\chi_{a,(1/2)}$	Spin magnitude	Uniform	$[0, 0.99]$
$\chi_{\theta,(1/2)}$	Spin polar angle	Solid angle	$[0, 2\pi]$
$\chi_{\phi,(1/2)}$	Spin azimuthal angle	Solid angle	$[0, \pi]$
ϕ_c	Reference phase	Uniform	$[0, 2\pi]$
V_C	Comoving volume	Uniform	$[5000, 92918664351]$ Mpc ³

TABLE II. Varied parameters in ringdown models and their associated priors. The t_c , i , and sky location priors used in ringdown analyses were identical to those shown in Table I. The amplitude priors indicate the amplitudes of the corresponding quasinormal modes at the start of the ringdown model (i.e., at the merger). The third column indicates the sampling method for each prior. Here, \log_{10} indicates a uniform distribution over the base 10 logarithm of the given parameter.

Parameter	Description	Prior distribution	Prior range
M_f	Source frame final mass	Uniform	$[10, 200]M_\odot$
χ_f	Final spin	Uniform	$[-0.99, 0.99]$
A_{220}	Initial (2, 2, 0) mode amplitude	\log_{10}	$[10^{-25}, 8 \times 10^{-17}]$
ϕ_{220}	(2, 2, 0) mode phase	Uniform	$[0, 2\pi]$
A_{221}/A_{220}	Initial (2, 2, 1) mode amplitude (as ratio of A_{220})	Uniform	$[0, 5]$
ϕ_{221}	(2, 2, 1) mode phase	Uniform	$[0, 2\pi]$

$t_c = 1126259462.408$, in accordance with the maximum likelihood values in [23].

As a preliminary test, the sky location posteriors of the analyses with variable sky/ t_c parameters on real data were plotted. As seen in Fig. 3, the analyses were able to recover most of the posterior for the full IMR analysis conducted in [19].

V. AREA THEOREM

The simplest area theorem test is to compare the progenitor horizon areas A_1 and A_2 to the remnant horizon area A_f to check that

$$A_1 + A_2 \equiv A_i \leq A_f. \quad (6)$$

In natural units ($G = c = 1$) the area of each black hole is given by [1]

$$A = 8\pi M^2 \left(1 + \sqrt{1 - \chi^2}\right), \quad (7)$$

where M the black hole's mass and $\chi = J/M^2$ is its dimensionless spin.

A more robust test can be performed by comparing the measured change in area to the expected change in area [2]

$$H = \frac{A_{f,\text{measured}} - A_i}{A_{f,\text{expected}} - A_i}. \quad (8)$$

Here, $A_{f,\text{measured}}$ is the area of the final black hole inferred from the postmerger analysis and A_i is the sum of the initial areas inferred from the premerger analysis. The expected area of the final black hole $A_{f,\text{expected}}$ is derived from the initial masses and spins measured from the premerger analysis. These are converted into an estimated final mass and spin using fits from numerical relativity, which is then converted into an area via Eq. (7).

Since $A_{f,\text{expected}}$ is evaluated with numerical relativity using A_1 and A_2 , the denominator of Eq. (8) is positive definite. Therefore, if $H < 0$, then $A_{f,\text{measured}} < A_i$ and the area theorem is violated. The confidence interval in favor of the area theorem is the fraction of the posterior for




FIG. 3. Sky position posteriors for GW150914 analyses. The colored contours represent analyses done using the normalization methods described in Sec. II. The blue contour represents the analysis with a QNM postmerger model, and the red contour depicts the analysis with an IMR postmerger model. The black contour represents the posterior from a full IMR analysis conducted in [19]. All contours mark the 90th percentile of each distribution.

which $H > 0$. More robustly, $H = 1$ indicates that signal is consistent with GR specifically, not just the more generic class of theories that satisfy the area increase law.

Figure 4 shows the posteriors of the final mass M_f and final spin χ_f of the QNM postmerger models on real data. Specifically, the M_f and χ_f posteriors from the analysis with variable sky/ t_c parameters are compared with the corresponding posteriors from fixed gate analysis. All posteriors contain within them the distribution from a full IMR analysis from 4-OGC [19]. Furthermore, both posteriors from the variable sky/ t_c run overlap almost completely with their corresponding fixed parameter posteriors.

A similar plot is shown in Fig. 5, except with the analyses that used an IMR postmerger model on real data. Again, the premerger and postmerger posteriors from the analysis with variable sky/ t_c parameters are compared to their fixed counterparts as well as the full IMR posterior from [19]. Notably, while still in very close agreement, the variable sky/ t_c posteriors have a greater discrepancy from their fixed counterparts than seen in Fig. 4. Specifically, the postmerger posterior tends towards slightly higher masses and spins, while the premerger posterior includes lower spin values.

Directly comparing the premerger posteriors in Figs. 4 and 5 shows a slight discrepancy in final mass and spin estimates. Namely, the analysis with a QNM postmerger model tends towards higher final mass and spin estimates than the IMR postmerger model. This indicates minor




FIG. 4. Final mass and final spin posteriors of GW150914 analyses using a QNM postmerger approximant and real waveform data. All masses are in the detector frame. Contours represent the 90th percentile of each distribution. The darker colored lines correspond to the model with variable sky location and t_c , and the lighter lines represent the model where these parameters were fixed. Blue lines correspond to postmerger results, while red lines correspond to premerger results, obtained using numerical relativity with the IMR parameters. The innermost black contour represents the 4-OGC posterior from [19].




FIG. 5. Final mass and final spin posteriors of GW150914 analyses using an IMR postmerger approximant and real waveform data. All masses are in the detector frame. Contours represent the 90th percentile of each distribution. The darker colored lines correspond to the model with variable sky location and t_c , and the lighter lines represent the model where these parameters were fixed. Blue lines correspond to postmerger results, while red lines correspond to premerger results, obtained using numerical relativity with the IMR parameters. The innermost black contour represents the 4-OGC posterior from [19].

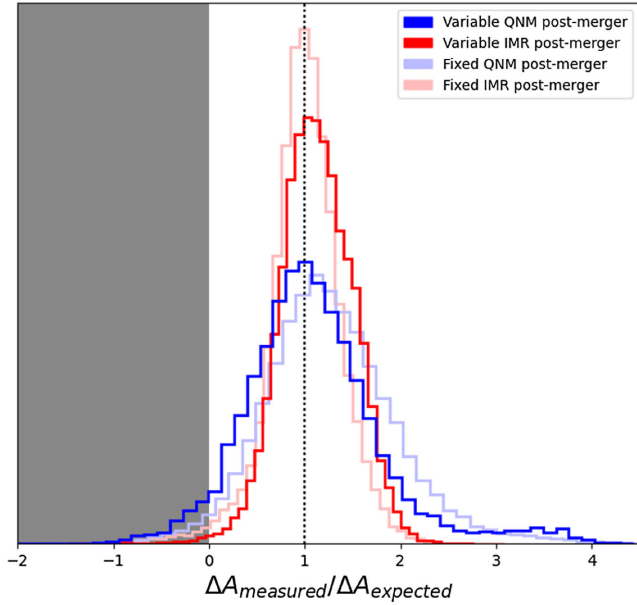


FIG. 6. Posteriors of $H = (A_{f,\text{measured}} - A_i) / (A_{f,\text{expected}} - A_i)$ for the GW150914 analyses with real waveform data. Blue histograms correspond to QNM postmerger models, while red lines correspond to IMR postmerger models. Darker histograms correspond to models with variable sky/t_c parameters, while lighter histograms correspond to models where these parameters are fixed to the maximum likelihood values from [19]. The expected value $H = 1$ is shown with a vertical black dotted line. Points in the shaded region correspond to $H < 0$ and therefore disagree with the area theorem.

coupling between the pre and postmerger models caused by their shared sky and time parameters. However, as made evident by these two plots, this is not a major effect, and both premerger models maintain strong consistency with the full IMR posterior.

Figure 6 compares the H posteriors between the four analyses of the real GW150914 data. All posteriors agree with the expected value $H = 1$, with each distribution peaking near this value. While the QNM postmerger posteriors are almost in exact agreement with each other, there is a slight discrepancy between the IMR postmerger posteriors. Specifically, the fixed sky/t_c posterior has a much sharper peak at a lower value than the variable gate counterpart. This may be explained by Fig. 5, where the fixed IMR postmerger posterior was observed to contain lower masses and spins than the variable gate model. This could lead to the slight bias to lower H values seen here.

Table III summarizes the H credible intervals and area theorem confidence intervals for all eight analyses. All credible intervals are in agreement with the expected value $H = 1$. Additionally, all analyses of the real GW150914 signal have $H > 0$ at greater than 95% confidence, indicating very high agreement with the area theorem. Results from the analysis of the simulated signal are largely consistent with these results (see Appendix C for equivalent figures).

The area theorem confidence interval of any given analysis appears to be uncorrelated to whether or not the sky location and coalescence time were allowed to vary. However, IMR postmerger models tended to have slightly higher agreement with the area theorem than corresponding QNM postmerger models. This pattern may be explained by trends in the pre and postmerger mass and spin results. The IMR postmerger models in Fig. 5 had postmerger measurements that were either in agreement with or greater than the corresponding premerger estimates. This corresponds to a very high concentration of H measurements approximating or exceeding 1, with very few points skewed to negative H values. Conversely, the postmerger posteriors shown in Fig. 4 have significant amounts of points with higher and lower mass and spin values than their premerger

TABLE III. Summary of posteriors for $H = (A_{f,\text{measured}} - A_i) / (A_{f,\text{expected}} - A_i)$ from various GW150914 analyses. The first column lists the properties of each model, namely whether or not sky location and t_c were allowed to vary, the postmerger model used, and the waveform data used. The approximant IMRPhenomXPHM [22] was used to model the postmerger signals of analyses labeled “IMR postmerger”, as were all premerger signals. Analyses with injected waveforms used a zero-noise injection of the GW150914 waveform as input data. The second column lists the median and the 90% credible interval for each H posterior. The third column lists the percentage of posterior points for which $H > 0$, thereby agreeing with the area theorem.

GW150914 analysis	Credible interval	$P(H > 0)$
Variable sky/t_c , QNM postmerger, real waveform	$1.0_{-1.0}^{+1.4}$	95.4%
Variable sky/t_c , IMR postmerger, real waveform	$1.1_{-0.6}^{+0.6}$	99.5%
Fixed sky/t_c , QNM postmerger, real waveform	$1.2_{-1.0}^{+1.2}$	97.4%
Fixed sky/t_c , IMR postmerger, real waveform	$1.0_{-0.6}^{+0.6}$	98.4%
Variable sky/t_c , QNM postmerger, zero-noise injection	$0.9_{-1.3}^{+1.5}$	89.1%
Variable sky/t_c , IMR postmerger, zero-noise injection	$0.9_{-0.6}^{+0.6}$	98.4%
Fixed sky/t_c , QNM postmerger, zero-noise injection	$1.1_{-0.9}^{+0.9}$	97.3%
Fixed sky/t_c , IMR postmerger, zero-noise injection	$1.0_{-0.5}^{+0.5}$	98.8%

counterparts. This leads to a wider distribution of H values to both higher and lower values, allowing for a higher concentration of negative measurements.

VI. CONCLUSIONS

This paper presented a novel method of marginalizing over sky location and coalescence time when performing a pre or postmerger analysis, which allows for a full accounting of uncertainties in parameter estimates. Tests of the area theorem were conducted using data from GW150914 using this method. It was found that the data for this event agrees very well with the area theorem regardless of whether sky position and t_c were allowed to vary. The only noticeable changes in agreement in the area theorem were caused by the postmerger approximant used, though in general this was a very minor change.

While we focus on tests of the area theorem here, our method for marginalizing over sky location and t_c applies to any analysis in which only a portion of the signal is modelled. In particular, black hole spectroscopy involves analyzing the postmerger signal using a QNM template in order to perform a test of the no-hair theorem [24–26]. As with previous tests of the area theorem, previous black hole spectroscopy studies have fixed the sky location and t_c for the reasons discussed in Sec. II [23,27,28]. Our finding that a signal model is needed for the entire observable signal applies equally well to black hole spectroscopy, even though the premerger signal is purposely excluded in such analyses. The hierarchical method we develop in Sec. III is therefore equally relevant for marginalizing over sky location and t_c when doing black hole spectroscopy. In Ref. [29] we use our QNM analysis of GW150914 to investigate the evidence for the presence of the (2, 2, 1) mode, which has been hotly contested in the literature [5,23,27,30].

In all our analyses here, we end the premerger analysis when the postmerger analysis begins. This is necessitated by our finding that some model must exist for the entire observable signal when the gate time is allowed to vary. Since there was no gap between our pre and postmerger template, we obtained agreement with the area theorem in excess of 95%. A more rigorous test of the area theorem is to excise the merger from the data, since any biases introduced by including the merger as part of the premerger model would be omitted. References [2,4] did this additional test with fixed sky location and t_c . However, this is not possible when marginalizing over sky location and t_c for the reasons highlighted in Sec. III. Introducing a gap between the pre and postmerger will once again favor points that excise as much of the signal as possible, even when both the pre and postmerger were modeled simultaneously.

Introducing a gap between the pre and postmerger models could be achieved by using three subdomains rather than two: one each for the inspiral, merger, and ringdown. The gate for one domain would start/end at the end/start of the next. For the merger domain, an arbitrary

signal model using wavelets could be used, similar to what Finch and Moore used in Ref. [10]. This would ensure that pre and postmerger parameters are fully unbiased by the merger while ensuring that the merger itself is not arbitrarily gated out. The initial black hole areas could then be estimated using an inspiral model and the final area using a QNM model. This would also be useful in black hole spectroscopy studies involving fundamental angular QNMs. These modes are not expected to become relevant until $\sim 10M$ after merger, necessitating a gap between the merger and the start of the QNM model. We plan to investigate this in a future study.

This research was conducted using PyCBC [31]. Our data is available at [32].

ACKNOWLEDGMENTS

A. C. was supported by funds from the Massachusetts Space Grant Consortium. C. C. acknowledges support from NSF Grant No. PHY-2309356. All computations were performed on Unity, a collaborative, multi-institutional high-performance computing cluster managed by UMass Amherst Research Computing and Data. This research has made use of data or software obtained from the Gravitational Wave Open Science Center [33], a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation, as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. K. A. G. R. A. is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan.

APPENDIX A: LIKELIHOOD FUNCTION DERIVATION

To calculate the posterior from Bayes’ theorem [Eq. (1)], one requires a model for both the signal h and the noise n .

While h can be determined from the linearized EFE wave solution, n requires more statistical considerations.

To start, consider a network of K gravitational wave detectors that each sample at a rate Δt over a total time length T . Defining the total number of samples $N = T/\Delta t$, the full data series s_{net} can be expressed as a series of K N -dimensional vectors s^K such that $s^K = \{s_0^K, s_1^K, \dots, s_N^K\}$ and $s_{\text{net}} = \{s^1, s^2, \dots, s^K\}$. To simplify calculations, assume that the signal is zero such that $s_{\text{net}} = \mathbf{n}$. Additionally, assume that noise model is a stochastic Gaussian distribution and is uncorrelated between detectors. Under these assumptions, the noise likelihood function is

$$p(s_{\text{net}}|n) = \frac{\exp\left[-\frac{1}{2}\sum_{d=1}^K s_d^T \Sigma_d^{-1} s_d\right]}{[(2\pi)^{NK} \prod_{d=1}^K \det \Sigma_d]^{1/2}}, \quad (\text{A1})$$

where Σ_d is the covariance matrix of the noise model for detector d , defined using the ensemble average,

$$\Sigma_d[j, k] = \langle s_d[j] s_d[k] \rangle. \quad (\text{A2})$$

(By dropping the d subscripts, one obtains an equivalent expression for the full covariance matrix and data. We do so in the following steps for brevity.) This is the exact expression of the likelihood function for noise. However, this function is infeasible to calculate analytically due to the inverse covariance matrix in the numerator.

To do this, one may expand the covariance matrix definition in Eq. (A2) as follows, defining $\Delta_{kj} = k - j$:

$$\begin{aligned} \Sigma[j, k] &= \langle s[j] s[k] \rangle \\ &= \langle s[j] s[\Delta_{kj} + j] \rangle \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=0}^{n-1} s^l[j] s^l[\Delta_{kj} + j], \end{aligned} \quad (\text{A3})$$

where in the last step the ensemble average is written out fully.

In general, this expression is dependent on time $t_j = j\Delta t$ and displacement $\tau_{kj} = \Delta_{kj}\Delta t$. However, one can make the assumption that the noise is wide sense stationary, where the mean and variance are both constant in time. Under this assumption, any constant can be added to the indices in Eq. (A2) to obtain the same result. This makes Σ symmetric [since the factors in Eq. (A3) commute] and Toeplitz (since the elements along the diagonals are equal) [34]. Additionally, since $\Sigma[0, \Delta_{kj}] = \Sigma[-\Delta_{kj}, 0] = \Sigma[0, -\Delta_{kj}]$, the elements of Σ are even functions of Δ_{kj} .

Additionally, one can assume that the data is ergodic, meaning that new realizations of s are obtained via time. Under this assumption and the properties of the elements of Σ , the ensemble averages in Eq. (A3) can be replaced with time averages,

$$\begin{aligned} \Sigma[j, k] &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=0}^{n-1} s^l[0] s^l[\Delta_{kj}] \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{l=0}^{n-1} s^l[l] s^l[\Delta_{kj} + l] \\ &= \lim_{n \rightarrow \infty} \frac{1}{2n} \sum_{l=-n}^{n-1} s^l[l] s^l[\Delta_{kj} + l] \\ &= \frac{1}{2} R_{ss}((k-j)\Delta t). \end{aligned} \quad (\text{A4})$$

The last step defines the autocorrelation function $R_{ss}(\tau)$, which describes the correlation between points in the time series s . If $R_{ss}(\tau)$ goes to zero in some finite time τ_{max} , then all diagonals with $|\Delta_{kj}| > \text{floor}(\tau_{\text{max}}/\Delta t) = \Delta_{\text{max}}$ will equal zero. This is similar to the form of a circulant matrix \mathbf{C} , a special case of a Toeplitz matrix where each row is a right-cycle permutation of the same vector (in this case, s). The eigenvectors of circulant matrices are well-known [34],

$$\mathbf{u}_p[k] = \frac{1}{\sqrt{N}} \exp(-2\pi i k p / N). \quad (\text{A5})$$

This generally is not true for Toeplitz matrices, but one may take advantage of Eq. (A5) by recognizing that the matrix described by Eq. (A4) asymptotes to a circulant matrix for large N ,

$$\lim_{N \rightarrow \infty} |\Sigma - \mathbf{C}| = \mathbf{0}. \quad (\text{A6})$$

Therefore, the eigenvalues λ_p of Σ can be evaluated using the usual eigenvalue equation as long as $\Delta_{\text{max}} \ll N/2$,

$$\Sigma \mathbf{u}_p \approx \lambda_p \mathbf{u}_p. \quad (\text{A7})$$

Therefore, using the fact that Σ is symmetric and $R_{ss}(l)$ is even,

$$\begin{aligned} \lambda_p &= \frac{1}{2} \text{Re} \left\{ \sum_{l=-N/2}^{N/2} R_{ss}(l) \exp(-2\pi i p l / N) \right\} \\ &= \frac{1}{2} \text{Re} \left\{ \sum_{l=0}^{N-1} R_{ss}(l) \exp(-2\pi i p l / N) \right\} \\ &= \frac{1}{2} \text{Re} \{ \tilde{R}_{ss}(p) / \Delta t \}, \end{aligned} \quad (\text{A8})$$

where $\tilde{R}_{ss}(p)$ is the discrete Fourier transform of the autocorrelation function,

$$\tilde{R}_{ss}(p) = \Delta t \sum_{k=0}^{N-1} R_{ss}(k) \exp(-2\pi i p k / N). \quad (\text{A9})$$

To simplify Eq. (A8), one may impose the Wiener-Khinchin theorem [35], which defines the power spectral density S_n as the Fourier transform of R_{ss} for a wide-sense stationary stochastic process,

$$\lambda_p = \frac{S_n[p]}{2\Delta t}. \quad (\text{A10})$$

One can then construct the inverse of Σ using the eigenvalues from Eq. (A10) and the eigenvectors of Eq. (A5),

$$\begin{aligned} \Sigma^{-1}[j, k] &= \frac{2\Delta t}{N} \sum_{p=0}^{N-1} \frac{\exp[-2\pi i p(j-k)/N]}{S_n(p)} \\ &= 2\Delta f(\Delta t)^2 \sum_{p=0}^{N/2-1} \frac{e^{-2\pi i p(j-k)/N} + e^{2\pi i p(j-k)/N}}{S_n(p)}, \end{aligned} \quad (\text{A11})$$

using the fact that S_n is symmetric about $N/2$ and defining the sample frequency $\Delta f = 1/T = 1/N\Delta t$. Since s is real and

$$(\Delta t)^2 \sum_{j,k=0}^{N-1} s[j]s[k] (e^{-2\pi i p(j-k)/N} + e^{2\pi i p(j-k)/N}) = 2|\tilde{s}|^2[p], \quad (\text{A12})$$

one can write

$$s^T \Sigma^{-1} s = 4\Delta f \sum_{p=0}^{N/2-1} \frac{2|\tilde{s}|^2[p]}{S_n(p)}. \quad (\text{A13})$$

Finally, if the inner product between two arbitrary vectors \mathbf{a} and \mathbf{b} is defined as

$$\langle \mathbf{a}, \mathbf{b} \rangle = 4\text{Re} \left\{ \Delta f \sum_{p=0}^{N/2-1} \frac{\tilde{\mathbf{a}}[p] \tilde{\mathbf{b}}[p]}{S_n(p)} \right\}, \quad (\text{A14})$$

Eq. (A13) can be written as an inner product,

$$s^T \Sigma^{-1} s = \langle s, s \rangle, \quad (\text{A15})$$

and the likelihood function can be expressed as

$$p(s_{\text{net}}|n) = \frac{\exp[-\frac{1}{2} \sum_{d=1}^K \langle s_d, s_d \rangle]}{[(2\pi)^{NK} \prod_{d=1}^K \det \Sigma_d]^{1/2}}. \quad (\text{A16})$$

APPENDIX B: GATING AND IN-PAINTING

Generally, BBH models do not require the entire waveform to be analyzed at once. For example, the analyses throughout this paper independently examine the pre and postmerger portions of the GW150914 waveform. To maintain independence between the models, any points not corresponding to the respective model (i.e., after t_c for the premerger model, or before t_c for the postmerger model) were excised, or ‘‘gated,’’ from the data. Here, a gate of length M applied starting at a sample a will excise all samples within the range $[a, a + M]$, corresponding to a gate

of time length $t = M\Delta t$. The gated time series will therefore take the form $s_{d,tr} = \{s_0, s_1, \dots, s_a, s_{a+M}, \dots, s_{N-1}, s_N\}$.

By doing this, the simplifications made to derive the likelihood function are no longer valid, since $\Sigma_{d,tr}$ is no longer Toeplitz (and, subsequently, no longer approximately circulant for large N). There are numerical methods to calculate the matrix inverse directly, but in general they can be unstable and time intensive. Therefore, a method known as ‘‘gating and in-painting’’ is employed for gated waveform analyses [12].

First, the method assumes that the noise time series \mathbf{n} is the sum of \mathbf{n}_g , the noise series with the gated times zeroed out, and \mathbf{x} , a time series containing only the gated samples in \mathbf{n} . The goal of in-painting is to solve the following equation in the gated region,

$$\Sigma^{-1}(\mathbf{n}_g + \mathbf{x}) = \mathbf{0}. \quad (\text{B1})$$

If the nonzero elements of \mathbf{x} are such that $(\Sigma^{-1}\mathbf{n})[k] = \mathbf{0}$ for all samples k in the gate $[a, a + M]$, then the inner product $\mathbf{n}^T \Sigma \mathbf{n}$ will be equal for the truncated and raw dataset. Since \mathbf{x} is zero outside of the gate, $\Sigma^{-1}\mathbf{x}$ will form an $M \times M$ Toeplitz matrix containing the $[a, a + M]$ rows and columns of Σ^{-1} . Therefore, Eq. (B1) can be rewritten within the gated region as

$$\Sigma^{-1}\mathbf{x} = -\Sigma^{-1}\mathbf{n}_g, \quad (\text{B2})$$

and adding \mathbf{x} to \mathbf{n}_g will give the same result as truncating \mathbf{n} and Σ . Unlike trying to solve for the inverse directly, this solution is readily found using a Toeplitz solver [18,36]. Given gated data s_g containing some gated signal \mathbf{h}_g , Eq. (B2) can be evaluated using $\mathbf{n}_g = s_g - \mathbf{h}_g$, and the value $\mathbf{x} + s_g - \mathbf{h}_g$ can be used to calculate the likelihood.

The analyses conducted in this paper utilize gating and in-painting to apply a gate starting/ending at t_c to the waveform template. In theory, the gates should extend to the edges of the analysis segment (i.e., the premerger gate starts at the segment start time, and the postmerger gate ends at segment end). However, the in-painting algorithm, which is the dominant cost in our analysis, is $\sim O(M^2)$ operations for a gate of M samples. Although the observable signal is only ~ 0.2 s long [13], we use an analysis segment that is 4 s in duration in order to resolve line artifacts in the PSD. We also use a sample rate of 2048 Hz, in order to fully capture all observable signal power. The in-painting would therefore need to span $M \sim 4096$ samples if the gates were to extend to the beginning/end of the analysis segment. This is computationally expensive; some of our analyses would take ~ 1 month to complete (utilizing 64 CPU cores).

To reduce the computational cost, we instead use a 1-second gate starting (ending) at t_c and ensuring that the remainder of the premerger (postmerger) template \mathbf{h} after (before) the gate is zero. This is equivalent to applying a full

gate to the template when a whitening filter is applied, since the early- and late-time template will be identically zero in both cases.

We do not zero the data s outside of the gates, however, since doing so would introduce additional ringing artifacts at the beginning/end of the analysis segment not accounted for by the in-painting. This means that the pre and postmerger models will share data outside of the gates (namely, greater than 1 second before and after t_c). However, the data is expected to be noise dominated in this region, yielding an insignificant contribution to the likelihood. This can be seen in Appendix C, which portrays the remnant posteriors for analyses using zero-noise waveform injections. Since the data in these analyses contain no noise, the shared data has identically zero contribution to the likelihood. These results are in strong agreement with Figs. 4, 5, and 6, indicating that any effects due to early- and late-time shared noise are indeed negligible.

APPENDIX C: SIMULATION RESULTS

This section contains additional plots showing results from the analysis of the GW150914-like simulated signal in zero noise. Figure 7 shows the M_f and χ_f posteriors for the QNM postmerger models, and Fig. 8 shows the same for the IMR postmerger models. Figure 9 shows the H posteriors for the zero-noise injection models.




FIG. 7. Final mass and final spin posteriors of GW150914 analyses using a QNM postmerger approximant and a zero-noise waveform injection. All masses are in the detector frame. Contours represent the 90th percentile of each distribution. The darker colored lines correspond to the model with variable sky location and t_c , and the lighter lines represent the model where these parameters were fixed. Blue lines correspond to postmerger results, while red lines correspond to premerger results, obtained using numerical relativity with the IMR parameters. The innermost black contour represents the 4-OGC posterior from [19].




FIG. 8. Final mass and final spin posteriors of GW150914 analyses using an IMR postmerger approximant and a zero-noise waveform injection. All masses are in the detector frame. Contours represent the 90th percentile of each distribution. The darker colored lines correspond to the model with variable sky location and t_c , and the lighter lines represent the model where these parameters were fixed. Blue lines correspond to postmerger results, while red lines correspond to premerger results, obtained using numerical relativity with the IMR parameters. The innermost black contour represents the 4-OGC posterior from [19].




FIG. 9. Posteriors of $H = (A_{f,\text{measured}} - A_i) / (A_{f,\text{expected}} - A_i)$ for the GW150914 analyses with zero-noise injected waveform data. Blue histograms correspond to QNM postmerger models, while red lines correspond to IMR postmerger models. Darker histograms correspond to models with variable sky/t_c parameters, while lighter histograms correspond to models where these parameters are fixed to the maximum likelihood values from [19]. The expected value $H = 1$ is shown with a vertical black dotted line. Points in the shaded region correspond to $H < 0$ and therefore disagree with the area theorem.

APPENDIX D: SAMPLER CONVERGENCE

To test the convergence of the sampler we repeated the analyses with 2000 and 4000 live points. Figure 10 shows the M_f and χ_f posteriors for QNM postmerger models on real data with 2000 and 4000 live points, while




FIG. 10. Final mass and spin posteriors for analyses of GW150914 data with a QNM postmerger model and a variable ringdown start time. The blue contour indicates the 90th percentile of an analysis where the DYNESTY sampler used 2000 live points. The orange contour represents the 90th percentile of an analysis with 4000 live points. Besides the number of sampler points used, both models were identical.

Fig. 11 depicts the same for the IMR postmerger models. Both the QNM and IMR models were able to converge to similar distributions regardless of the number of live points used. All sky/ t_c results are reported using 4000 live points.




FIG. 11. Final mass and spin posteriors for analyses of GW150914 data with a IMR postmerger model and a variable ringdown start time. The blue contour indicates the 90th percentile of an analysis where the DYNESTY sampler used 2000 live points. The orange contour represents the 90th percentile of an analysis with 4000 live points. Besides the number of sampler points used, both models were identical.

-
- [1] J. M. Bardeen, B. Carter, and S. W. Hawking, The four laws of black hole mechanics, *Commun. Math. Phys.* **31**, 161 (1973).
 - [2] 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).
 - [3] M. Cabero, C. D. Capano, O. Fischer-Birnholtz, B. Krishnan, A. B. Nielsen, A. H. Nitz, and C. M. Biwer, Observational tests of the black hole area increase law, *Phys. Rev. D* **97**, 124069 (2018).
 - [4] 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).
 - [5] 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](https://arxiv.org/abs/2310.19645).
 - [6] G. Carullo, W. D. Pozzo, and J. Veitch, Observational black hole spectroscopy: A time-domain multimode analysis of GW150914, *Phys. Rev. D* **99**, 123029 (2019).
 - [7] G. Carullo, W. Del Pozzo, and J. Veitch, pyRing: A time-domain ringdown analysis Python package, <https://git.ligo.org/lscsoft/pyring> (2023).
 - [8] M. Isi and W. M. Farr, maxisi/ringdown: Initial ringdown release, Zenodo, [10.5281/zenodo.5094068](https://zenodo.org/record/5094068) (2021).
 - [9] M. Isi and W. M. Farr, Analyzing black-hole ringdowns, [arXiv:2107.05609](https://arxiv.org/abs/2107.05609).
 - [10] E. Finch and C. J. Moore, Searching for a ringdown overtone in GW150914, *Phys. Rev. D* **106**, 043005 (2022).
 - [11] 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).
 - [12] 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. D* **104**, 063034 (2021).
- [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] D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial* (Oxford University Press, Oxford, 2006).
- [15] L. S. Finn, Detection, measurement, and gravitational radiation, *Phys. Rev. D* **46**, 5236 (1992).
- [16] V. Pan, Y. Yu, and C. Stewart, Algebraic and numerical techniques for the computation of matrix determinants, *Comput. Math. Appl.* **34**, 43 (1997).
- [17] J. R. Taylor, *An Introduction to Error Analysis: The Study of Uncertainties in Physical Measurements*, 2nd ed. (University Science Books, Sausalito, CA, 1996).
- [18] P. Virtanen *et al.*, SciPy1.0: Fundamental algorithms for scientific computing in PYTHON, *Nat. Methods* **17**, 261 (2020).
- [19] A. H. Nitz, S. Kumar, Y. Wang, S. Kasta, 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).
- [20] J. S. Speagle, DYNESTY: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [21] R. Abbott *et al.*, Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, *SoftwareX* **13**, 100658 (2021).
- [22] 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).
- [23] 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).
- [24] C. V. Vishveshwara, Scattering of gravitational radiation by a Schwarzschild black-hole, *Nature (London)* **227**, 936 (1970).
- [25] S. Chandrasekhar and S. L. Detweiler, The quasi-normal modes of the Schwarzschild black hole, *Proc. Roy. Soc. Lond. A* **344**, 441 (1975).
- [26] O. Dreyer, B. J. 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).
- [27] R. Cotesta, G. Carullo, E. Berti, and V. Cardoso, Analysis of ringdown overtones in GW150914, *Phys. Rev. Lett.* **129**, 111102 (2022).
- [28] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), Tests of general relativity with GWTC-3, arXiv:2112.06861.
- [29] A. Correia, Y.-F. Wang, J. Westerweck, and C. D. Capano, companion paper, Low evidence for ringdown overtone in GW150914 when marginalizing over time and sky location uncertainty, *Phys. Rev. D* **110**, 041501 (2024).
- [30] M. Isi and W. M. Farr, Revisiting the ringdown of GW150914, arXiv:2202.02941.
- [31] A. Nitz *et al.*, GWASTRO/PyCBC: v2.3.3 release of PyCBC (2024).
- [32] A. Correia and C. Capano, Sky marginalization in black hole spectroscopy (v1), Zenodo, 10.5281/zenodo.10493247 (2024).
- [33] gwosc.org.
- [34] R. Gray, Toeplitz and circulant matrices: A review, *Found. Trends Commun. Inf. Theory* **2**, 155 (2001).
- [35] C. Chatfield, *The Analysis of Time Series: An Introduction*, 6th ed. (CRC Press, Florida, USA, 1996).
- [36] C. R. Harries *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).