Weakly parametric approach to stochastic background inference in LISA

Federico Pozzoli,[1,*] Riccardo Buscicchio,[2,3,4,5] Christopher J. Moore,[3] Francesco Haardt,[1,4,5] and Alberto Sesana[2,4,5]

[1]*Dipartimento di Scienza e Alta Tecnologia, Università dell'Insubria, via Valleggio 11, I-22100 Como, Italy*
[2]*Dipartimento di Fisica "G. Occhialini," Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, 20126 Milano, Italy*
[3]*Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, United Kingdom*
[4]*INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, 20126 Milano, Italy*
[5]*INAF—Osservatorio Astronomico di Brera, via Emilio Bianchi 46, 23807 Merate, Italy*

(Received 11 December 2023; accepted 28 March 2024; published 22 April 2024)

Detecting stochastic gravitational wave backgrounds (SGWBs) with the Laser Interferometer Space Antenna (LISA) is one of the mission's scientific objectives. Disentangling SGWBs of astrophysical and cosmological origin is a challenging task, further complicated by the noise level uncertainties. In this study, we present a Bayesian methodology for inferring SGWBs, drawing inspiration from Gaussian stochastic processes. We assess the effectiveness of this approach for signals with unknown spectral shapes by systematically exploring the model hyperparameters—a preliminary step toward a more efficient trans-dimensional exploration. To validate our method, we apply it to a representative astrophysical scenario: the inference of the astrophysical background of extreme mass ratio inspirals, as recently estimated [F. Pozzoli *et al.*, Phys. Rev. D **108**, 103039 (2023)]. Our findings indicate that the algorithm is capable of recovering the injected signal even with uninformative priors, simultaneously providing an estimate of the noise level.

DOI: 10.1103/PhysRevD.109.083029

I. INTRODUCTION

The Laser Interferometer Space Antenna (LISA) [1] is a groundbreaking mission for the detection of gravitational waves (GWs) from space. Throughout the nominal 4 years of duration of operations, an abundance of individually resolvable sources will be available. A primary focus of LISA will be the detection of several tens of massive black-hole (MBH) binary mergers [2,3], spanning masses from approximately 10^4 to 10^7 times that of the Sun. These mergers are anticipated to have high signal-to-noise ratios (SNRs), reaching up to several thousands. Additionally, LISA will identify several to a few hundred extreme mass ratio inspirals (EMRIs) annually [4], along with tens of thousands of Galactic white dwarf binaries [5].

Besides individual sources, the Universe is pervaded by a stochastic gravitational wave background (SGWB), which arises from an incoherent superposition of GWs originating from numerous unresolved or weak sources. At nanohertz frequencies, the first evidence for an unresolved GW signal of astrophysical and/or cosmological origin [6,7] has been recently been reported by the European pulsar timing array [8], NANOGrav [9], the Parkes pulsar timing array [10] and the Chinese pulsar timing array [11]. The SGWB in LISA may come from both cosmological and astrophysical sources. Cosmologically [12], these waves could stem from primordial quantum fluctuations [13], potentially amplified during cosmic inflation. Furthermore, processes such as first-order phase transitions [14,15], interacting cosmic strings [16], and primordial black holes [17] may also contribute to stochastic GW emissions. Astrophysically, the Galactic foreground [18] from double white dwarf binaries (DWDs) is expected to dominate, particularly in the 0.5–3 mHz range. Additionally, both stellar-origin black-hole binaries (SOBHBs) [19] and EMRIs have the potential to generate non-negligible backgrounds [20,21].

The persistence of the SGWB in the collected data ties inextricably to the estimation of instrumental noise. Various efforts have been made to construct a null channel in LISA, that is a data stream effectively insensitive to GWs: however, in realistic instrumental setups, this will not be possible across the whole sensitivity band [22]. Disentangling different SGWB components and instrumental noise provides a challenging data-analysis task. Existing approaches in literature can be broadly categorized into three main classes: (i) template-based for signal and noise [15,23], (ii) template-based for signal and template-free for noise [24,25], and (iii) template-based for noise and template-free for signal [26,27]. The above algorithms may be suitable for different scenarios. While the first class of algorithms is

*fpozzoli@uninsubria.it

expected to have better resolving power—and subsequently shorter run-time—the second and third ones are important to assess the significance of results: by relaxing astrophysical and cosmological prior knowledge to less informative ones, they serve as competing models for Bayesian model selection. In this study, we introduce a weakly parametric approach by implementing an inference based on stochastic Gaussian processes, which can be applied in various contexts with great flexibility in the SGWB spectral shape.

The paper is organized as follows. In Sec. II we provide a brief overview of the putative astrophysical SGWBs potentially observable with LISA. In Sec. III we describe the underlying assumptions of our analysis. In Sec. IV a detailed exposition of the core statistical model is presented. In Sec. V, we report our parameter estimation results on a set of simulated SGWB signals. Finally, in Sec. VI, we summarize our findings and outline the prospects for future developments toward a highly flexible inference on SGWBs. Unless otherwise specified, summation over repeated indices is assumed throughout.

II. ASTROPHYSICAL BACKGROUNDS

Astrophysical SGWBs are expected to form from the superposition of GWs from compact binary systems. However, other putative astrophysical SGWBs are expected to arise, e.g. from asymmetric supernova explosions [28] or rapidly rotating neutron stars [29].

Signals typically manifest as either isolated and faint, falling below the detection threshold, or as overlapping ones. They pile up incoherently, hence resulting in a "confused" time series, whose morphology is too complex to tell individual signals apart. Moreover, the different astrophysical (and cosmological) SGWBs will overlap in the LISA data making it difficult to study them, individually [30]. However, the different spectral shapes of each background component and the different sky distributions of Galactic (localized near the plane or bulge of the Milky Way) versus extragalactic (more isotropic) backgrounds can be used to separate them. The dominant contributors to astrophysical SGWBs in the LISA sensitivity band are expected to originate mainly from three populations. We summarize briefly their properties below.

Galactic DWDs. The white dwarf binary Galactic population is expected to be formed by around $\tilde{\times}10^7$ systems [31]. A significant fraction (up to ∼10%) contributes with GW signals in the LISA band, i.e. at GW frequencies above 0.1 mHz [5]. Thousands of such highly monochromatic systems will be resolved individually, some of which are already observed through their electromagnetic counterparts; these are known as verification binaries [32–34]. Ongoing and future missions such as GAIA [35] and the Vera Rubin Observatory [36] have the potential to unveil up to a few thousand more DWD systems. The largest fraction of the DWDs will remain unresolved by LISA—only ∼0.1% of DWDs are going to be resolved individually [37]—and form a stochastic signal, commonly referred to as *confusion noise*. This is expected to be larger than LISA instrumental noise in the frequency range of 0.5–3 mHz. The bulk of such foreground arises from systems located toward the Galactic Center; hence, the confusion noise is expected to be highly anisotropic. Leveraging the SGWB sky distribution and the induced signal modulation due to LISA motion may help in characterizing the Galactic morphology [38], i.e. its disk, bulge, halo, and streams.

SOBHBs. SOBHBs will be detected by LISA during their early inspiral phase [39]. A portion of these binaries will enter the final years of their inspiral and eventually transition into the frequency range detectable by LIGO, Virgo, and KAGRA (LVK), thereby allowing for multiband gravitational-wave astronomy [40,41]. Based on current LVK constraints, it is anticipated that at least a few of these binaries will be individually detected and characterized [42–44], while the majority will contribute to the formation of a SGWB. The background is expected to be isotropic and power-law shaped with spectral index 2/3 for the energy density power spectrum (or, equivalently, $-7/3$ for the power spectral density), for quasicircular binaries at leading-order post-Newtonian approximation [45]. Recently, accurate detectability of the SOBHB background in LISA has been demonstrated [19], accounting for instrumental noise uncertainties and coexistence of DWDs foreground.

EMRIs. A stellar mass object orbiting around a MBH constitutes an EMRI. These sources have not been observed through either electromagnetic or GW radiation, yet. Nonetheless, GWs emitted by EMRIs are primary targets for LISA, due to their potential association with fast radio bursts in the electromagnetic spectrum [46]. Because of their very small mass ratios $\sim 10^{-6} - 10^{-4}$, EMRIs evolve slowly: a considerable number of these systems will persist in the LISA band throughout the entire mission. A large number of orbital cycles $\sim 10^4$–10^5 provides the opportunity to probe, in the test-particle limit, highly curved spacetime close to supermassive black holes. More recently, EMRIs have been proposed as potential targets for strong gravitational lensing [47]. As most of the population will remain unresolved, a SGWB is expected to form from these GW sources. Unlike the SOBHB population, the EMRI background spectral shape prediction is uncertain due to (i) the spread of each gravitational-wave signal over multiple harmonics of the orbital frequency and (ii) the uncertainty related to the populations. A systematic characterization of EMRI background from astrophysically motivated model populations developed in [4] was recently presented in [20]: the majority of models yield a bright SGWB with SNRs ranging between tens to thousands, contributing mostly to higher frequencies than the DWD background, namely 1–10 mHz.

III. STOCHASTIC SIGNALS IN LISA

The LISA constellation is composed of three spacecrafts arranged in a nearly equilateral configuration, respectively separated by 2.5×10^6 km. Spacecrafts will trail Earth motion by an angle of about 20°, lying on a plane inclined by 60° with respect to the orbital plane. Distances between spacecrafts will be modulated by incident GWs, inducing variations at the picometer level. Such variations will be detected by monitoring the frequency (or phase, equivalently) of laser beams exchanged between spacecrafts and comparing them to local reference lasers. The primary source of noise will come from laser frequency fluctuations, many orders of magnitude higher than the target sensitivity. To suppress it, six raw Doppler measurements are suitably delayed and linearly combined in a set of time-delay interferometric (TDI) variables [48]. Various TDI scheme exist in literature, associated to different approximations of the LISA constellation orbits; in this work we use TDI 1.5 combinations [49], assuming constant and unequal-length LISA arms. In principle, this configuration provides nonzero correlation between TDI channels [22]. However, due to computational reasons, we decided to neglect these terms.

We perform inference on such data, $\tilde{d} = (\tilde{A}, \tilde{E}, \tilde{T})$, the Fourier transforms of TDI channels. We model the data streams as the superposition of two independent stochastic processes: the SGWB and the instrumental noise. In this work we do not consider the contribution from the nonstationary Galactic foreground, whose addition we leave for future study. At each frequency f the correlation between data reads

$$\langle \tilde{d}_\alpha(f) \tilde{d}^*_\beta(f) \rangle = S_{h,\alpha\beta}(f) + S_{n,\alpha\beta}(f), \tag{1}$$

where α and β denote each TDI channel. We assume throughout isotropic, Gaussian, stationary SGWBs with equal-weighted and uncorrelated polarizations.

The first component in the previous equation reads

$$S_{h,\alpha\beta}(f) = R_{\alpha\beta}(f) S_h(f). \tag{2}$$

$S_h(f)$ represents the one-sided SGWB strain power spectral density (PSD) [50], defined as

$$\langle h_p(f;\Omega) h^*_{p'}(f',\Omega') \rangle = \frac{1}{2} S_h(f) \delta(f-f') \delta(\Omega,\Omega') \delta_{pp'} \tag{3}$$

and $R_{\alpha\beta}$ is a 3×3 matrix obtained by combining the TDI transfer matrix $M_{\mathrm{TDI}}(f)$ (3×6) and the single-link response matrix $G(f, t_0)$ (6×6) as follows:

$$R_{\alpha\beta}(f) = M_{\alpha i,\mathrm{TDI}}(f) G_{ij}(f, t_0) M^*_{j\beta,\mathrm{TDI}}(f). \tag{4}$$

For completeness, we provide full derivation of $G(f, t_0)$ in Appendix A, while M_{TDI} matches exactly the one available in literature for the chosen set of TDI variables [24].

The spectrum $S_{n,\alpha\beta}$ represents the single-sided noise power spectra in a single laser link. Following [24], we further assume no correlations between individual links and the same PSD in each of them. Hence, PSDs in each channel depend uniquely on M_{TDI} and on a common spectral model $S_n(f)$:

$$S_{n,\alpha\beta}(f) = \frac{1}{2} S_n(f) M_{\alpha i,\mathrm{TDI}}(f) M^*_{i\beta,\mathrm{TDI}}(f). \tag{5}$$

Specifically, we include in S_n secondary noise sources, namely the test mass (TM) and the optical metrology system (OMS) noise. The overall PSD is $S_n = 2S_n^{\mathrm{TM}} + S_n^{\mathrm{OMS}}$ [22], where

$$S_n^{\mathrm{TM}}(f) = A^2 \left[1 + \left(\frac{0.4\ \mathrm{mHz}}{f}\right)^2\right] \left[1 + \left(\frac{f}{8\ \mathrm{mHz}}\right)^4\right] \times \left(\frac{1}{2\pi f c}\right)^2 (10^{-30}\ \mathrm{m^2/s^3}) \tag{6}$$

and

$$S_n^{\mathrm{OMS}}(f) = P^2 \left[1 + \left(\frac{2\ \mathrm{mHz}}{f}\right)^4\right] \left(\frac{2\pi f}{c}\right)^2 (10^{-24}\ \mathrm{m^2\,s}) \tag{7}$$

with $A = 3$ and $P = 15$ [51].

Through TDI channels signal and noise PSDs, the SNR for an SGWB is readily defined as

$$\mathrm{SNR}^2 = T \int_0^\infty df \frac{S^2_{h,\alpha\alpha}}{S^2_{n,\alpha\alpha}} \tag{8}$$

with T the LISA observation time [52].

IV. INFERENCE MODEL

In the present study we only focus on the parameter estimation of SGWB and noise; thus, we assume that all nonstochastic GW sources have been perfectly subtracted from the data. We choose to model both spectra—introduced in Eqs. (2) and (5)—with a weakly parametric model, which allows us to explore a large family of spectral shapes. We do so by using expectation values of Gaussian processes. A Gaussian process (GP) is a stochastic model, formally describing distributions over functions $g(x)$. It is parametrized by a mean function $m(x)$ and positive definite covariance function $k(x, x')$. Realizations from such process are denoted as follows:

$$g(x) \sim \mathcal{GP}(m(x), k(x, x')). \tag{9}$$

The defining property of GPs is that for any finite subset $X = \{x_1 \ldots x_n\}$ of the domain for x, the marginal distribution is a multivariate Gaussian

$$\boldsymbol{g}(X) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \tag{10}$$

with mean vector and covariance matrix defined by $\mu = m(x_i)$ and $\Sigma_{ij} = k(x_i, x_j)$, respectively. Typically, GPs are used to flexibly incorporate some observed data in the distribution over $\boldsymbol{g}(X)$ and make predictions on new domain points X_* as $\boldsymbol{g}(X_*)$. The joint probability distribution $p(\boldsymbol{g}(X_*), \boldsymbol{g}(X))$ is a multivariate normal distribution. In order to employ this formalism as a regression model, one needs the conditional probability $p(\boldsymbol{g}(X_*)|\boldsymbol{g}(X))$. Upon suitable marginalization, the conditional distribution is also a multivariate normal distribution:

$$p(\boldsymbol{g}(X_*)|\boldsymbol{g}(X)) = \mathcal{N}(\boldsymbol{\mu}(X_*|X), \boldsymbol{\Sigma}(X_*|X)), \tag{11}$$

where

$$\boldsymbol{\mu}(X_*|X) = \boldsymbol{\mu}(X_*) + \boldsymbol{\Sigma}(X_*|X)\boldsymbol{\Sigma}(X,X)^{-1}(\boldsymbol{g}(X) - \boldsymbol{\mu}(X)) \tag{12}$$

and

$$\boldsymbol{\Sigma}(X_*|X) = \boldsymbol{\Sigma}(X_*, X_*) + \boldsymbol{\Sigma}(X_*, X)\boldsymbol{\Sigma}^{-1}(X,X)\boldsymbol{\Sigma}(X, X_*) \tag{13}$$

are specified by the single kernel function $k(x, x')$. The matrix $\boldsymbol{\Sigma}$ is also referred to as *kernel matrix*, and it models the covariance between each pair of its two arguments through the definition of the bivariate function k. In order to be a valid covariance for the multivariate Gaussian in Eq. (11), the kernel matrix must be symmetric and positive definite. A variety of kernels are available in literature, to capture different processes peculiarities. In this work, we consider the radial basis kernel function (RBF) defined by

$$\Sigma^{\mathrm{RBF}}_{ij}(X, Y) = k_{\mathrm{RBF}}(x_i, y_j) = \exp\left(-\frac{|x_i - y_j|^2}{2\sigma^2}\right), \tag{14}$$

where σ is a model positive hyperparameter.

In the present study, we exploit GPs in an unconventional way: we use each GP expectation value from Eq. (12) (hereinafter EGP) as a proposed spectrum and set $\boldsymbol{g}(X)$ as free parameters. The GP covariance is not used in the inference model. The dimensionality and domain location of X are to be considered hyperparameters.

Henceforth, we model the logarithm of noise and signal PSDs as two independent EGPs over the logarithmic frequencies, i.e. $X = \{\log_{10} f_1, \ldots, \log_{10} f_n\}$. We illustrate both in Secs. IV A and IV B, respectively. We will refer to the parameters X and to the mean function $m(X)$ as *knots* and *baseline*, respectively.

A. Signal

We assume the response matrix $R_{\alpha\beta}(f)$ to be exactly known. Thus, we need to specify $S_h(f)$ only. We use an EGP, with baseline and knots left free to vary simultaneously. We do so to be able to capture both global spectral shapes and fine structure in specific frequency regions: the choice of baseline influences globally the spectral-shape proposals, while knots control the fine-local structure. In Fig. 1, we show the proposed EGP models for two baseline families, i.e. power-law and *bump* spectra. The knot variability can compensate easily for misidentification of the truth baseline family. In our inference, we use a power-law baseline family for $\log_{10} S_h(f)$, parametrized by logarithmic amplitude $\log_{10} A$ at a reference frequency $f_\star = 10^{-2.5}$ Hz and slope γ. Thus, $\boldsymbol{\mu}(X_*)$ in Eq. (12) reads $\log_{10} A + \gamma \log_{10}(f/f_\star)$. Fine-structure deviations with respect to the baseline are parametrized through knots $\boldsymbol{\delta}^h$, defined by

$$\begin{aligned} g(X)_k - \mu(X)_k &= \log_{10}(S_h(f_k) \cdot \delta^h_k) - \log_{10} S_h(f_k) \\ &= \log_{10} \delta^h_k. \end{aligned} \tag{15}$$

The number of knots is fixed in each inference, and we choose the associated frequencies f_k equally log-spaced within the data frequency range. Overall the model is summarized by functions from the parametric family

$$\begin{aligned} &\log_{10} S_h(f; \log_{10} A, \gamma, \boldsymbol{\delta}^h) \\ &= \log_{10} A + \gamma \log_{10}(f/f_\star) \\ &\quad + k(f, f_j) k(f_j, f_k)^{-1} \log_{10} \delta^h_k. \end{aligned} \tag{16}$$

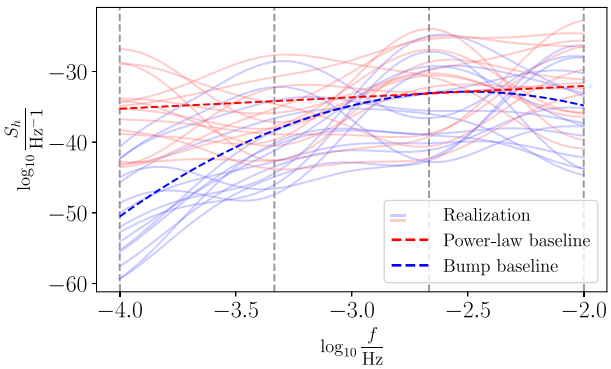

FIG. 1. Realizations from two EGP models, with fixed baseline and knots amplitudes sampled from their respective parameter space. Dashed gray lines represent the knot locations. EGP realizations are shown with solid lines, while dashed red (blue) lines correspond to the fixed power-law (bump) baseline.

B. Noise

We follow a similar approach to parametrize and infer upon the noise uncertainties. Again, we assume a well-known TDI transfer matrix $M_{\rm TDI}$, so the modeling freedom is left for the single-link noise PSD. Specifically, we introduce a set of parameter $\log_{10}\boldsymbol{\delta}^n$. Contrary to the signal case discussed in Sec. IV A, the baseline is fixed to the nominal PSD level from ESA's science requirement document [51]. Thus, the final noise model reads

$$\log_{10} S_n(f;\boldsymbol{\delta}^n) = \log_{10} S_n^{\rm SciRD} + k(f,f_j)k(f_j,f_k)^{-1}\log_{10}\delta_k^n. \quad (17)$$

With this parametrization, we augment the noise reference model with considerable flexibility to vary across a large functional space. In all our inference the injected noise in the data is generated according to the fixed baseline, so we anticipate recovering zero values for $\log_{10}\boldsymbol{\delta}^n$. In Fig. 2, we show a comparison between GP- and EGP-like instrumental noise modeling, illustrating the different degree of flexibility each one exhibits.

C. Likelihood

Inference is performed simultaneously on SGWB and instrumental noise. We construct joint posteriors on parameters $\boldsymbol{\theta}$:

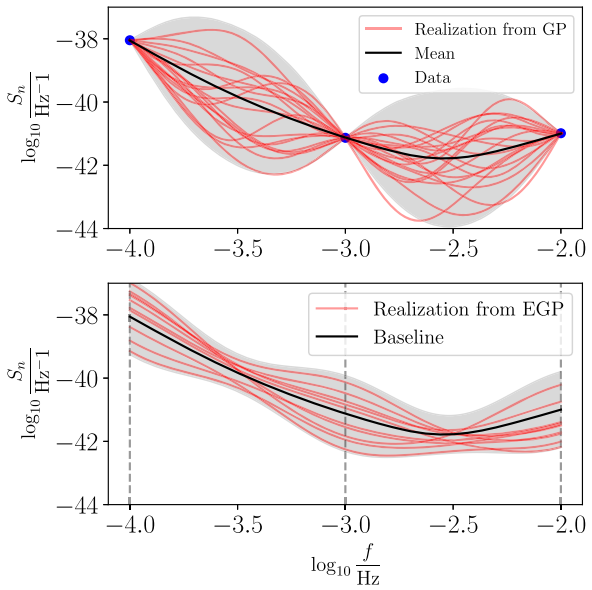

FIG. 2. Noise realizations from GP and EGP model. Top panel: realizations from GP. The gray shaded region represents the 90% credible interval. Bottom panel: realizations from EGP. The gray area represents the prior 90% credible interval. While GP allows for straightforward inclusion of observed data, EGP smoothly explores a large set of spectral shapes, with the *data* (i.e. knot amplitudes at chosen locations shown as gray dashed vertical lines) effectively proposed as inference parameters.

$$p(\boldsymbol{\theta}|\tilde{d}) \propto \mathcal{L}(\tilde{d}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}) \quad (18)$$

through stochastic sampling of the likelihood $\mathcal{L}(\tilde{d}|\boldsymbol{\theta})$ under chosen priors $\pi(\boldsymbol{\theta})$. To do so, we use Balrog, a large codebase for simulation and inference on LISA signals. In this study, we use it alongside a nested sampling algorithm as implemented in Nessai [53] to obtain each posterior and evidence. Data are distributed according to the Gaussian likelihood [54]

$$\log\mathcal{L}(\tilde{d}|\boldsymbol{\theta}) = -\frac{1}{2}\sum_{\alpha,\beta}\sum_{f=f_{\rm min}}^{f_{\rm max}} \times \frac{\tilde{d}_\alpha(f)\tilde{d}^*_\beta(f)}{S_{n,\alpha\beta}(f,\boldsymbol{\delta}^n)+S_{h,\alpha\beta}(f,\log A,\gamma,\boldsymbol{\delta}^h)}\Delta f, \quad (19)$$

with $\boldsymbol{\theta} = (\log_{10}A, \gamma, \log_{10}\boldsymbol{\delta}^h, \log_{10}\boldsymbol{\delta}^n)$. We choose uniform priors for each parameter in $\boldsymbol{\theta}$, and the following prior ranges:

(i) power-law amplitude $\log_{10}A$: $[-70,-35]$;
(ii) power-law slope γ: $[-5,5]$;
(iii) signal knot amplitudes $\log_{10}\boldsymbol{\delta}^h$: $[-2,2]$;
(iv) noise knot amplitudes $\log_{10}\boldsymbol{\delta}^n$: $[-0.5,0.5]$.

Additionally, we will consider a prior uniform in $[-5,5]$ for $\log_{10}\boldsymbol{\delta}^h$, instrumental to the results discussion in Sec. V.

We highlight once again that the number of knots $\boldsymbol{\delta}^h$ and $\boldsymbol{\delta}^n$ can be further optimized, as much as the remaining hyperparameters (i.e. the kernel length scale, its functional form, and the knot locations). This represent a potential high degree of flexibility to our algorithm, which can be leveraged to find unexpected spectral features upon inference. In this study, we focus on a finite set of kernel lengths σ and number of knots. Noise parameter estimation is instead fixed to three knots, given the baseline choice matching the noise model used for injection in the data.

In absence of a transdimensional sampling framework [55] in our codebase, our study is constrained to perform inference for each choice of hyperparameters and select the most suitable via Bayesian model selection. In order to compare the different choices, we use the log-Bayes factor

$$\log_{10}\mathcal{B}^{\sigma,n}_{\sigma',n'} = \log_{10}\mathcal{Z}_{\sigma,n} - \log_{10}\mathcal{Z}_{\sigma',n'}, \quad (20)$$

where (σ,n) and (σ',n') are labels identifying competing models with different kernel length scale and number of knots, respectively, and

$$\mathcal{Z} = \int d\boldsymbol{\theta}\mathcal{L}(\tilde{d}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}) \quad (21)$$

is the Bayesian evidence. In Sec. V we present our results, together with a map of evidence: they are to be interpreted as the model marginal likelihoods, conditioned on hyperparameters, only. This is a core building block toward a transdimensional exploration of the full EGP

model for SGWB inference and its inclusion in a larger global fit scheme [56,57].

V. INFERENCE RESULTS

In this section, we present parameter estimation results performed using the model introduced in Sec. IV. First, we use a simple toy model to control the correct recovery of a SGWB with known spectral shape (Sec. V A). Then, we apply our formalism to the parameter estimation of an SGWB from a population of unresolved EMRIs (Sec. V B). Throughout the analysis, LISA data are simulated for 4 years of observation time. We proceed with the assumption that nonstochastic GWs sources and glitches have been subtracted perfectly. Additionally, we consider the absence of data gaps, resulting in the generation of an idealized residual dataset.

A. Toy model

We conduct a test retrieving an EGP model with a power-law baseline and nonzero knot amplitudes.

Although the injected signals lack an astrophysical interpretation, they serve as a compelling test for our method's capability of recovering complex spectral shapes. We inject separately two SGWBs, with SNR of 209 and 5, respectively. We choose the following parameters for each signal:

(i) power-law baseline with $\log_{10}A = -43$, $\gamma = -2$ (SNR $= 209$) and $\log_{10}A = -44.5$, $\gamma = -1.5$ (SNR $= 5$);
(ii) knot amplitudes $\log_{10}\boldsymbol{\delta}^h = [0.6, -0.2, -0.3, 0.4]$, identical in both cases.

We choose the parameters as representative of two SGWBs, above and below the nominal power-law sensitivity curve [58] to an SGWB with SNR of 10 after 4 years of observation, respectively. We emphasize once more that the injected model is an EGP with local features arising with respect to a simple power law. In this specific context, the Power-law sensitivity is not a faithful indicator of a signal detectability but should be interpreted as a rough reference level.

Spectral inference results and parameter posteriors are shown in Figs. 3 and 8 in Appendix B, respectively. The EGP model is able to capture the spectral shape of the high SNR injected signal, and posterior distributions are consistent with the injected values within their 90% credible intervals. The posterior distributions for the noise parameters are consistent with zero, as expected. If the signal SNR is too low, as for the second simulated SGWB (SNR $= 5$), the posterior distribution for its parameters is effectively an upper limit, only.

B. Astrophysical case: EMRI SGWB

We further apply our method to the astrophysical case of an SGWB from EMRIs. We inject a background signal chosen from a set available in literature [20]: twelve background signals from astrophysically motivated populations [4] were generated (therein referred to as M1–M12). The stochastic signal is constructed through suitable processing of whole population signals and subtraction of resolvable sources via an iterative algorithm [59]. Each source GW strain is computed using the state-of-the-art EMRI waveforms [60] and the resulting LISA signal is simulated through an accurate response model. Resulting SGWBs exhibit a wide range of amplitudes, though the majority of them provide an SGWB potentially detectable with LISA over 4 years of observation time.

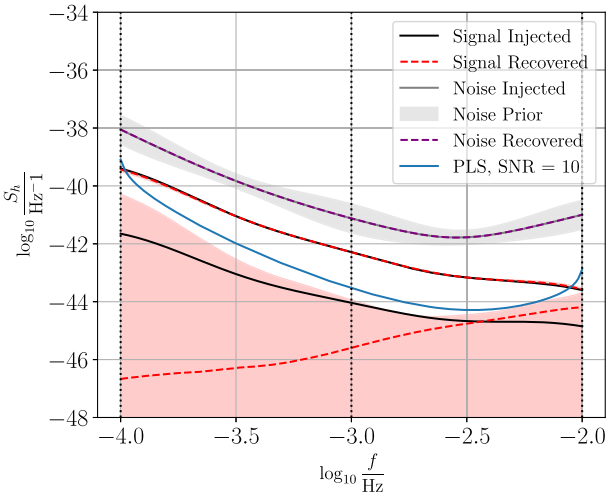

FIG. 3. Spectral reconstruction of two toy model SGWBs. Injected SGWBs are denoted by black solid lines. They have SNR of 209 and 5, respectively. This places them above and below the nominal power-law sensitivity curve (solid blue line) to an SGWB with SNR of 10 after 4 years of observation. Noise injection and prior 90% confidence intervals are shown as gray solid and purple dashed lines, respectively. Posterior median and 90% confidence intervals on each signal—analyzed independently—are denoted by red dashed lines and light-red shared areas, respectively. The EGP flexibility captures features in the PSD shape.

In this study, we use M1 as our fiducial model, as it provides an intermediate SGWB amplitude across the populations studied in [20]. Because of uncertainties on spectral morphology and population distribution, our flexible inference method is suitable to analyze such a signal. In order to inject coherently the signal in the Balrog simulator, we divide the A channel signal realization obtained by [20] by the response function, following the formalism in Eq. (2). Then, we smooth the resulting spectrum it to obtain a reference model PSD. Subsequently, we regenerate each TDI channel data and use it for inference.

1. Exploring hyperparameter space

We perform a set of Bayesian inferences, exploring two hyperparameters for the signal model: the number of knots,

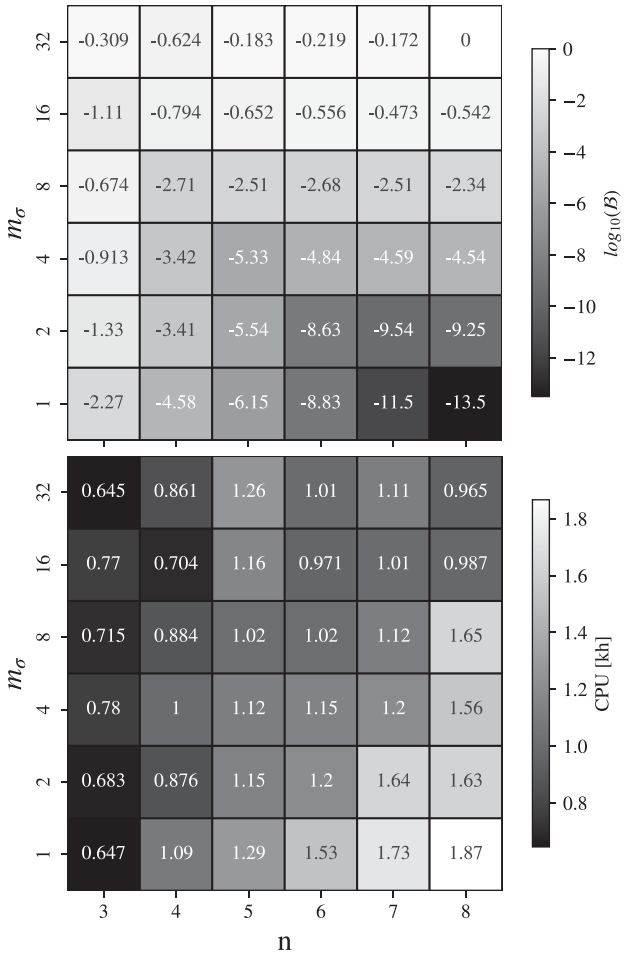

FIG. 4. Top panel: grid of logarithmic Bayes factor for different pairs of number of signal knots (n) and integers of kernel length scale σ. The numbers are computed in relation to the model that has the highest evidence, specifically $(n, m_\sigma) = (8, 32)$. We observe a diagonal pattern: as m_σ decreases and n increases, the model becomes progressively less preferred. Bottom panel: computational cost (i.e. sampling time as reported by Nessai) of parameter estimations as function of integer multiple of σ, m_σ (with $\sigma = 0.6$) and the number of knots of the signal model, n. The numbers are in unit of kilohours. Each inference is run in multithreading over 40 cores.

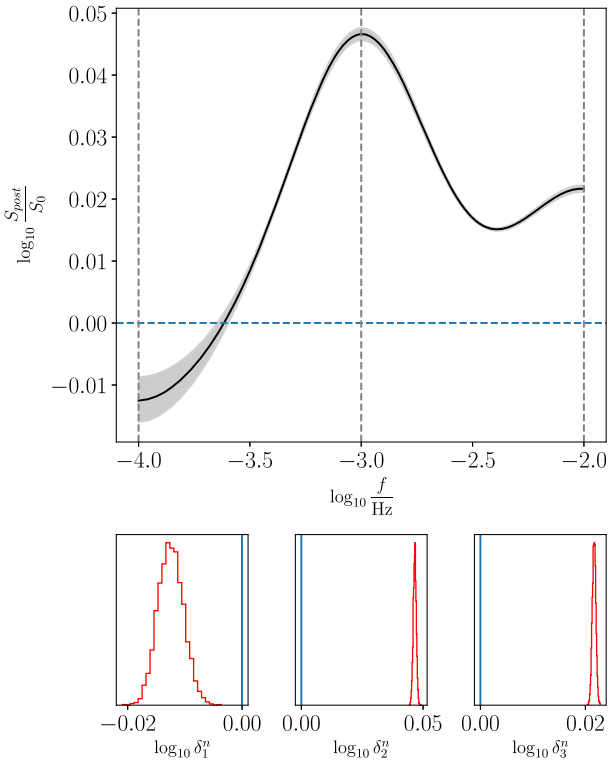

FIG. 5. Posterior distribution for the inference on EMRI background, with a noise-only model. The noise estimation employs three knots, with frequencies corresponding to vertical dashed gray lines in the main panel. Bottom subpanels show posteriors on individual knot amplitudes. The injected spectral noise model corresponds to $\boldsymbol{\delta}^n = \mathbf{0}$ and is denoted by solid blue lines. The noise model can compensate with significant biases the presence of an SGWB in the data. Nonetheless, its evidence is significantly lower when compared to the noise and signal models, as illustrated in Fig. 4.

$n = 3, 4, \ldots, 8$, and the kernel length scale, as integer m_σ multiples of $\sigma = 0.6$, with $m_\sigma = \{1, 2, 4, 8, 16, 32\}$.

We evaluate each evidence through stochastic nested sampling. Results are presented in Fig. 4 (top panel), as Bayes factors $\log_{10}\mathcal{B}$ with respect to the highest-evidence model, i.e. $(n, m_\sigma) = (8, 32)$. They reveal a prevailing trend: as m_σ increases, the number of knots become less effective at influencing each model evidence, resulting in Bayes factors increasingly close to each as function of m_σ.

For completeness, in Fig. 4 (bottom panel) we also provide the computational cost of each inference, measured in CPU core kilohours (kh). The observed trend is the result of two competing factors: as the number of parameters (i.e. the number of knots n) increases, the inference process would take longer. However, if the increased model complexity (e.g. shorter kernel length scale) is not required by the simulated data, exploration of parameter space is typically fast. As a result, we observe a pattern in computational cost broadly similar to the log Bayes factors: as the number of knots (m_σ) increases (decreases), the computational cost diminishes.

As a further check we perform inference for the null-hypothesis $\mathcal{H}_0$, where we exclusively model the data as instrumental noise. Results reveal decisive evidence, in favour of any signal-including model (n, m_σ) with respect to the noise-only model $\mathcal{H}_0$, with log-Bayes factors largely greater than 10^3. As expected, the spectral noise reconstruction in Fig. 5 shows significant biases to compensate for the unmodeled SGWB signal in the data.

2. Spectral reconstruction

We discuss here in greater detail the signal spectral reconstruction of each model. Our findings are best

described by a comparison between the inferences with (n, m_σ) equal to (8,1) and (8,32), whose results are shown in top and bottom panels of both Figs. 6 and 7, respectively. We show the reconstruction of the fiducial EMRI background with and without T channel data in Figs. 6 and 7, respectively. The T channel is known to be less sensitive to GWs than the A and E channels at the frequencies of our analyses, therefore acting as a potential noise estimator. We observe however that the spectral reconstructions with and without the T channel are broadly consistent, with broader posteriors when the T channel is excluded, as expected. This is due to the common noise modeling of the single-link PSD spectrum, which is then propagated coherently to each channel through $M_{\rm TDI}$ Eq. (5). Relaxing the exact knowledge of the transfer matrix is expected to significantly broaden the noise PSD posterior.

A comparison between top and bottom panels reveals an important result, originating from the competition in Bayes factors between excess model complexity and the effectiveness of our chosen parametrization. The inference presented in this study is inherently integrated across a large

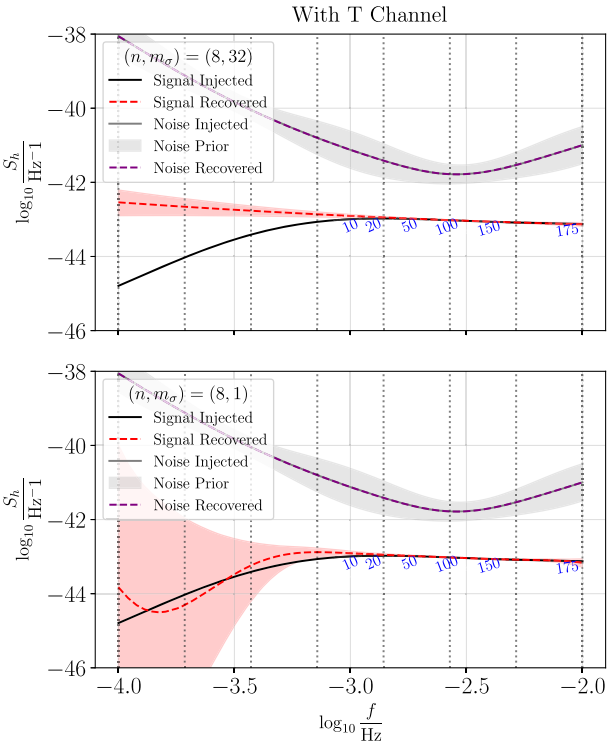

FIG. 6. Signal reconstruction of the M1 EMRI background carried out using the EGP parametrization. In both figures, the noise estimation employs three knots. Top panel: the model features eight knots for the signal and a kernel length scale that is 32 times larger than σ. This configuration aligns with the model exhibiting the highest evidence. Bottom panel: the model features eight knots for the signal, and the kernel length scale is set equal to σ. This configuration corresponds to the model with the lowest evidence. Blue numbers denote the cumulative SNR integrated from 10^{-4} Hz up to the frequency they are located at.

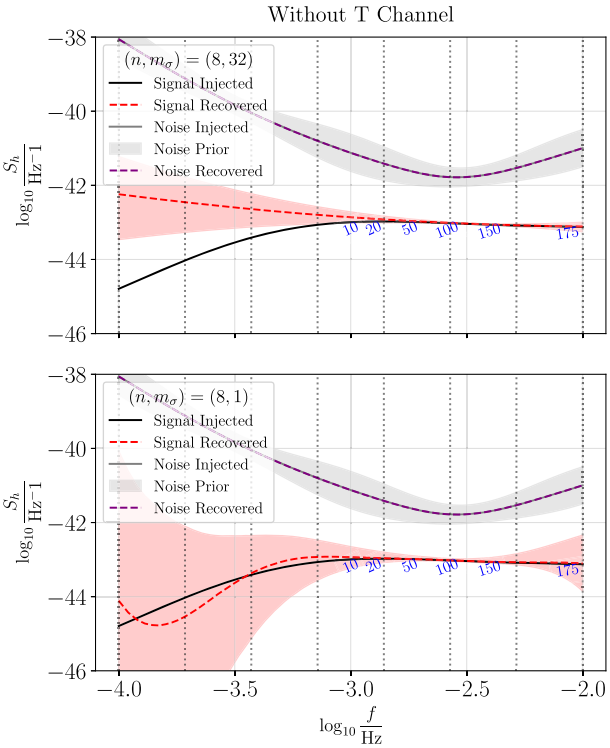

FIG. 7. Analog illustration of Fig. 6 where the inference has been conducted without incorporating the T channel in the data. Consequently, the resulting posterior regions exhibit greater dispersion, signifying reduced informativeness. Blue numbers denote the cumulative SGWB SNR integrated from 10^{-4} Hz up to the frequency they are located at.

frequency band: we annotate the injected SGWB with blue values for cumulative SNRs with respect to the true noise level on a set of representative frequencies. It is readily apparent that 80% of the SNR is accumulated between 1 and 4 mHz: in this frequency range the reconstruction is consistent with the injected spectrum. The dashed red lines and red shaded area denote the posterior median and 90% confidence interval, respectively.

At the lowest frequencies, the preferred model (top panel, $n, m_\sigma = 8, 32$) exhibits larger biases as compared to the disfavored one (bottom panel, $n, m_\sigma = 8, 1$). This is due to the narrow-band ($m_\sigma = 1$) fluctuations being highly disfavored by the data in the high SNR region, where the spectrum is close to a power-law model, a shape already described by the baseline parameters. On the contrary broadband ($m_\sigma = 32$) fluctuations can be absorbed by an adjustment of the baseline parameters, while keeping an accurate spectral reconstruction. However, the preference for $m_\sigma = 32$ has an important side effect: m_σ is fixed and equal for every knot in each analysis, including the four at the lowest frequencies. Therefore, the model with the highest evidence is not capable of introducing local fluctuations to capture significant tilts with respect to the baseline power law. These are instead provided by the disfavored model.

Figures 9 and 10 further support our interpretation, where we compare the posterior distributions for the two runs shown in Fig. 6 (red contours) with two more differing only by a broader prior for the knot amplitudes $\boldsymbol{\delta}^h$, uniform in $[-5, 5]$ (blue contours). The inference results for the noise are largely unaffected and uncorrelated to the signal parameters (and hyperparameters). However, we find that relaxing the knot priors affects the posterior on the baseline parameters: a sign of strong correlation. Moreover, posteriors on $\boldsymbol{\delta}^h$ for the preferred runs (Fig. 9) are prior dominated, while the likelihood becomes informative in the disfavored models $(n, m_\sigma) = (8, 1)$ (Fig. 10). In conclusion we argue that in a realistic scenario models more flexible than EGP should be developed, alongside astrophysically motivated priors to avoid introducing—even in an evidence-based model selection—unwanted biases.

3. T channel influence

Given its reduced sensitivity to the SGWB, we further investigate the impact of the T channel on parameter estimation. We anticipate that excluding the T channel from data analysis will result in less informative posterior distributions (as depicted in Fig. 7). To be more specific, we seek to assess whether hyperparametrization modeling remains influential, as depicted in Fig. 4.

To achieve this, we examine two models characterized by the highest and lowest values of evidence (or hyperlikelihood), denoted as $(n, m_\sigma) = (8, 32)$ and $(n, m_\sigma) = (8, 1)$, respectively. The calculated logarithmic Bayes factor stands at -10.95, underscoring a substantial strength of evidence in favor of one model over the other. In other words, the outcome suggests that the trend observed in Fig. 4 is robust and not significantly affected by variations in our ability to estimate the noise level.

VI. CONCLUSION

In this study we present a novel method to estimate parameters of an SGWB observed by LISA. We implement our analysis in Balrog, a broader data analysis framework for LISA. This is a building block toward the completion of a full global fit of deterministic and stochastic GW signals, alongside a flexible noise model. We employ a versatile parametrization to infer the SGWB of uncertain spectral shape using EGPs. Our method is able to capture unexpected features in the signal shape, while assessing their significativity in a purely Bayesian formalism. We test our algorithm using a simplified toy model, where the true injected model is by construction within the parameter space explored.

Following, we apply the analysis to the expected signal from a population of unresolved EMRIs. We leverage the model flexibility to explore the hyperparameter space and optimize them based on evidence. We also investigate the impact of T channel on the inference results. We find that the EGP model has great flexibility at reconstructing the SGWB spectral shape. However, while Bayesian evidence might be suitable to select models in the high SNR region, it may disfavor less biased models in the low SNR region, at the price of excess complexity where it is not needed. Therefore, the model inherent nonlocality should be treated with caution when interpreting spectral results. We also find that the T channel has a significant impact on the inference results, and its exclusion leads to naturally less constrained results.

Our findings suggest that parametric or weakly parametric models (as the EGP) should be considered a first step toward even more agnostic and realistic approaches. Including cross-correlation terms between TDI channels is expected to provide a more coherent inference given the orbits and TDI generation chosen, and it is a natural extension of our work. Finally, for specific SGWBs a parameter estimation modeling nonstationarities and non-Gaussianities would extract valuable information from the data. This is a crucial improvement, to allow one to reduce biases and help disentangle superimposed SGWBs in real LISA data.

ACKNOWLEDGMENTS

We wish to thank Jonathan Gair, Antoine Klein, Golam Shaifullah, Germano Nardini, Chiara Caprini, Mauro Pieroni, Nikos Karnesis, Fabio Rigamonti, and all Balrog developers for fruitful discussions. We also wish to thank the University of Geneva for the kind hospitality. R. B. acknowledges the support of the ICSC National Research Center funded by NextGenerationEU, and the Italian Space Agency Grant *Phase A activity for LISA mission*, Agreement No. 2017-29-H.0. C. J. M. acknowledges the support of the United Kingdom Space Agency, Grant No. ST/V002813/1. A. S. acknowledges the financial support provided under the European Union's H2020 ERC Consolidator Grant "Binary Massive Black Hole Astrophysics" (B Massive, Grant Agreement No. 818691). Computational resources were provided by University of Birmingham BlueBEAR High Performance Computing facility, by CINECA through EuroHPC Benchmark access call grant "LISAFIT" (Proposal ID No. EHPC-BEN-2023B08-024), and by the Google Cloud Research Credits program with Award No. GCP19980904. We acknowledge usage of *Mathematica* [61] and of the following Python [62] packages for modeling, analysis, postprocessing, and production of results throughout: Nessai [53], Matplotlib [63], NumPy [64], and SciPy [65].

APPENDIX A: SGWB RESPONSE FUNCTION

For completeness and consistency of definitions, we provide here a derivation of the one-arm LISA response—induced by a GW propagating along $\hat{k}$—as experienced

along a single link identified by the sender and receiver spacecraft s and r and the unit vector pointing from the former to the latter, i.e. $\hat{r}_{sr}$. Henceforth, no summation over repeated s and r is assumed. Each fractional arm length variation reads

$$W_{sr}(t,\hat{k}) = \frac{\delta L_{sr}(t)}{L_{sr}} = \frac{1}{2}\hat{r}^a_{sr}\hat{r}^b_{sr}\int_{-\infty}^{\infty}\tilde{h}_{ab}(f)\tau_{sr}(f,t,\hat{k})e^{2\pi i f(t-\Delta t)}df, \tag{A1}$$

where

$$\Delta t = \frac{1}{2}\left[\hat{k}\cdot(\vec{x}_s+\vec{x}_r)+L_{sr}\right] \tag{A2}$$

and

$$\tau_{sr}(f,t,\hat{k}) = \mathrm{sinc}(\pi L_{sr} f[1-\hat{k}\cdot\hat{r}_{sr}(t)]), \tag{A3}$$

where $\vec{x}_s$ and $\vec{x}_r$ denote the sender and receiver spacecraft, respectively. Due to the motion of the satellites, the arm direction vector $\hat{r}$ varies with time. The GW strain can be written in TT gauge as

$$h_{ab}(t,\hat{k}) = h_+(t)\epsilon^+_{ab}(\hat{k}) + h_\times(t)\epsilon^\times_{ab}(\hat{k}), \tag{A4}$$

where $\epsilon^{+,\times}$ denote a linear polarization-tensor basis. The stochastic strain is the linear superposition of plane tensor waves from each direction of the sky; therefore, we define an integrated response to each polarization $p=+,\times$, as follows:

$$W^p_{sr}(t) = \int_{\hat{k}} W^p_{sr}(t,\hat{k})d^2\hat{k}, \tag{A5}$$

where

$$W^p_{sr}(t,\hat{k}) = \frac{1}{2}\hat{r}^a_{sr}\hat{r}^b_{sr}\int_{-\infty}^{\infty}\tilde{h}_p(f,\hat{k})\epsilon^p_{ab}(\hat{k})\tau(f,t,\hat{k})e^{2\pi i f(t-\delta t)}df = \frac{1}{2}\int_{-\infty}^{\infty}\tilde{h}_p(f,\hat{k})\xi^p(\hat{k},\hat{r}_{sr})\tau(f,t,\hat{k})e^{2\pi i f(t-\delta t)}df. \tag{A6}$$

In Eq. (A5) we introduced the antenna pattern functions $\xi^p(\hat{k},\hat{r}_{sr})$. They read as

$$\xi^+(\hat{k},\hat{r}_{sr}) = (\hat{u}_k\cdot\hat{r}_{sr})^2-(\hat{v}_k\cdot\hat{r}_{sr})^2,$$
$$\xi^\times(\hat{k},\hat{r}_{sr}) = 2(\hat{u}_k\cdot\hat{r}_{sr})(\hat{v}_k\cdot\hat{r}_{sr}). \tag{A7}$$

We rewrite Eq. (A6) equivalently as

$$W^p_{sr}(t,\hat{k}) = \int_{-\infty}^{\infty}\tilde{h}_p(f,\hat{k})\bar{G}^p_{sr}(f,t,\hat{k})e^{2\pi i f t}df, \tag{A8}$$

where $\bar{G}^p_{sr}(f,t,\hat{k})$ reads

$$\bar{G}^p_{sr}(f,t,\hat{k}) = \frac{1}{2}\xi^p(\hat{k},\hat{r}_{sr})e^{-i\pi f(L_{sr}+\hat{k}\cdot(x_s+x_r))}\tau_{sr}(f,t,\hat{k}). \tag{A9}$$

Noteworthy, $\bar{G}^p_{sr}(f,t,\hat{k})$ is a complex quantity unlike $\tau(f,t,\hat{k})$. Equivalently, we recast Eq. (A9) as a response in fractional frequency domain:

$$G^p_{sr}(f,t,\hat{k}) = 2\pi i L_{sr} f\bar{G}^p_{sr}(f,t,\hat{k}). \tag{A10}$$

The Fourier transform of (A8) is then

$$\begin{aligned} W^p_{sr}(f,\hat{k}) &= \int_{-\infty}^{\infty}dt e^{-2\pi i f t}W^p_{sr}(t,\hat{k}) \\ &= \int_{-\infty}^{\infty}dt e^{-2\pi i f t}\int_{-\infty}^{\infty}\tilde{h}_p(f',\hat{k})\bar{G}^p_{sr}(f',t,\hat{k})e^{2\pi i f' t}df' \\ &= \int_{-\infty}^{\infty}dt\int_{-\infty}^{\infty}\tilde{h}_p(f',\hat{k})\bar{G}^p_{sr}(f',t,\hat{k})e^{-2\pi i t(f-f')}df' \\ &= \int_{-\infty}^{\infty}df'\bar{G}^p_{sr}(f',f-f',\hat{k})\tilde{h}_p(f',\hat{k}); \end{aligned} \tag{A11}$$

note that

$$\bar{G}^p_{sr}(f',f,\hat{k}) = \int_{-\infty}^{\infty}e^{-2\pi i f t}\bar{G}^p_{sr}(f',t,\hat{k})dt. \tag{A12}$$

As described in Eq. (3) an isotropic, stationary, zero-mean background is characterized entirely by the strain covariance which allows for the definition of a response matrix for two given links sr and $s'r'$ reads

$$\begin{aligned} \Sigma_{sr,s'r'}(f) &= \langle W^p_{sr}(f)W^{*p}_{s'r'}(f)\rangle \\ &= \int_{\hat{k}}d^2k\int_{-\infty}^{\infty}df'\bar{G}^p_{sr}(f',f-f',\hat{k}) \\ &\quad\times\bar{G}^{*p}_{s'r'}(f',f-f',\hat{k})S_h(f'). \end{aligned} \tag{A13}$$

For an isotropic distribution of sources, integrating the product of $\bar{G}$'s and evaluating it at a reference time t_0—without loss of generality given the assumption of perfect stationarity—we obtain, thus,

$$\Sigma^p_{sr,s'r'}(f) = S_h(f)G^p_{sr,s'r'}(f), \tag{A14}$$

with

$$G^p_{sr,s'r'}(f) = \int_{\hat{k}}d^2k\bar{G}^p_{sr}(f,t_0,\hat{k})\bar{G}^{*p}_{s'r'}(f,t_0,\hat{k}), \tag{A15}$$

which corresponds to G_{ij} in Eq. (4) upon remapping of the indices $sr, s'r'$ to ij.

APPENDIX B: POSTERIOR DISTRIBUTION FOR THE TOY MODEL

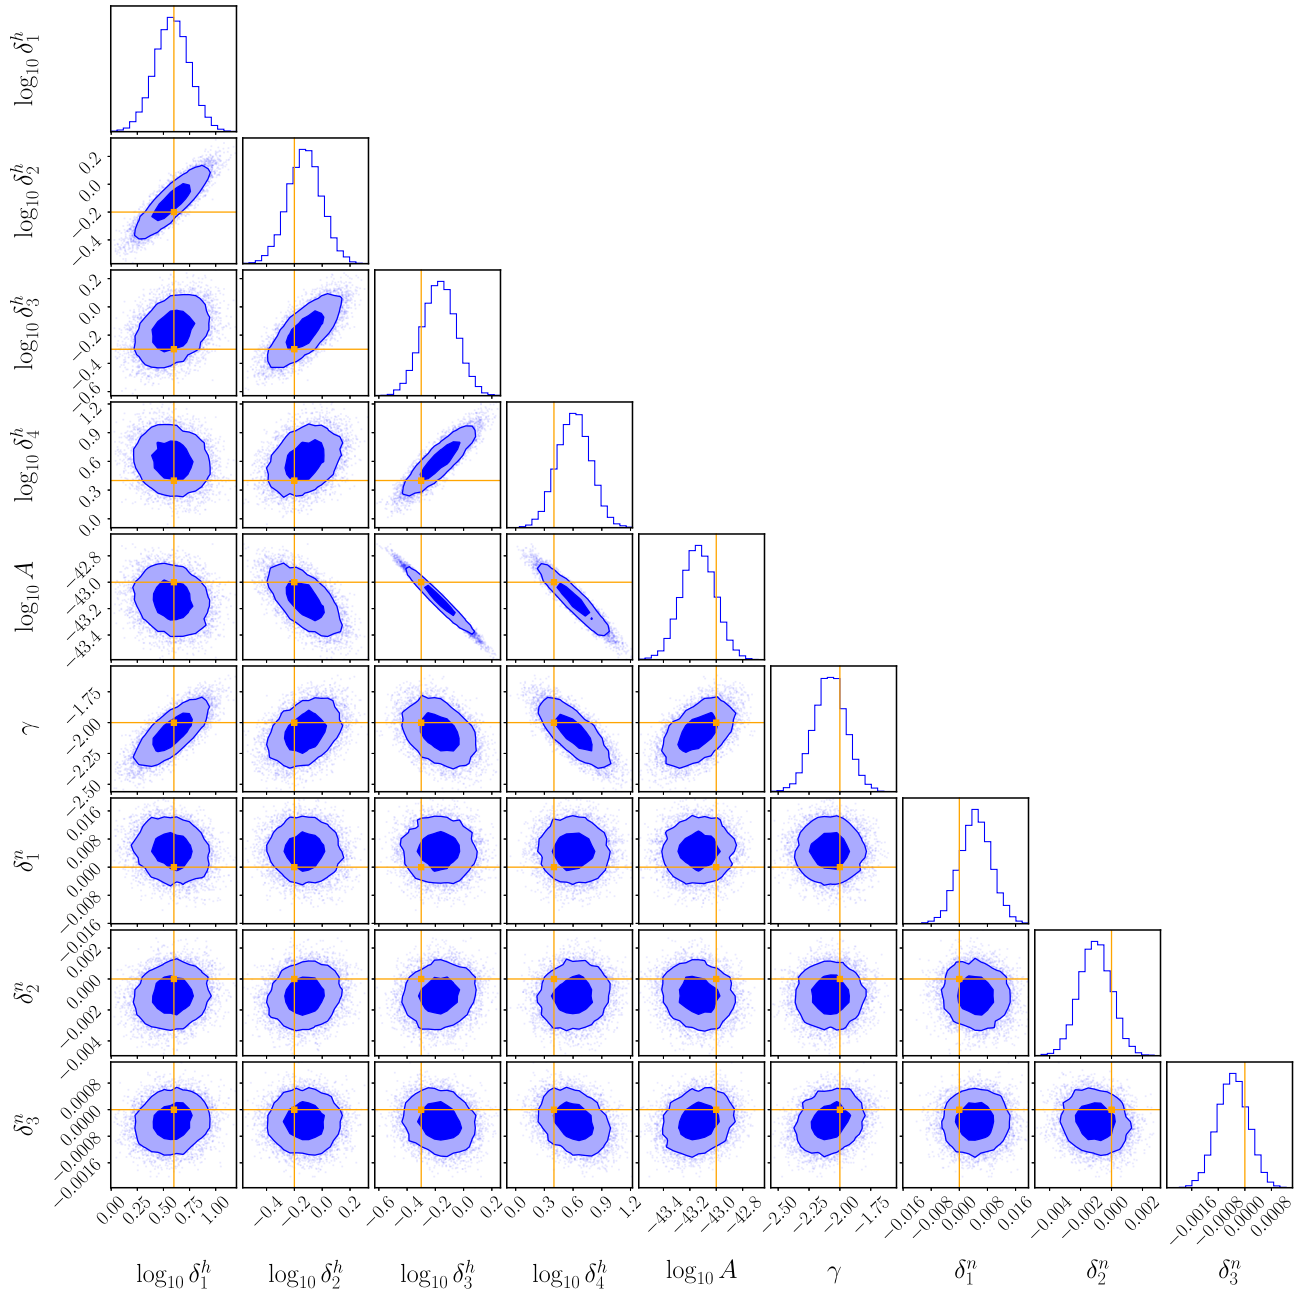

FIG. 8. Posterior parameter distribution for the toy-model SGWB with SNR = 209 described in Sec. V A, obtained with an EGP model inference. The signal injected parameters are $\{\log_{10}\boldsymbol{\delta}^h, \log_{10}A, \gamma\} = \{[0.6, -0.2, -0.3, 0.4], -43, -2\}$. The injected noise level matches the nominal LISA sensitivity curve [51]. Darker (lighter) shaded areas denote 90% (50%) credible regions, and solid orange lines indicate the injected values. Posteriors are consistent with the injected values at the 90% credible level, and the noise parameters are consistent with zero with accuracies between 10^{-3} and 10^{-2}.

APPENDIX C: POSTERIOR DISTRIBUTIONS FOR MODEL M1 SGWB

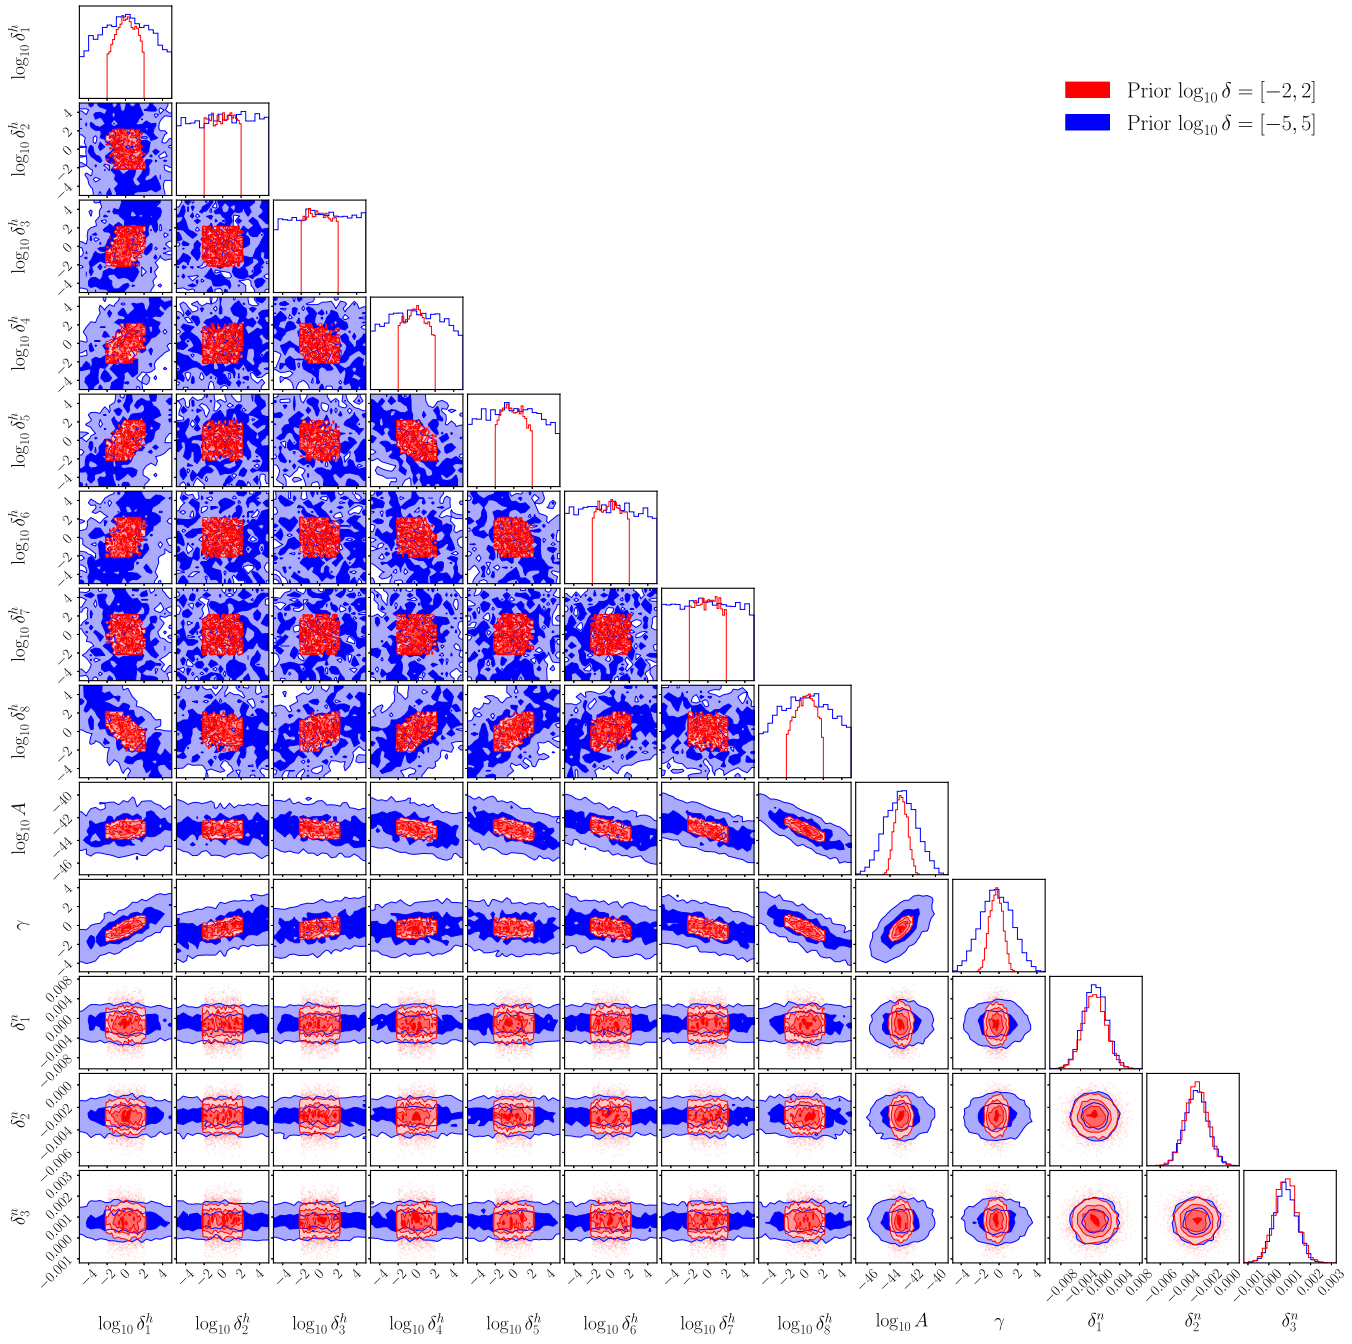

FIG. 9. Posterior parameter distribution for M1 EMRI SGWB for model $(n, m_\sigma) = (8, 32)$ using a larger and narrower prior on $\boldsymbol{\delta}^h$. Posteriors shown in this plot correspond to the spectral reconstructions in top panels of Figs. 6 and 7, respectively. They support the result interpretation provided in Sec. V B 2.

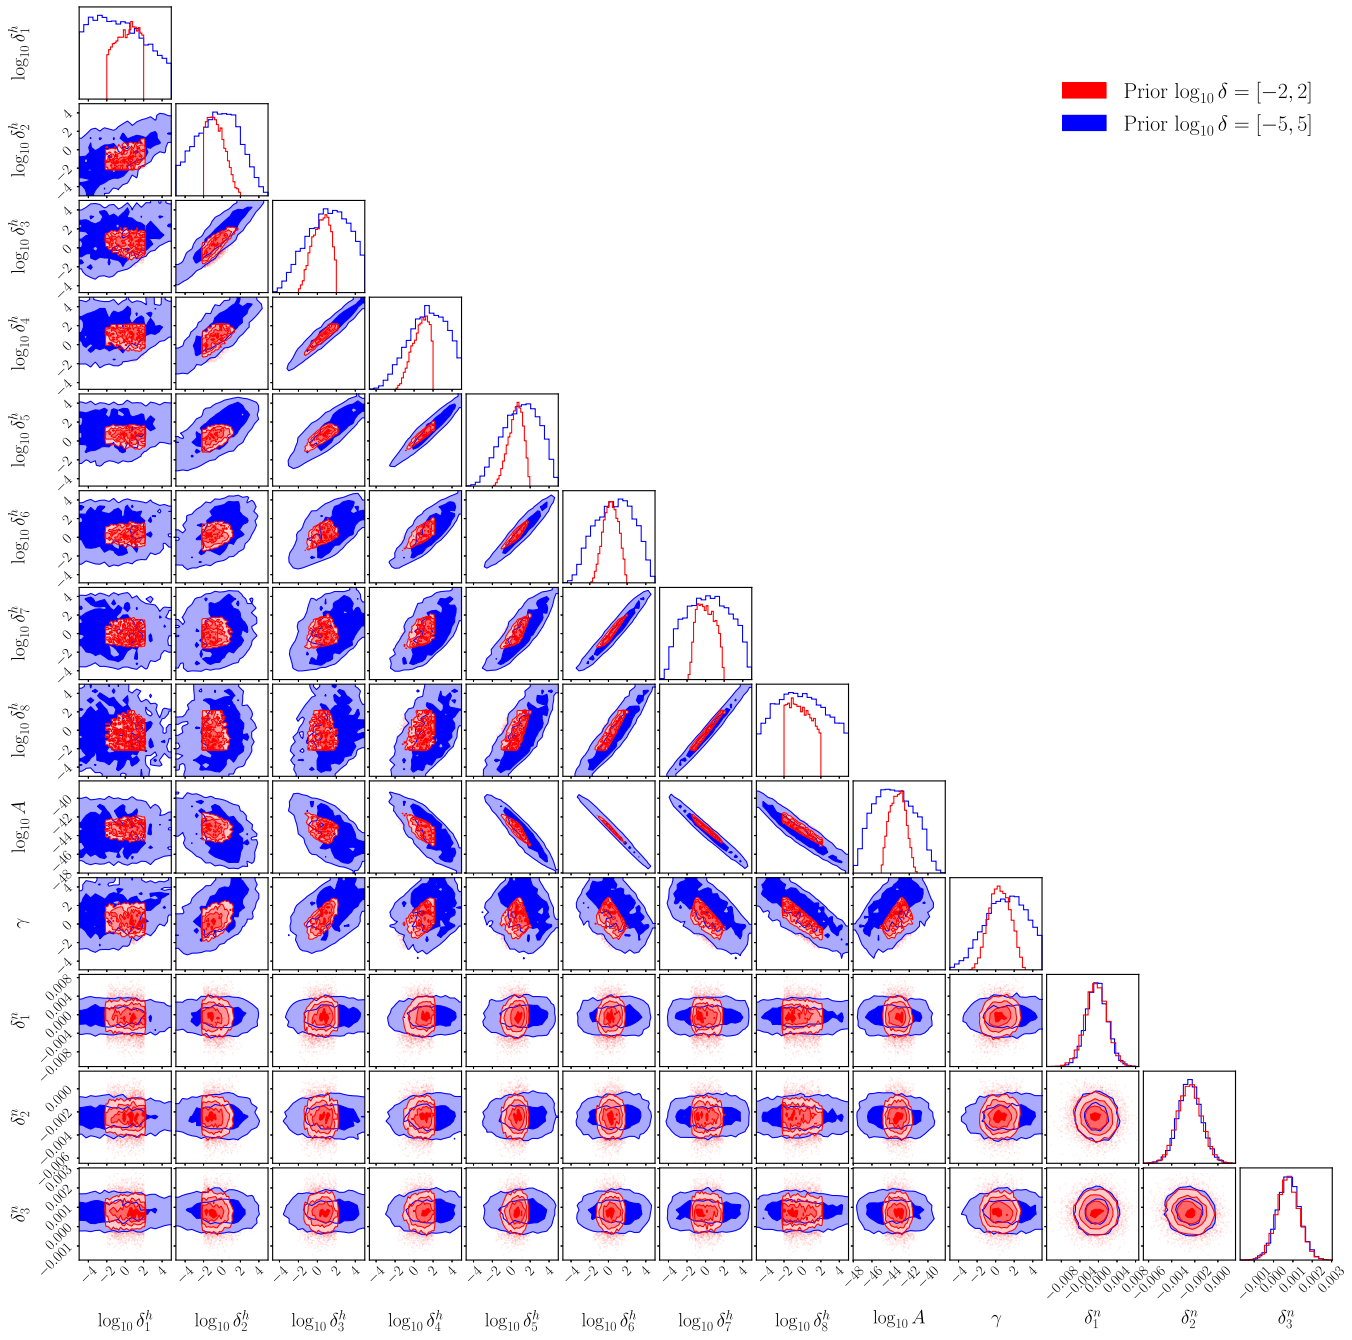

FIG. 10. Posterior parameter distribution for M1 EMRI SGWB for model $(n, m_\sigma) = (8, 1)$ using a larger and narrower prior on $\boldsymbol{\delta}^h$. Posteriors shown in this plot correspond to the spectral reconstructions in bottom panels of Figs. 6 and 7, respectively. They support the result interpretation provided in Sec. V B 2.

[1] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker *et al.*, arXiv:1702.00786.
[2] A. Klein, E. Barausse, A. Sesana, A. Petiteau, E. Berti, S. Babak, J. Gair, S. Aoudia, I. Hinder, F. Ohme, and B. Wardell, Phys. Rev. D **93**, 024003 (2016).
[3] A. Sesana, F. Haardt, P. Madau, and M. Volonteri, Astrophys. J. **611**, 623 (2004).
[4] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. L. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, Phys. Rev. D **95**, 103012 (2017).
[5] V. Korol, S. Toonen, A. Klein, V. Belokurov, F. Vincenzo, R. Buscicchio, D. Gerosa, C. J. Moore, E. Roebber, E. M. Rossi, and A. Vecchio, Astron. Astrophys. **638**, A153 (2020).
[6] J. Antoniadis, P. Arumugam, S. Arumugam, P. Auclair *et al.*, arXiv:2306.16227.
[7] A. Afzal, G. Agazie, A. Anumarlapudi, A. M. Archibald *et al.*, Astrophys. J. Lett. **951**, L11 (2023).
[8] J. Antoniadis, P. Arumugam *et al.* (EPTA and InPTA Collaborations), Astron. Astrophys. **678**, A50 (2023).
[9] G. Agazie, A. Anumarlapudi, A. M. Archibald, Z. Arzoumanian *et al.*, Astrophys. J. Lett. **951**, L8 (2023).
[10] D. J. Reardon, A. Zic, R. M. Shannon, G. B. Hobbs *et al.*, Astrophys. J. Lett. **951**, L6 (2023).
[11] H. Xu, S. Chen, Y. Guo, J. Jiang *et al.*, Res. Astron. Astrophys. **23**, 075024 (2023).
[12] C. Caprini, J. Phys. Conf. Ser. **610**, 012004 (2015).
[13] A. Ricciardone, J. Phys. Conf. Ser. **840**, 012030 (2017).
[14] C. Gowling and M. Hindmarsh, J. Cosmol. Astropart. Phys. 10 (2021) 039.
[15] G. Boileau, N. Christensen, C. Gowling, M. Hindmarsh, and R. Meyer, J. Cosmol. Astropart. Phys. 02 (2023) 056.
[16] P. Auclair, J. J. Blanco-Pillado, D. G. Figueroa, A. C. Jenkins *et al.*, J. Cosmol. Astropart. Phys. 04 (2020) 034.
[17] S. Clesse and J. García-Bellido, Phys. Dark Universe **18**, 105 (2017).
[18] G. Nelemans, Classical Quantum Gravity **26**, 094030 (2009).
[19] S. Babak, C. Caprini, D. G. Figueroa, N. Karnesis, P. Marcoccia, G. Nardini, M. Pieroni, A. Ricciardone, A. Sesana, and J. Torrado, J. Cosmol. Astropart. Phys. 08 (2023) 034.
[20] F. Pozzoli, S. Babak, A. Sesana, M. Bonetti, and N. Karnesis, Phys. Rev. D **108**, 103039 (2023).
[21] M. Bonetti and A. Sesana, Phys. Rev. D **102**, 103023 (2020).
[22] O. Hartwig, M. Lilley, M. Muratore, and M. Pieroni, Phys. Rev. D **107**, 123531 (2023).
[23] M. R. Adams and N. J. Cornish, Phys. Rev. D **89**, 022001 (2014).
[24] Q. Baghi, N. Karnesis, J.-B. Bayle, M. Besançon, and H. Inchauspé, J. Cosmol. Astropart. Phys. 04 (2023) 066.
[25] M. Muratore, J. Gair, and L. Speri, Phys. Rev. D **109**, 042001 (2024).
[26] C. Caprini, D. G. Figueroa, R. Flauger, G. Nardini, M. Peloso, M. Pieroni, A. Ricciardone, and G. Tasinato, J. Cosmol. Astropart. Phys. 11 (2019) 017.
[27] R. Flauger, N. Karnesis, G. Nardini, M. Pieroni, A. Ricciardone, and J. Torrado, J. Cosmol. Astropart. Phys. 01 (2021) 059.
[28] R. Rosca-Mead, M. Agathos, C. J. Moore, and U. Sperhake, Phys. Rev. D **107**, 124040 (2023).
[29] P. A. Rosado, Phys. Rev. D **86**, 104007 (2012).
[30] G. Boileau, A. Lamberts, N. Christensen, N. J. Cornish, and R. Meyer, arXiv:2105.06793.
[31] S. E. Timpano, L. J. Rubbo, and N. J. Cornish, Phys. Rev. D **73**, 122001 (2006).
[32] A. Stroeer and A. Vecchio, Classical Quantum Gravity **23**, S809 (2006).
[33] E. Finch, G. Bartolucci, D. Chucherko, B. G. Patterson, V. Korol, A. Klein, D. Bandopadhyay, H. Middleton, C. J. Moore, and A. Vecchio, Mon. Not. R. Astron. Soc. **522**, 5358 (2023).
[34] T. Kupfer, V. Korol, T. B. Littenberg, S. Shah *et al.*, Astrophys. J. **963**, 100 (2024).
[35] T. Prusti, J. H. J. de Bruijne, A. G. A. Brown *et al.* (Gaia Collaboration), Astron. Astrophys. **595**, A1 (2016).
[36] P. A. Abell, J. Allison, S. F. Anderson *et al.* (LSST Science Collaboration), arXiv:0912.0201.
[37] M. Georgousi, N. Karnesis, V. Korol, M. Pieroni, and N. Stergioulas, Mon. Not. R. Astron. Soc. **519**, 2552 (2023).
[38] K. Breivik, C. M. F. Mingarelli, and S. L. Larson, Astrophys. J. **901**, 4 (2020).
[39] A. Sesana, A. Lamberts, and A. Petiteau, Mon. Not. R. Astron. Soc. **494**, L75 (2020).
[40] A. Sesana, Phys. Rev. Lett. **116**, 231102 (2016).
[41] A. Klein, G. Pratten, R. Buscicchio, P. Schmidt, C. J. Moore, E. Finch, A. Bonino, L. M. Thomas, N. Williams, D. Gerosa, S. McGee, M. Nicholl, and A. Vecchio, arXiv:2204.03423.
[42] R. Buscicchio, A. Klein, E. Roebber, C. J. Moore, D. Gerosa, E. Finch, and A. Vecchio, Phys. Rev. D **104**, 044065 (2021).
[43] A. Toubiana, S. Babak, S. Marsat, and S. Ossokine, Phys. Rev. D **106**, 104034 (2022).
[44] M. C. Digman and N. J. Cornish, Phys. Rev. D **108**, 023022 (2023).
[45] E. S. Phinney, arXiv:astro-ph/0108028.
[46] R.-N. Li, Z.-Y. Zhao, Z. Gao, and F.-Y. Wang, Astrophys. J. Lett. **956**, L2 (2023).
[47] M. Toscani, O. Burke, C. Liu, N. B. Zamel, N. Tamanini, and F. Pozzoli, Phys. Rev. D **109**, 063505 (2024).
[48] M. Tinto and S. V. Dhurandhar, Living Rev. Relativity **8**, 4 (2005).
[49] M. Tinto, F. B. Estabrook, and J. W. Armstrong, Phys. Rev. D **69**, 082001 (2004).
[50] C. J. Moore, R. H. Cole, and C. P. L. Berry, Classical Quantum Gravity **32**, 015014 (2015).
[51] LISA Science Study Team, LISA science requirements document, Technical Report No. 1.0, ESA, 2018.
[52] T. L. Smith and R. R. Caldwell, Phys. Rev. D **100**, 104055 (2019).
[53] M. J. Williams, J. Veitch, and C. Messenger, Phys. Rev. D **103**, 103006 (2021).
[54] C. Cutler and É. E. Flanagan, Phys. Rev. D **49**, 2658 (1994).
[55] P. J. Green, Biometrika **82**, 711 (1995).
[56] N. Karnesis, M. L. Katz, N. Korsakova, J. R. Gair, and N. Stergioulas, Mon. Not. R. Astron. Soc. **526**, 4814 (2023).
[57] N. J. Cornish and T. B. Littenberg, Classical Quantum Gravity **32**, 135012 (2015).

[58] E. Thrane and J. D. Romano, Phys. Rev. D **88**, 124032 (2013).
[59] N. Karnesis, S. Babak, M. Pieroni, N. Cornish, and T. Littenberg, Phys. Rev. D **104**, 043019 (2021).
[60] M. L. Katz, A. J. K. Chua, L. Speri, N. Warburton, and S. A. Hughes, Phys. Rev. D **104**, 064047 (2021).
[61] Wolfram Research Inc., *Mathematica* (2022).
[62] G. Van Rossum and F. L. Drake, *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).
[63] J. D. Hunter, Comput. Sci. Eng. **9**, 90 (2007).
[64] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers *et al.*, Nature (London) **585**, 357 (2020).
[65] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland *et al.*, Nat. Methods **17**, 261 (2020).