

**Fortifying the characterization of binary mergers in LIGO data**Tyson B. Littenberg,<sup>1</sup> Michael Coughlin,<sup>2</sup> Benjamin Farr,<sup>1</sup> and Will M. Farr<sup>1</sup><sup>1</sup>*Department of Physics and Astronomy and Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, 2131 Tech Drive, Evanston Illinois 60208, USA*<sup>2</sup>*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, United Kingdom*

(Received 5 August 2013; published 25 October 2013)

The study of compact binary inspirals and mergers with gravitational wave observatories amounts to optimizing a theoretical description of the data to best reproduce the true detector output. While most of the research effort in gravitational wave data modeling focuses on the gravitational waveforms themselves, here we will begin to improve our model of the instrument noise by introducing parameters that allow us to determine the background instrumental power spectrum while simultaneously characterizing the astrophysical signal. We use data from the fifth LIGO science run and simulated gravitational wave signals to demonstrate how the introduction of noise parameters results in the resilience of the signal characterization to variations in an initial estimation of the noise power spectral density. We find substantial improvement in the consistency of Bayes factor calculations when we are able to marginalize over uncertainty in the instrument noise level.

DOI: [10.1103/PhysRevD.88.084044](https://doi.org/10.1103/PhysRevD.88.084044)

PACS numbers: 04.30.-w, 04.80.Nn

**I. INTRODUCTION**

The cornerstone sources of gravitational radiation targeted by the LIGO/Virgo observatories are mergers of binary systems comprised of compact stellar remnants such as neutron stars and/or black holes. Signal processing for ground-based gravitational wave detectors is a daunting data mining problem as merger signals are rare events of short duration, lasting tens to hundreds of seconds, buried in years worth of instrument noise. Extracting astrophysical information from this mountain of data requires an elaborate analysis pipeline, which in turn is dependent on accurate theoretical predictions for the signals being pursued and on the instrument noise with which those signals compete.

While noise characterization is important for any measurement, it is especially vital for gravitational wave (GW) detection because most events will contribute relatively little power to the data. In this noise-dominated regime, small changes to the instrument background result in fractionally significant changes to the relative strength of the signal (or signal-to-noise ratio) and, consequently, inferences that can be made about the system parameters. For a detailed theoretical account of the role noise plays in the process of drawing inferences from the data, see Vallisneri's 2011 work [1]. For studies investigating the impact of real LIGO/Virgo noise on astrophysical inferences using simulated gravitational wave signals, see Refs. [2–4].

In Gaussian noise, the data analysis strategy for compact mergers relies on Weiner matched filtering [5,6], which, in turn, assumes an accurate model for the instrument data. In the matched-filtering paradigm, we first require gravitational wave simulations, or waveforms,  $h$ , which can be computed for arbitrary astrophysical parameters  $\theta$  over the

prior volume of the search. Then, assuming the data  $d$  are comprised of a gravitational wave signal and additive Gaussian noise  $n$ , a matched-filtering statistic is adopted such as the signal-to-noise ratio (SNR)  $\rho \equiv (d|h)/\sqrt{\langle h|h \rangle}$  or chi-squared residual,

$$\chi^2 \equiv (d - h|d - h), \quad (1)$$

where we have used the standard noise weighted inner product

$$(a|b) \equiv 2 \int_0^\infty \frac{a(f)b(f)^* + a^*(f)b(f)}{S_n(f)} df \quad (2)$$

with  $S_n(f)$  representing the one-sided noise spectral density, i.e., the expectation value of the instrumental noise's Fourier power.

In the analysis of LIGO/Virgo data, the search for candidate signals is performed by convolving the simulated instrument response to a grid of precomputed trial waveforms  $h(\theta)$ , or templates, with the data computing a statistic that is maximized when the template identically matches the signal, and keeping events, or “triggers,” which exceed a threshold for statistical significance [7,8]. The statistic employed in GW searches is similar in principle to  $\rho$  but modified using statistical studies of the data to satisfy the Neyman–Pearson criterion of minimizing the false dismissal rate for a fixed false alarm probability. Candidate events that exceed the statistical threshold and pass consistency checks across the network are further analyzed using stochastic samplers based on the nested sampling [9–12] and/or Markov chain Monte Carlo (MCMC) [2,3,13,14] algorithms to characterize the probability density function of  $\theta$  from which inferences about the astrophysical system are then made. The latter portion of the data analysis procedure, generically referred to as

“parameter estimation” (PE), is the focus of the work presented here.

Parameter estimation procedures, like the search, amount to comparing predictions for the GW data with those that were actually collected by the interferometers. Unlike the search phase, which uses a grid of waveforms to make a point estimate of the GW signal’s parameters, PE is performed over the continuous parameter space, and the goal of the sampler is to find all of the locations in the prior volume that sufficiently minimize the  $\chi^2$  statistic or maximize the likelihood

$$p(d|\boldsymbol{\theta}) \propto e^{-\chi^2/2}, \quad (3)$$

which, when weighted by our prior expectations for the model parameters  $p(\boldsymbol{\theta})$ , yields the posterior distribution function

$$p(\boldsymbol{\theta}|d, M) \equiv \frac{p(d|\boldsymbol{\theta}, M)p(\boldsymbol{\theta}|M)}{p(d|M)}, \quad (4)$$

where  $p(d|M)$  is the marginalized likelihood, or evidence, computed formally by integrating the numerator of Eq. (4) over the full parameter space. The posterior distribution function encodes our knowledge about the astrophysical system, and it is from this distribution that all inferences are ultimately made. Here and throughout this paper,  $p(a|b)$  denotes the conditional probability density of  $a$  given  $b$  so, for example, the likelihood  $p(d|\boldsymbol{\theta}, M)$  is the probability (density) that we would measure data realization  $d$  given model  $M$  and parameters  $\boldsymbol{\theta}$ . Implicit in this strategy are all of our assumptions about the data, represented by  $M$ , regarding both the signal (the gravitational waves themselves) and the instrument noise. Improving any aspect of  $M$  effectively improves the sensitivity of the detectors, and software is a very cost-effective upgrade. The majority of development work in this regard has focused on the numerator of Eq. (2) through the study of gravitational wave source modeling (e.g., Refs. [15–18]), in which the goal has been to develop approximate waveforms that are both accurate, precise, and computationally efficient. The importance of waveform development cannot be overstated, with key challenges remaining as we approach the advanced detector era.

Having an accurate noise model is of similar value to waveform simulations, but has historically received less attention, with the instrumental background assumed to be stable and easily characterized. The standard picture of gravitational wave detector data assumes that, in the Fourier domain representation, the noise correlation matrix is diagonal (i.e., the noise distribution is stationary [19]) and that the noise samples are Gaussian distributed in each bin. The GW detector noise is characterized by the one-sided power spectral density (PSD)  $S_n(f) \equiv \frac{2}{T} \langle \tilde{n}^2(f) \rangle$ . A typical additional assumption on the data model is that the PSD is both constant over the observation time and precisely known. Figure 1 shows the average PSD for

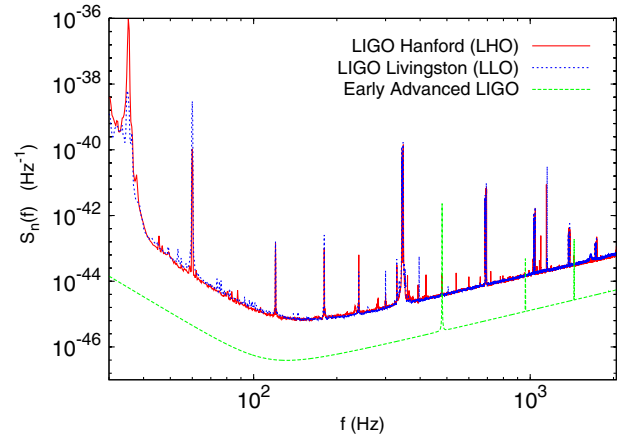


FIG. 1 (color online). Average LIGO PSDs during the S5 run [38], with a prediction for the early Advanced LIGO PSD for comparison [40].

LIGO from the initial design sensitivity configuration and the expected performance of Advanced LIGO when the facilities reach full sensitivity.

The three main assumptions about the noise model—stationarity, Gaussian distributed, and known PSD—are overly restrictive. The instrumental background is at best slowly varying and at worst subject to short-duration, high-amplitude excursions of excess power, or “glitches.” Even in the absence of the (very problematic) glitches, failure to account for long-term changes in the noise PSD can introduce systematic errors in the signal characterization, which may be comparable to the fundamental statistical measurement uncertainties. Examples in which PSD fluctuations have significantly biased the parameter estimation were encountered during the parameter estimation analysis of simulated signals added to data from the sixth LIGO and second Virgo science runs (S6 and VSR2, respectively) [20].

Parameter estimation is only half of the story, as the follow-up analyses can be employed to quantifiably assess the probability that a candidate detection is a GW event by comparing the evidence for competing models of the data. In this context, the alternative model to a GW detection is that the data contain only instrument noise. Thus, having an adequate noise model is a prerequisite for performing meaningful evidence comparisons.

Previous considerations for extending GW data analyses beyond the strict assumptions about the noise came with Allen *et al.* in Refs. [6,21], to which most subsequent studies owe their heritage. The papers by Allen *et al.*, while considering generic GW data analysis, are ultimately focused on measuring a stochastic GW background, which was more recently addressed in a space-based detector context [22] in which the noise must be modeled as the usual instrumental background plus a *foreground* of blended gravitational-wave signals from the Galaxy. There are several studies promoting the use of a student’s  $t$  distribution in lieu of Gaussian noise [19,23] as well as

model selection recipes that employ additional coherence tests between detectors in the global interferometer network [24]. Finally, several proof-of-principle studies using simulated data and parametrized noise models [22,25–27] appear in the literature, building from techniques developed in response to the Mock LISA Data Challenges (e.g., Refs. [28,29]).

In this work, we will employ a simple modification to the noise model by upholding our classic assumptions about the statistical properties of the noise but adding parameters to our model of the data that fit for the PSD in each detector. The additional degrees of freedom that we will introduce to the model are not strongly correlated with astrophysical parameters, thus avoiding marginalization penalties for the physical quantities of interest (e.g., mass, spin, location, etc.). The parametrized noise model demonstrated here still operates completely in the Fourier domain without modifying the underlying functional form of the noise distribution and therefore cannot account for short-duration transient and non-Gaussian noise events. However, we find significant improvement in the consistency of the parameter estimation and model selection results when using real LIGO data. Mitigating the effects of detector glitches and therefore relaxing the assumptions about stationary and Gaussian noise will be left to future work in which we intend to build off of the theoretical progress and simulations made in Refs. [19,21,27].

## II. LIGO POWER SPECTRAL DENSITY

GW observatories are (nearly) continuously collecting data as all-sky monitors awaiting transient events. Without the ability to make accurate, real-time estimates of the instrument background, previous incarnations of the parameter estimation analysis have leveraged the scarcity of GW events by using “off-source” segments of data that are temporally “nearby” a significant trigger produced by the matched-filtering search. The off-source data are assumed to contain only instrument noise and are divided into smaller segments used to estimate the PSD via Welch filtering. The terms *off source* and *nearby* are not well defined or motivated and often come down to an *ad hoc* procedure of looking at the data around a trigger and choosing times that do not have any obvious pathologies (glitches, highly irregular PSDs, etc.). To date, the PE procedures have had to make trade offs between long-duration averages of the PSD for precision and shorter contemporaneous estimates of the noise level to combat the slow drift of the PSD. Short-duration data segments are not ideal for PSD estimation, as the Welch averaging is noisy in its own right. Consequently, the prescription used in past analysis can be limited by the availability of data for PSD estimation, perhaps due to gaps in the data collection or obvious glitches. The challenges of adequately estimating the PSD from off-source data will be exacerbated in the Advanced LIGO/Virgo data because, due to the improved

sensitivity at low frequencies, the amount of time during which GW inspiral signals are in the sensitive band of the detector will be increased substantially over previous incarnations of the detectors. For low-mass sources, such as binary neutron star mergers, several hours worth of data will be necessary for PSD estimation. Figure 2 shows how the goals of PSD estimation are in tension with the protocol for characterizing the instrument background. The grey trace is the superposition of four different PSD estimates made from successive 1024 s segments of data. The solid (red) line is the noise spectrum of the same data, for which the entire 4096 s were used for its determination. Both example PSDs are compared to the average noise level for the entire S5 run, shown in the dashed (blue) line. Intuitively, it is clear that the random fluctuations in the noise estimate—which effectively add noise to the measurement—are suppressed by using more data, but at the expense of accuracy as the PSD slowly drifts.

During times in which the GW detector’s noise characteristics satisfy the assumptions of Gaussianity and long-term stationarity, the outlined procedure for noise estimation is adequate for detection, characterization, and, to some extent, model selection. Nonetheless, data analysis procedures must be prepared for the inevitability of discrepancies between the estimated instrument background and the actual noise in the data being analyzed. Figure 3 shows an example of such variations occurring over several days using data collected when LIGO was operating at its initial design sensitivity. The figure focuses on the lower frequencies of the sensitivity band because, at




FIG. 2 (color online). Comparison of PSD estimates over a small fraction of the full LIGO bandwidth made with the prescription applied in LIGO/Virgo parameter estimation software and using different durations of data to characterize the instrument background. The gray line is a superposition of four consecutive 1024 s blocks of data, divided into 16 s segments for which the PSD is computed and then averaged across the block. 1024 s is typical for the PE follow-up of triggers from the search pipeline. The red (solid) line shows the average PSD if the full 4096 s of data are averaged. The blue dotted line is the average PSD for the entire science run.




FIG. 3 (color online). Example PSDs averaged over 4096 s of initial LIGO science data from the Livingston Observatory. Secular drift in the noise is subtle but does significantly impact the characterization of gravitational wave signals. Dates are in dd/mm/yy format, and times are in UTC, with the GPS time for the start of the data segment in parenthesis.

high frequency, the noise behavior is much more stable. Dates are in the dd/mm/yy format, and times are in UTC, with the GPS time for the beginning of the data segment in parenthesis. The data used to estimate the PSDs were collected at the LIGO Livingston Observatory (local time is UTC-6 hours). By allowing the follow-up analysis to fit for the PSD level while simultaneously characterizing the GW signal, the restrictions that require sufficiently quiescent data near a candidate event can be relaxed, instead allowing for long-term estimates of the PSD to be used. In this paper, we will demonstrate such an adaptation to the model on real detector data—a first for this type of noise modeling, which has, until now, been restricted to proof-of-principal studies on simulated data.

Before delving into the supporting examples, we must first understand how the noise enters into the data analysis formalism by considering discretely sampled noise-only data from a single interferometer. The probability of collecting  $N$  noise samples  $\tilde{\mathbf{n}}$  when noise has zero bin-to-bin correlations is

$$p(d) \equiv \prod_i^N \frac{1}{\sqrt{\pi}\sigma_i} e^{-\frac{|\tilde{n}_i|^2}{\sigma_i^2}}, \quad (5)$$

where  $\sigma_i^2 = \frac{1}{2}S_{n,i}$ , the PSD is  $S_{n,i}$ , and the summation is performed over positive frequencies only. From this fundamental assumption about the character of the instrument noise using Eq. (1) and the discretized version of Eq. (2) it is clear how we arrive at the likelihood function (3) ubiquitously applied in gravitational wave data analysis.

To relax the assumption that we precisely know the PSD, we introduce scale parameters  $\boldsymbol{\eta}$ , which act as piecewise multipliers to the PSD, modifying the likelihood as follows [26]:

$$S_{n,i} \rightarrow \eta_j S_{n,i}, \quad i_j < i \leq i_{j+1}$$

$$\ln p(d|\boldsymbol{\theta}, \mathcal{N}_{\mathcal{P}}) = -\frac{1}{2} \left( \chi^2 + \sum_j N_j \ln \eta_j \right) + \text{const.} \quad (6)$$

Each  $\eta_j$  spans  $N_j$  Fourier bins from  $i_j$  to  $i_{j+1}$ . The logarithmic term in Eq. (6) comes from the normalization in Eq. (5), which is normally left out of any calculations because it does not depend on waveform parameters and because *relative*, rather than absolute, likelihoods are employed in the PE procedures. There is an independent set of  $\boldsymbol{\eta}$  for each detector in the network, and an additional summation over interferometers must be included in Eq. (6).

### III. METHODOLOGY

To test the merits of the PSD-fitting noise model, we use data from the end of LIGO’s fifth science run (S5). S5 began on November 5, 2005, and was completed on September 30, 2007, yielding a series of novel, though null, studies pertaining to limits placed on gravitational wave emission by nearby pulsars and event rates for compact binary mergers and cosmological backgrounds [7,30–33]. Since then, the LIGO and Virgo facilities have undergone a series of upgrades, and additional science runs (S6 and VSR2), to substantially improve the prospects of directly detecting gravitational radiation and realizing the promises of this long-anticipated new field in astronomy.

Virgo data from VSR1 were also available during part of S5, but we have restricted this study to LIGO-only data for the following reasons. Having only two detectors in the network puts stress on the parameter estimation methods due to strong degeneracies between extrinsic parameters (spatial location and orientation of the source), and an auxiliary goal of this work was to improve efficiency in such a network configuration. Furthermore, it is possible that early detections will come from a two-detector network as, during any given science run, the individual detectors have duty cycles below 100%. For example,  $\sim 75\%$  of the “coincident time,” defined as having more than one interferometer producing good science data, during the S6/VSR2 run had only two detectors operating [8]. Additionally, the LIGO and Virgo facilities are operating on different schedules to complete the upgrades and commissioning of the advanced detectors.

We use the Markov chain Monte Carlo software found in the LIGO Algorithm Library (LAL) package LALINFERENCE, which was one of the samplers used for the parameter estimation studies following significant events from the S6 science run [20], in which we have included optional noise-model parameters. Complete details of the sampler’s specific implementation in LALINFERENCE will be available in a forthcoming publication. The MCMC algorithm implemented in LAL leans heavily on parallel tempering [34] to efficiently sample

the complicated posterior distribution function we have come to expect from gravitational wave detections. In parallel tempering, several Markov chains are run simultaneously, each sampling the target distribution tempered by a parameter  $T$  such that  $p(d|\boldsymbol{\theta}, M, T) \rightarrow p(d|\boldsymbol{\theta}, M)^{1/T}$ . Chains can exchange parameters while preserving detailed balance, thereby improving the efficiency with which the  $T = 1$  chain samples the target posterior distribution function. We further use the high-temperature chains to compute model evidences using thermodynamic integration [27,35] via

$$p(d|M) = \int_0^1 \langle p(d|\boldsymbol{\theta}, M) \rangle_{\beta} d\beta, \quad (7)$$

where  $\beta = 1/T$  is the inverse temperature of the chain. For this application, we use 24 parallel chains with temperatures  $T$  spaced geometrically, i.e.,  $T_i = T_{i-1}^{\delta T}$  with a step size of  $\delta T = 1.3$  for chains 1–12 ( $T = 1$ –25) and  $\delta T = 1.55$  for chains 13–24 ( $T = 40$ –5000). We followed the prescription for diagnosing parallel tempering and thermodynamic integration laid out in Appendix A of Ref. [27] to ensure an adequate number of chains and temperatures.

The exact implementation of the PSD scale parameters can be adjusted; however, we did not find a strong influence on the results for the GW parameters based on different choices for the noise parametrization. The most obvious adjustable features for the PSD parameters are the number of scale factors per interferometer and the priors used for these same parameters. We use eight scale parameters per interferometer, spaced logarithmically in frequency giving the noise fitting more fidelity at lower frequencies, at which the PSD fluctuations are most prevalent. Our choice of prior in this work is the same as originally described in Refs. [25,26], in which we assume that the priors are Gaussian with widths determined by the number of Fourier bins over which each scale factor is applied. The scheme for proposing a new scale parameter  $\eta_{j,y}$  from the current value  $\eta_{j,x}$  employs a draw from a normal distribution with variance proportional to the width of the prior. To wit,

$$p(\eta_j|M) \propto e^{-\frac{(1-\eta_j)^2}{2\sigma_j^2}} \quad \eta_{j,y} = \eta_{j,x} + N(0, \sigma_j^2/10) \quad (8)$$

$$\sigma_j = \alpha \times 1/\sqrt{N_j},$$

where  $N_j$  is the number of Fourier bins over which  $\eta_j$  is applied to  $S_n(f)$  and  $\alpha$  is used to broaden the prior distribution, as  $\alpha = 1$  would be used for the case in which we expect no long-term variation in the noise level. For the results reported here, we used  $\alpha = 10$  but found that the GW parameter estimation was not strongly affected by this choice for reasonable values of  $\alpha$ . The scale parameters are independent across interferometers as well as frequency blocks, so the joint prior is simply the product of  $p(\eta_j|M)$ .

## IV. RESULTS

To test the model's performance on real LIGO instrument noise, we selected from the S5 data 300 segments with a duration of 4096 s. As the LIGO interferometers are susceptible to noise artifacts (i.e., glitches), data quality flags identifying times of poor detector behavior were developed and used in GW searches during S5 [36]. We have chosen our segments to fall within times that pass all of the available data quality cuts. For each segment of data, we estimated the PSD using the Welch filtering as implemented in LALINFERENCE. We then selected as the data to be analyzed a 16 s segment from September 13, 2007, (GPS time 873750000) and performed various tests on that data using the PSDs from throughout S5. We will refer to the different models being tested using the following notation:

- (i) Model  $\mathcal{N}_{\mathcal{D}}$ : Gaussian noise with fixed PSD.
- (ii) Model  $\mathcal{N}_{\mathcal{P}}$ : Gaussian noise with variable PSD.
- (iii) Model  $\mathcal{N}_{\mathcal{T}}$ : Student's t-distributed noise with fixed PSD and degrees-of-freedom parameter  $\nu$ .

Each run of the MCMC could also be performed with (model  $\mathcal{S}_1$ ) or without (model  $\mathcal{S}_0$ ) including a gravitational wave template in the model.

### A. Noise-only model comparison

Our first study compared the performance of the model on noise-only data, focusing on four PSDs estimated at times relatively nearby the data being analyzed (August 13, two segments from September 12, and one from September 13, all in 2007).

The top panel of Fig. 4 shows a representative posterior for the whitened noise distribution  $p(n/\sqrt{S_n(f)})$  over a frequency block with a common value for  $\eta$ , spanning  $\sim 15$  Hz and encompassing 240 data samples. The grey shaded region spans the 90% credible interval for the noise model, and the red (solid) line shows the median noise distribution. Effectively, what we show here is the superposition of the different noise models used by the Markov chain over a 15 Hz band in the data. By comparison, the noise distribution using the true PSD would be a unit variance, zero mean Gaussian. Using the median distribution (red, solid curve), we compute the relative error  $(p_{\text{med}}(n) - p_{\text{best}}(n))/p_{\text{best}}(n)$ , shown in the lower panel of Fig. 4. From this figure, we can see how marginalizing over the PSD effectively introduces broader tails to the theoretical noise distribution, despite still using a Gaussian distribution for the noise at any point in the Markov chain.

We repeat the analysis using four different initial PSD estimates and compute the Bayes factor between  $\mathcal{N}_{\mathcal{P}}$  and  $\mathcal{N}_{\mathcal{D}}$  to discern whether or not adding noise parameters sufficiently improves our fit to the data to overcome the ‘‘Occam factor’’ incurred by enlarging the prior volume of the model by introducing new degrees of freedom. Our results are to be found in Fig. 5, where the marginalized

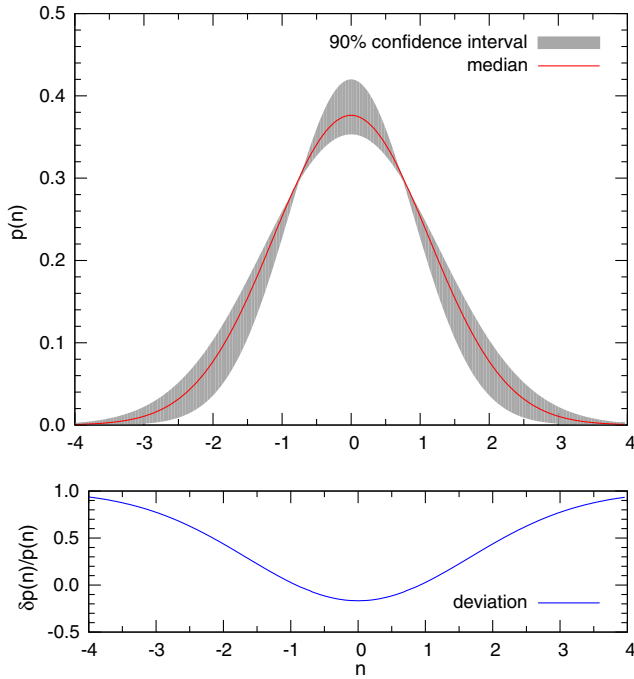


FIG. 4 (color online). Posterior distribution for the whitened noise  $n \rightarrow n/\sqrt{S_n(f)}$  from 40–55 Hz. In the top panel, the line represents the median noise probability density function from an MCMC analysis, while the grey shaded region encompasses the 90% confidence region. The bottom panel shows the residual when the best-fit Gaussian is subtracted from the median. We find that including PSD-fitting effectively broadens the tails of the underlying expected noise distribution.

likelihood ratio for model  $\mathcal{B}_{\mathcal{N}_p, \mathcal{N}_D}$  is shown as a function of the GPS time for the PSD estimation. We generally find very strong support for the PSD-fitting model ( $\mathcal{N}_p$ ). The only example in which the fixed-PSD model was favored, 873750000, used the data being analyzed when making the initial noise estimation. The Bayes factor comparing  $\mathcal{N}_p$  and  $\mathcal{N}_D$  in this example resulted in 5:1 odds in favor of the fixed-PSD model since the PSD estimation in this case is accurate and does not require the additional complexity in the model. However, from a model selection standpoint, the two models are similarly supported by the data—a Bayes factor of 5:1 is not typically considered a strong discriminator between models. Thus, the noise-fitting model is roughly on par with the standard procedure when the initial PSD estimate is accurate. On the other hand, if we move off source to estimate the PSD, as is done in the real analysis, we see *overwhelming* support for the PSD-fitting model (notice the y axis in Fig. 5 is a logarithmic Bayes factor, so even though the support for  $\mathcal{N}_p$  in example 873739840 seems small, the odds ratio is larger than  $2 \times 10^7:1$ ). To be fair, PSD estimation is usually performed at times much closer to the candidate event, but that requirement is born out of the knowledge that there is secular evolution of the PSD and that it ends up restricting the amount of data that can be used to characterize the

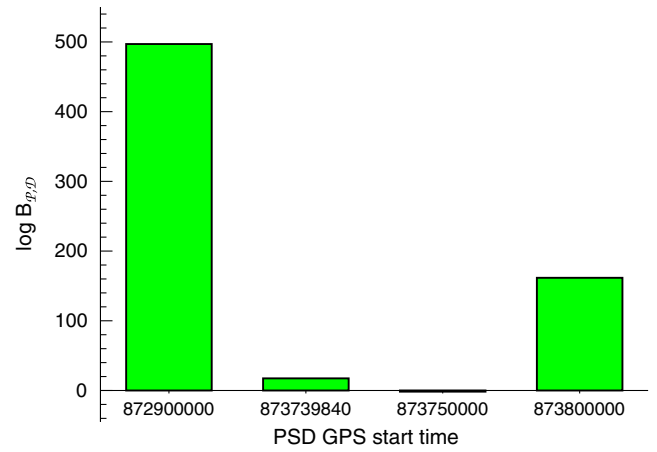


FIG. 5 (color online). Bayes factors computed for noise-only data comparing  $\mathcal{N}_p$  and  $\mathcal{N}_D$  using different PSDs marked by the GPS time of the noise estimate. Models that include PSD parameters are heavily favored except for 873750000, which included the data being analyzed in the noise estimation. In this example,  $\mathcal{N}_D$  is favored with 5:1 odds—not sufficient to discriminate between models. This implies that the PSD fitting does no harm to the analysis when it is not needed and makes a substantial improvement otherwise.

instrument noise, thereby increasing stochastic fluctuations in the PSD estimation itself. The capability that we are demonstrating here frees the analysis of that constraint and opens the possibility of using long-term, smooth estimates for the noise.

## B. Statistical improvement in parameter estimation

The true motivation of this study is not to compare noise-only models but to see how the PSD fitting can impact the characterization of GW detections. To that end, we now simulate a gravitational wave detection by adding a binary inspiral waveform to the data collected at GPS time 873750000 and repeat the analysis including the signal model  $S_1$ .

We selected as our signal model quasicircular, nonspinning, stationary phase waveforms computed out to the 3.5th post-Newtonian order in the phase, taking only the leading-order amplitude—the so-called TaylorF2 waveforms from Ref. [37]. Because this study is focused on the instrument noise, and is not a waveform study, the TaylorF2 templates were chosen because of their computational efficiency. We did not anticipate significantly different results for more physically detailed waveforms. We then compared posterior distribution functions from the various data models as well as the evidence for detection in the form of the Bayes factor between models with and without noise parameters. All injections were set to have a signal-to-noise ratio between 1 and 20, depending on the particular example.

The LALINFERENCE MCMC sampler was run on the same data using the PSDs estimated from the other 300

S5 data segments, with and without noise fitting. Our figure of merit for this study is the distribution of the medians, and the 90% credible intervals, from each run. We focus on the chirp mass  $\mathcal{M} \equiv (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$  because it shows the most sensitivity. In an ideal setting, the PSD would be known *a priori*, and both the median and credible interval would be identical from run to run—the data are not changing from one example to the next—but long-term variations in the PSD, nonstationary and/or non-Gaussian features in the noise used to determine the PSD, and the finite amount of data available for noise estimation all conspire to add systematic errors to our GW measurement and thus modify the shape of the posterior distribution from one example to the next.

Encouragingly, we find significant improvement in the consistency of the parameter estimation when we use the  $\mathcal{N}_{\mathcal{P}}$  model. Figure 6 shows the distribution for the median (left panel) and the width of the 90% credible intervals (right panel) over all 300 runs using different times to estimate the PSD. If our PSD model were perfect, there would be no difference from run to run. In both figures, the red (solid) distribution is for the fixed-PSD model ( $\mathcal{N}_{\mathcal{D}}$ ), and the blue (dotted) histogram comes from the results of the variable-PSD model ( $\mathcal{N}_{\mathcal{P}}$ ). The vertical dashed grey line in each example is the value obtained when using the PSD, which was estimated from the data being analyzed (prior to us adding the gravitational wave signal), and is therefore assumed to be the “true” value for the median and the size of the statistical error.

The fixed-PSD (red) distribution shows two distinct modes in the 90% credible intervals. The LIGO detectors improved throughout S5, with an increase in sensitivity at low frequencies roughly between April 2006 and April 2007 [38]. The two modes seen in Fig. 6 correspond to GPS times before and after the improvements in the LIGO

detectors. While it would never be under consideration to use months-old data to estimate the PSD for a candidate detection, we find this an encouraging example of how including PSD scale parameters improves the consistency from run to run for the location of the posterior distribution function (the median) and its width (the 90% credible interval) are more narrowly peaked around the true value.

### C. Comparison with Student’s t likelihood

We will now take a slight detour in our study to address an alternative method for mitigating the effects of real detector noise—using a student’s t distribution as the underlying noise model. This was originally advocated in a gravitational wave context by Röver *et al.* [23] and has shown promise in studies in which it has been applied to both simulated and real LIGO data [19]. In this paradigm, the likelihood becomes

$$\ln p(d|\theta, \mathcal{N}_{\mathcal{T}}) = -\sum_j \frac{\nu_j + 1}{2} \ln \left( 1 + \frac{\chi_j^2}{\nu_j} \right) + C(\nu_j), \quad (9)$$

where the sum on  $j$  is performed over all of the data,  $\chi_j^2$  is the whitened residual power in a single bin of data, and  $C(\nu_j)$  is the normalization that is computed using Gamma functions in  $\nu$ . In the limit of  $\nu_j \rightarrow \infty$ , the student’s t likelihood recovers the standard Gaussian likelihood in Eq. (3). The student’s t distribution is the ideal theoretical choice for the likelihood when the noise is stationary and the PSD is unknown and must be estimated by a finite amount of data. Under such circumstances, the degrees-of-freedom parameter  $\nu_j$  is, in principal, a constant over  $j$  determined by the number of segments of data used to make the initial estimate of the PSD. However, small values of  $\nu_j$  substantially broaden the tails of the

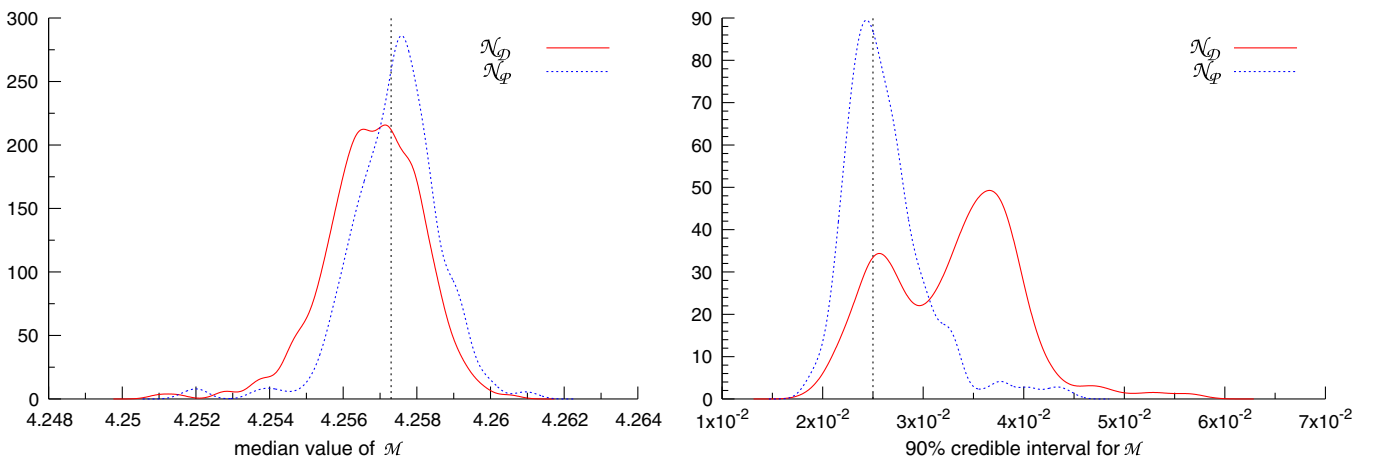


FIG. 6 (color online). Distributions of median (left) and width of 90% credible intervals (right) obtained from the MCMC samples after analyzing the same data using 300 different initial PSD estimates. The red (solid) histograms come from the model that assumes a perfectly known PSD, while the blue (dotted) histogram is for a model with PSD-scale parameters. The vertical dashed line denotes the value when the PSD estimation came from the data being analyzed (prior to the GW signal being injected). Using the PSD scale parameters, we find more consistency in the posterior distribution function across different times used to estimate the PSD.

theoretical noise distribution, making the noise model robust against non-Gaussianity in the data. Röver found  $\nu \sim 10$  to work well for real detector data despite using  $\sim 30$  data segments to estimate the PSD.

We ran a few of our cases using the student's t distribution  $\mathcal{N}_{\mathcal{T}}$  to compare against models  $\mathcal{N}_{\mathcal{D}}$  and  $\mathcal{N}_{\mathcal{P}}$ . Results are shown in Fig. 7, where we have again chosen to focus on the marginalized posterior distribution function for the chirp mass  $\mathcal{M}$ , as it shows the largest dependence to the PSD time for this example. Notice  $\mathcal{M}$  in this example is larger than before—we used a higher-mass injection to put the “loudest” part of the waveform in the data where PSD fluctuations were most noticeable. The top left panel shows the results using the usual Gaussian likelihood with no PSD-scale parameters  $\mathcal{N}_{\mathcal{D}}$ . The independent variable is the displacement away from the injected value. Different histograms correspond to the GPS time when the PSD was estimated. Time 873750000 (blue, dashed) contains the data segment being analyzed. The differences between chirp-mass

distributions may not seem dramatic, but given that the data are the same for each run, we would like to mitigate any differences due to the PSD estimation as much as possible. If we redo the analysis but now include the PSD-fitting parameters  $\mathcal{N}_{\mathcal{P}}$  (bottom left panel), we find improved consistency. The top right panel uses no noise parameters but adopts a student's t distribution with constant  $\nu_j$  determined by the number of data segments used to estimate the PSD. In this example, we find that the student's t likelihood does improve the convergence for one example (873739840, green dashed line) but could not similarly improve results using the PSD from GPS time 871000000 (red, solid line). This does not imply that an optimized application of the student's t-distribution likelihood, by choosing fewer degrees of freedom  $\nu$ , is not capable of improving this example. The point of this demonstration is to show that differences in PE results are not simply attributed to sampling errors in the PSD estimation (which the student's t distribution should perfectly accommodate) and the PSD

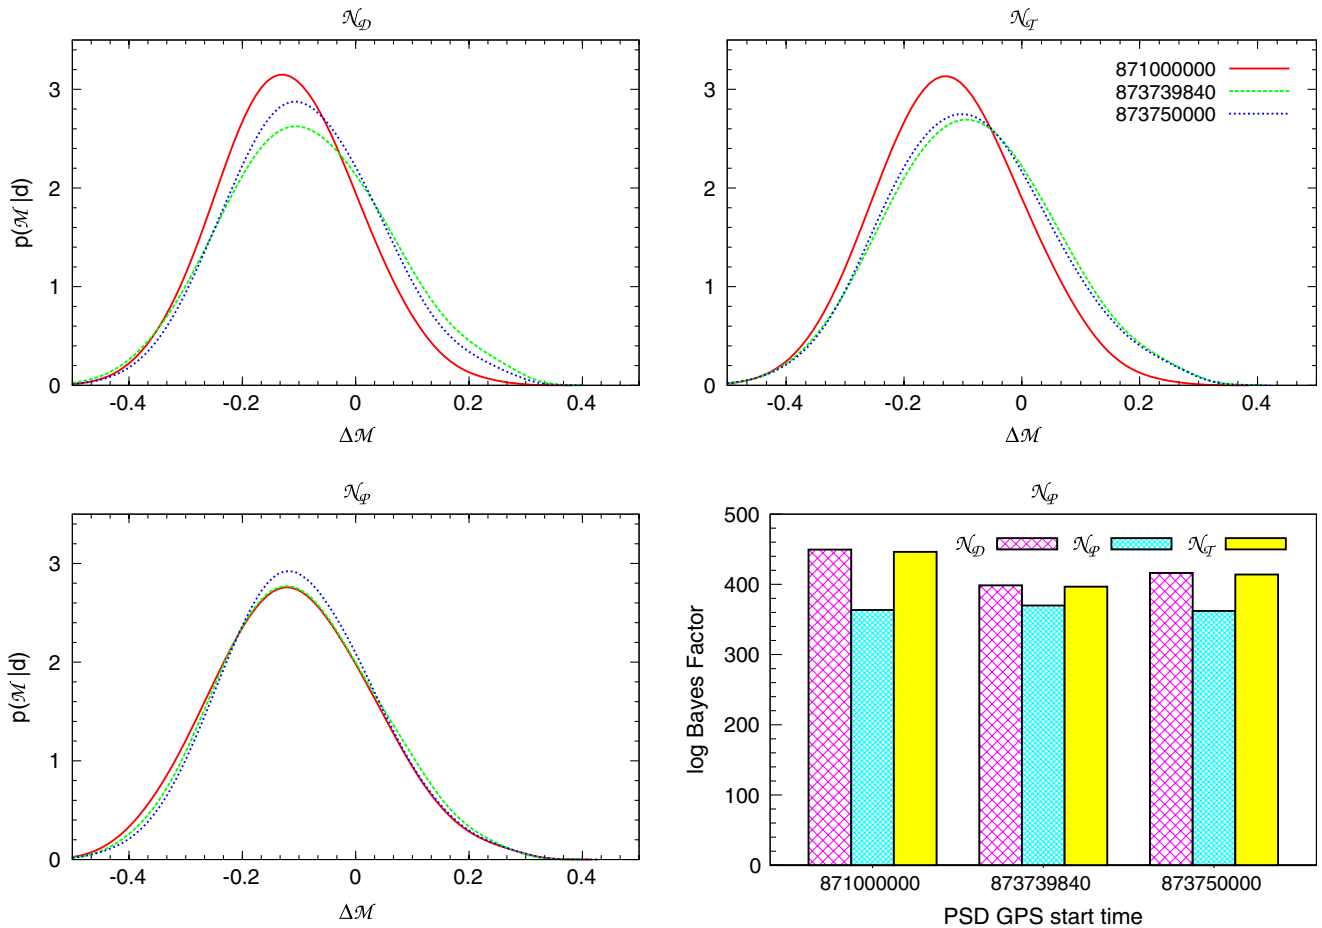


FIG. 7 (color online). Posteriors for  $\mathcal{M}$  for different PSDs and models. The top left panel shows the results of the standard version of the MCMC sampler. The top right panel uses one implementation of the student's t likelihood. The bottom left panel uses the PSD-fitting model. In these examples, the PSD-fitting model is the least sensitive to different PSD estimates. The bottom right panel shows the Bayes factor for detection as a function of PSD estimation. In these examples, the PSD-fitting model  $\mathcal{N}_{\mathcal{P}}$  yields the most stable evidence ratios.



scale parameters can readily ameliorate differences between the data being analyzed and the estimated noise level.

The bottom right panel of Fig. 7 shows an arguably more significant divergence between the standard likelihood, the student’s t, and the PSD scale parameters. Here, we show how the Bayes factor we would compute to determine the odds that a GW has been detected depends strongly on the PSD estimation. In this example, the PSD-fitting model  $\mathcal{N}_P$  (cyan, dotted) shows the least variability in the detection confidence based on the PSD estimation, while both the fixed-PSD model  $\mathcal{N}_D$  (magenta, cross-hatched) and student’s t likelihood  $\mathcal{N}_T$  (yellow, solid) exhibit similar sensitivity to the noise estimation, despite  $\mathcal{N}_T$  being able to “correct” the posterior for GPS time 873739840. The Bayes factors in this example are sufficiently large ( $\sim e^{400}:1$ ) that these fluctuations would not impact any conclusions drawn from the data. However, we will now investigate the role  $\eta$  can play in model selection for marginal detections, when precision and accuracy in the Bayes factor computation will be paramount in evaluating candidate GW events.

We would be remiss if we failed to reemphasize that the student’s t distribution in these examples uses the strict interpretation of the degrees of freedom  $\nu$  and that by artificially setting that to a lower value, or marginalizing over it in the way we treat the PSD scale parameters, we may temper the differences between the models. The PSD-fitting method advocated in this work is in fact closely related to the student’s t distribution likelihood in Refs. [19,23]. By replacing our Gaussian prior on  $\eta$  with an inverse-chi-squared distribution and allowing each Fourier bin an independent  $\eta$ , we can analytically marginalize over the noise parameters and arrive at the student’s t distribution for the likelihood. If we were to adopt the inverse-chi-squared prior for the piecewise PSD scale parameters and integrate the posterior over  $\eta$ , we arrive at a likelihood that is qualitatively similar to the student’s

t distribution but maintains the broad flexibility needed to handle large-scale drift in the LIGO/Virgo PSD. The derivation and discussion of the analytic marginalization for our PSD-fitting model can be found in the Appendix.

**D. Discerning between signal and noise models**

For our last study, we will use the same GW system and S5 data as before but now change the distance to the source to get a desired SNR. In doing so, we can monitor how the Bayes factor changes as we increase the signal strength and compare our two Gaussian-noise models  $\mathcal{N}_D$  and  $\mathcal{N}_P$  with regards to their model selection capabilities. The Bayes factor vs SNR plots can be found in Fig. 8. Here, we see substantial variation in the Bayes factors when using noise model  $\mathcal{N}_D$  (left-hand panel), to the point where different conclusions about low SNR gravitational wave detections (see inset of left-hand panel) could potentially be made. For example, the SNR = 6 injection produced Bayes factors between  $\sim 50:1$  and  $\sim 2 \times 10^5:1$  under model  $\mathcal{N}_D$ . Bayesian model selection will not be a useful tool for discriminating between GW detection and noise if it is subject to such variability. On the other hand, our model selection calculation is substantially more stable when we include the PSD scale parameters (right-hand panel).

**V. SUMMARY AND FUTURE WORK**

We have demonstrated, using LIGO data from the fifth science run, that adding scale parameters that modify the PSD makes inferences drawn from the data more robust against slowly varying instrument noise. The added parameters effectively allow us to marginalize our results over uncertainty in the PSD. This work comprehensively demonstrates the importance, and the benefits, of modeling detector noise in the context of parameter estimation and model selection.

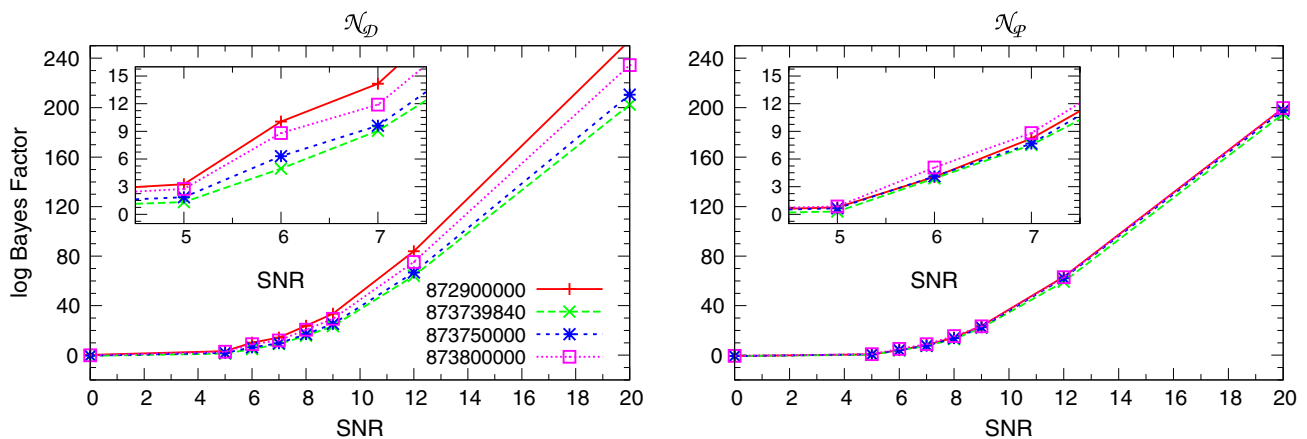


FIG. 8 (color online). Bayes factors as a function of SNR for different noise models and PSDs. The left-hand panel shows substantial variation depending on PSD estimation. Bayes factors that show this much variability are of limited use. The right-hand panel shows significant improvement gained by including the PSD-fitting parameters.

We show in Fig. 5 that the noise model including scale parameters  $\eta$  receives overwhelming support when the PSD estimation comes from times separate from the data being analyzed but garners comparable support to the fixed PSD noise model when the instrument background is estimated from the data itself (which is, of course, not possible when a potential GW event is contained in the data).

We show in Fig. 6 that parameter estimation—in the way of a point estimate (the median) and the width of the posterior distribution—is sensitive to the PSD estimation and that sensitivity can be significantly reduced by marginalizing over the noise level as prescribed here.

We performed a small comparison between the PSD-fitting model and the student’s  $t$  likelihood in Fig. 7 and find the student’s  $t$  distribution, as we employed it here, was not as effective as our parametrized PSD. That being said, we see potential in combining the assets of these different approaches to solving the same problem and anticipate a hybrid marginalized-PSD/student’s  $t$  likelihood may prove optimal, and we will pursue this and other strategies for noise modeling as we continue to improve the LIGO/Virgo noise model, such as using the analytically marginalized likelihood in the Appendix and marginalizing over hyperparameters characterizing the prior.

Arguably the most important impact of the noise estimation is in computing the Bayes factor to select between models suggesting a GW detection and those favoring instrument noise. Figure 8 shows strong dependence on the PSD estimate upward of 20% to 50% variation in  $\log \mathcal{B}_{S_1, S_0}$ , which is reduced to  $\lesssim 1\%$  to 20% when the PSD-fitting parameters are included. This improvement does not come at the expense of diminished sensitivity to the GW signals (which is always a risk when the dimension of a model is increased). Specifically, if we look at the  $\text{SNR} = 7$  injection in Figure 8, without the PSD-fitting parameters, we could conclude that the Bayes Factor for the signal model was anywhere between  $8 \times 10^3:1$  and  $1.2 \times 10^6:1$ , which are sufficiently different odds to warrant alternative interpretations of the detection if we were to interpret the odds ratio as a sort of false alarm rate. When the PSD parameters are included, the Bayes factors are found to lie between  $2 \times 10^3:1$  and  $8 \times 10^3:1$ , telling a more consistent story about the significance of this (simulated) candidate.

The noise model promoted here frees the parameter estimation follow-up analysis from its requirement to find data near to a candidate event to estimate the noise level. Such a requirement can impede the PE results because data-quality shortcomings near a GW trigger can restrict the duration of data available for the noise estimation. Finding enough data to estimate the PSD will become increasingly problematic in the advanced detector era as low-mass binary inspiral signals will be  $\sim 1000$  s in

duration, requiring several hours of data for adequate Welch filtering—a *very* long time to hope for stationary noise. The ability to use long-term estimates of the instrument noise relaxes that restriction, mitigates the impact of transient noise events in the PSD estimation, and reduces random fluctuations introduced in the residual by the finite duration of data used to estimate the PSD.

Our demonstrations using real instrument data in this paper show great promise in making parameter estimation results more robust in the presence of real instrument noise, but they are not perfect. Alternative functional forms for the PSD scaling parameters, such as using parametrized power laws [22] or a physical model for the noise power spectral density, may further improve noise modeling while ultimately liberating the analysis from using data outside of that which is being analyzed to provide information about the instrument background.

While mitigating the impact of the initial PSD estimation is certainly an important technology for gravitational wave data analysis, it does not address the more fundamental assumptions about the instrument noise—that the noise is Gaussian distributed and stationary over the duration of the data being analyzed. This study is our first foray into applying more sophisticated noise modeling techniques into real detector data, and we are bolstered by the results here to pursue alternative, “heavy-tailed” distributions for the noise such as the student’s  $t$  [19] or a sum of two Gaussians [27], as well as incorporating the Bayesian model selection on glitches into the LALINFERENCE software as demonstrated by the *Bayes Wave* algorithm first introduced in Ref. [27].

## ACKNOWLEDGMENTS

We would like to thank Nelson Christensen for thorough comments on this manuscript, specifically in regard to our use of the student’s  $t$  distribution, and Neil Cornish for suggesting the inclusion of Fig. 4 as a way to understand the effect of marginalizing over the PSD. We also acknowledge Vicky Kalogera and Ilya Mandel for discussions and suggestions related to this work. T.L. and B.F. are supported by NSF LIGO Grant No. PHY-0969820. B.F. receives additional support from a NSF graduate research fellowship, Award No. DGE-0824162. M.C.’s work is funded by the Winston Churchill Foundation of the United States of America. The MCMC runs were carried out using computing resources from the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation Grant No. OCI-1053575 and the Nemo cluster given by National Science Foundation Awards No. PHY-0923409 and No. PHY-0421416 to UW-Milwaukee. T.L. would like to thank the The Center for Space Plasma and Aeronomic Research (CSPAR) at The University of Alabama in Huntsville for its hospitality.

### APPENDIX: ANALYTIC MARGINALIZATION OVER $\boldsymbol{\eta}$

It is possible to analytically marginalize over the noise parameters that appear in Eq. (6) if a slightly different prior than Eq. (8) is used. Consider the likelihood in Eq. (6),

$$p(d|\boldsymbol{\theta}, \mathcal{N}_{\mathcal{P}}) = \prod_j \prod_{i_j < i \leq i_{j+1}} \frac{1}{\sqrt{2\pi\eta_j S_{n,i}}} \times \exp\left[-\frac{\sum_{i_j < i \leq i_{j+1}} \chi_i^2}{2\eta_j}\right], \quad (\text{A1})$$

where

$$\chi_i^2 \equiv \frac{|\tilde{n}_i|^2}{S_{n,i}}. \quad (\text{A2})$$

If we impose a prior on each  $\eta_j$  of the form [39]

$$p(\eta_j|\kappa_j, \mathcal{N}_{\mathcal{P}}) = \frac{(1 + \kappa_j)^{1+\kappa_j}}{\Gamma(1 + \kappa_j)} \eta_j^{-(\kappa_j+3)} \exp\left[-\frac{1 + \kappa_j}{\eta_j}\right], \quad (\text{A3})$$

where  $\kappa_j > 0$  is a hyperparameter, then we can integrate the dependence on  $\eta_j$  out of the posterior (i.e., marginalize over the  $\eta_j$ ), obtaining

$$\begin{aligned} p(d|\boldsymbol{\theta}', \boldsymbol{\kappa}, \mathcal{N}_{\mathcal{P}}) &\equiv \int d\boldsymbol{\eta} p(d|\boldsymbol{\theta}, \mathcal{N}_{\mathcal{P}}) p(\boldsymbol{\eta}|\boldsymbol{\kappa}, \mathcal{N}_{\mathcal{P}}) \\ &= \prod_j \frac{\Gamma(\frac{N_j}{2} + 2 + \kappa_j)}{\Gamma(1 + \kappa_j)} (1 + \kappa_j)^{-\frac{N_j}{2}+1} \\ &\quad \times \left[ \prod_{i_j < i \leq i_{j+1}} \frac{1}{\sqrt{2\pi S_{n,i}}} \right] \\ &\quad \times \left( 1 + \frac{1}{2(1 + \kappa_j)} \sum_{i_j < i \leq i_{j+1}} \chi_i^2 \right)^{-\frac{N_j}{2} + \kappa_j + 2}, \end{aligned} \quad (\text{A4})$$

where  $\boldsymbol{\theta}' = \boldsymbol{\theta} \setminus \boldsymbol{\eta}$  are the non-noise parameters, and recall that  $N_j = i_{j+1} - i_j$  is the number of frequency bins in the  $j$ th group.

The prior in Eq. (A3) is chosen so that the *prior* mean of  $\boldsymbol{\eta}$  is **1**, and the variance is

$$\text{Var}_{\text{prior}}(\boldsymbol{\eta}) = \frac{1}{\boldsymbol{\kappa}}. \quad (\text{A5})$$

The distribution in Eq. (A3) is not a Gaussian, but, provided that  $\kappa_j \ll N_j$ , the prior does not have a strong influence on the posterior. If, for example, we choose  $\kappa_j$  to match the variance in Eq. (8), that is,

$$\kappa_j = \frac{N_j}{\alpha^2}, \quad (\text{A6})$$

with  $\alpha \simeq 10$ , then we do not expect that the difference between the Gaussian prior and the inverse-chi-squared prior will matter in the posterior. Intuitively, as long as the prior allows for relative fluctuations in the noise level that are greater than  $1/\sqrt{N_j}$ , the uncertainty in the posterior from making  $N_j$  independent measurements of the noise level in the frequency bins from  $i_j$  to  $i_{j+1}$  will be much smaller than the uncertainty in the prior, and the precise shape of the prior will not matter.

Note that the likelihood in Eq. (A4) differs in general from the student's t likelihood discussed in Ref. [19] and used elsewhere in this paper. In that work, each frequency bin is allowed to fluctuate independently in noise level, and the prior on the noise is not constrained to have a mean of 1. The independence of the noise bins means that there is, in effect, one  $\eta$  parameter per bin in that work [see Eq. (9)]; because of the large number of noise parameters, measurements of the noise are correspondingly weak, and the marginalized likelihood is much more dependent on choice of prior. For the choice  $\nu_j \simeq 10$ , the prior variance on the noise (which is approximately the posterior variance on the noise, since there is only one bin per noise parameter) is about 10%; in the method discussed in this Appendix and throughout the paper, the posterior variance is controlled mostly by the  $N_j$  independent measurements of the noise level associated with one  $\eta_j$  parameter, not the wide prior. This method can therefore cope better with the relatively large noise variations from time to time (see Fig. 2) without large uncertainty on the noise estimate from the data.

- 
- [1] M. Vallisneri, *Phys. Rev. Lett.* **107**, 191104 (2011).  
 [2] M. van der Sluys, I. Mandel, V. Raymond, V. Kalogera, C. Röver, and N. Christensen, *Classical Quantum Gravity* **26**, 204010 (2009).  
 [3] V. Raymond, M. V. van der Sluys, I. Mandel, V. Kalogera, C. Röver, and N. Christensen, *Classical Quantum Gravity* **27**, 114009 (2010).

- [4] S. Vitale, W.D. Pozzo, T.G.F. Li, C. Van Den Broeck, I. Mandel, B. Aylott, and J. Veitch, *Phys. Rev. D* **85**, 064034 (2012).  
 [5] K.S. Thorne, *Gravitational Radiation in 300 Years of Gravitation* (Cambridge University Press, Cambridge, England, 1987), p. 330.  
 [6] B. Allen, J.D.E. Creighton, E.E. Flanagan, and J.D. Romano, *Phys. Rev. D* **65**, 122002 (2002).

- [7] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **82**, 102001 (2010).
- [8] J. Abadie *et al.* (LIGO Collaboration and Virgo Collaboration), *Phys. Rev. D* **85**, 082002 (2012).
- [9] J. Skilling, in *American Institute of Physics Conference Series*, edited by R. Fischer, R. Preuss, and U.V. Toussaint, American Institute of Physics Conference Series Vol. 735 (American Institute of Physics, New York, 2004), p. 395.
- [10] <http://www.inference.phy.cam.ac.uk/bayesys/>.
- [11] J. Veitch and A. Vecchio, *Phys. Rev. D* **78**, 022001 (2008).
- [12] F. Feroz, M.P. Hobson, and M. Bridges, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
- [13] N. Metropolis, A. Rosenbluth, R. M.N., A. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
- [14] W. Hastings, *Biometrika* **57**, 97 (1970).
- [15] L. Blanchet, *Living Rev. Relativity* **9**, 4 (2006).
- [16] P. Ajith *et al.*, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [17] A. Buonanno, Yi. Pan, J. Baker, J. Centrella, B. Kelly, S. McWilliams, and J. van Meter, *Phys. Rev. D* **76**, 104049 (2007).
- [18] L. Lindblom, B. J. Owen, and D. A. Brown, *Phys. Rev. D* **78**, 124020 (2008).
- [19] C. Rover, *Phys. Rev. D* **84**, 122004 (2011).
- [20] J. Aasi *et al.* (LIGO Collaboration and Virgo Collaboration), *Phys. Rev. D* **88**, 062001 (2013).
- [21] B. Allen, J.D. Creighton, E.E. Flanagan, and J.D. Romano, *Phys. Rev. D* **67**, 122002 (2003).
- [22] M.R. Adams and N.J. Cornish, *Phys. Rev. D* **82**, 022002 (2010).
- [23] C. Rover, R. Meyer, and N. Christensen, *Classical Quantum Gravity* **28**, 015010 (2011).
- [24] J. Veitch and A. Vecchio, *Phys. Rev. D* **81**, 062003 (2010).
- [25] N.J. Cornish and T.B. Littenberg, *Phys. Rev. D* **76**, 083006 (2007).
- [26] T.B. Littenberg and N.J. Cornish, *Phys. Rev. D* **80**, 063007 (2009).
- [27] T.B. Littenberg and N.J. Cornish, *Phys. Rev. D* **82**, 103007 (2010).
- [28] S. Babak *et al.* (Mock LISA Data Challenge Task Force), *Classical Quantum Gravity* **25**, 114037 (2008).
- [29] S. Babak *et al.* (Mock LISA Data Challenge Task Force), *Classical Quantum Gravity* **27**, 084009 (2010).
- [30] B. Abbott *et al.* (LIGO Scientific Collaboration), *Astrophys. J.* **681**, 1419 (2008).
- [31] B. Abbott *et al.* (LIGO Scientific Collaboration), *Astrophys. J. Lett.* **683**, L45 (2008).
- [32] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **80**, 102002 (2009).
- [33] B. Abbott *et al.* (LIGO Scientific Collaboration and The Virgo Collaboration), *Nature (London)* **460**, 990 (2009).
- [34] R. Swendsen and J. Wang, *Phys. Rev. Lett.* **57**, 2607 (1986).
- [35] P. Goggans and Y. Chi, *Bayesian Inference and Methods in Science and Engineering* (American Institute of Physics, New York, 2004).
- [36] LIGO Scientific Collaboration, Data Quality and Veto Choices of S5 Lowmass CBC Searches (2010), document LIGO-T1000056-v5 in <https://dcc.ligo.org/LIGO-T1000056-v5/public>.
- [37] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009).
- [38] LIGO Scientific Collaboration, Sensitivity to Gravitational Waves from Compact Binary Coalescences Achieved during LIGO's Fifth and Virgo's First Science Run (2010), document LIGO-T0900499-v19 in <https://dcc.ligo.org/LIGO-T0900499-v19/public>.
- [39] This is a scaled inverse-chi-squared distribution for  $\eta_j$ .
- [40] LIGO Scientific Collaboration, Advanced LIGO anticipated sensitivity curves. (2010), document LIGO-T0900288 in <https://dcc.ligo.org/LIGO-T0900288-v3/public>.