Flexible, GPU-accelerated approach for the joint characterization of LISA instrumental noise and stochastic gravitational wave backgrounds

Alessandro Santini,[1,*] Martina Muratore,[1] Jonathan Gair,[1] and Olaf Hartwig[1,2,3]

[1]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, 14476 Potsdam, Germany*

[2]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstraße 38, 30167 Hannover, Germany*

[3]*Leibniz Universität Hannover, Institut für Gravitationsphysik, Callinstraße 38, 30167 Hannover, Germany*

(Received 25 July 2025; accepted 25 September 2025; published 17 October 2025)

LISA data analysis represents one of the most challenging tasks ahead for gravitational wave (GW) astronomy. Characterizing the instrument's noise properties while fitting for all the other detectable sources is a key requirement of any robust inference pipeline. Noise estimation will also play a crucial role in searches and parameter estimation of cosmological and astrophysical stochastic signals. Previous studies have tackled this topic by assuming perfect knowledge of the spectral shape of the instrumental noise and of different possible types of stochastic GW backgrounds (SGWBs), usually resorting to parametrized templates. Recently, various works that employ template-agnostic methods have been presented. In this work, we take an additional step further, introducing flexible spectral shapes in both the instrumental noise and the stochastic signals. We account for the lack of knowledge of the exact shape of the individual contributions to the overall power spectral density by using splines to represent arbitrary perturbations of the noise and signal spectral densities. We implement a data-driven reversible jump Markov chain Monte Carlo algorithm to fit different components simultaneously and to infer the level of flexibility required under different scenarios. We test this approach on simulated LISA data produced under different assumptions. We investigate the impact of this increased flexibility on the reconstruction of both the injected signal and the noise level, and we discuss the prospects for claiming a successful SGWB detection.

DOI: 10.1103/csx9-9trp

I. INTRODUCTION

The Laser Interferometer Space Antenna (LISA) [1] is a joint ESA-NASA space-based gravitational wave (GW) detector that is expected to be launched in the mid-2030s. Operating in the millihertz region of the frequency spectrum, it will enable the observation of sources in a completely unexplored band, complementing ground-based GW detectors [2–5] and pulsar timing arrays [6–9] observations. Over the mission duration, LISA is expected to observe a wide range of individually resolvable sources, including massive black-hole binary (MBHB) mergers [10], Galactic white dwarf binaries (WDBs) [11], extreme mass ratio inspirals (EMRIs) [12], and stellar-mass black-hole binaries (SBHBs) in their early inspiral phase [13–15].

In addition, LISA will be sensitive to stochastic gravitational wave backgrounds (SGWBs) of astrophysical or cosmological origin. Astrophysical stochastic backgrounds are generated by the incoherent superposition of GW signals from a large number of individual sources. Astrophysical populations that could produce detectable stochastic backgrounds for LISA include Galactic [16] and extragalactic [17] WDBs, SBHBs [18], and EMRIs [19–21]. The WDB background, in particular, is expected to dominate over the expected instrumental noise [22] between $\sim$0.5 and $\sim$3 mHz. SGWBs in the LISA band could also be of cosmological origin [23]. Possible cosmological sources include first order phase transitions, which are expected in many theories that extend physics beyond the Standard Model [24–26], cosmic strings [27], primordial fluctuations in the early Universe [28], and primordial black holes [29–31].

The wide variety of sources populating the LISA band represents an exciting and unique opportunity to explore the GW universe with unprecedented precision. However, it also poses significant data analysis challenges. Besides the instrumental noise and stochastic backgrounds, which will

*Contact author: alessandro.santini@aei.mpg.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

always be present in data we receive from the instrument, many individually resolvable sources will be long lived and will remain in band for months or years. This scenario will require a global analysis of the data, where the parameters of all the transient and stochastic components will have to be estimated simultaneously. The development of the necessary data analysis pipeline, referred to in the LISA community as the "global fit", is currently an active area of research. Early implementations are described in Refs. [32–35].

In this work, we focus on the task of simultaneously characterizing the instrumental noise and possible SGWBs present in the data. Previous studies have tackled this problem by using relatively simple parametrized templates with 2–12 parameters to describe the spectral shape of the instrumental noise and of the different SGWBs considered (see, e.g., Refs. [36–38]). Past works also include a templated analysis of the WDB background anisotropy [39]. Nevertheless, we expect that the instrumental noise entering the LISA data streams will be significantly more complex and hard to fully model analytically. The LISA-Pathfinder (LPF) mission [40,41] has already shown a similar behavior, with a significant fraction of the total measured noise that could not be explained by on-board measurements of individual noise components. As the goal of LPF was to test the behavior of (part of) the key LISA hardware, we can already expect the same to happen for LISA, which is a significantly more complex system than LPF. Therefore, we must be prepared to account for uncertainties in the noise models we will ultimately use in data analysis pipelines, which will come at the cost of affecting the detection and characterization of SGWBs [42]. Moreover, the different SGWBs described are expected to have different spectral shapes. Fitting them would require the inclusion of models for all of them in our pipelines, which is not only impractical but would also require a complete SGWB template bank. Since our knowledge of the exact spectral shape of the different SGWBs is limited, we cannot compile a complete catalogue of all the possible sources and associated spectral shapes. For this reason, robust and flexible signal reconstruction algorithms are of pivotal importance to successfully extract the information carried by these persistent signals.

Some previous works have explored possible solutions to the two problems. In Refs. [43,44], the authors propose an automated pipeline based on a binned power-law signal reconstruction, where the frequency band is divided into bins. In each bin, the data are fitted with a power-law signal plus a model of the instrumental noise. The latter is parametrized and common across all the bins. Conversely, in Ref. [45], the authors propose a flexible model for the instrumental noise, where the spectral shape is described by a set of splines, maintaining a parametrized description of the SGWB signals. Finally, in Ref. [46], the authors propose a flexible model for both the instrumental noise and the SGWB signals, where their weakly parametric models rely on expectation values of Gaussian processes [47]. In all these works, the authors assume the perfect knowledge of how noise and SGWBs are converted in the time-delay interferometry (TDI) variables [48] used to represent, and analyze, the data. A description of time-delay interferometry can be found in Sec. II A.

With the methodology proposed in this paper, we take an additional step further, lifting the crucial assumption of knowing exactly how the different noise sources propagate in the chosen TDI basis. This is done by introducing spline components directly at the TDI level. Our model builds on the work of Ref. [42], where the authors include realizations of cubic splines in the noise model to account for possible deviations from the baseline spectral density. Those splines had a fixed number of knots placed at fixed frequencies, while their amplitudes were left free to vary. There, the goal was to assess the impact of the noise knowledge uncertainties on the measurement precision of the parameters of SGWBs with different shapes. This was done relying on the widely employed Fisher information matrix formalism [49,50]. Here, we extend the idea to a full Bayesian analysis pipeline, capable of estimating not only the amplitudes of the spline knots but also the required number and location, as well as the parameters of the injected noise model and SGWB signals.

This work is part of a broader, collaborative effort involving the authors of Refs. [43–46]. There, our goal is to compare the performance of different pipelines when analyzing the same datasets. Those results are presented separately.

The remainder of this paper explores the following questions:

(i) Can we successfully use flexible, data-driven approaches to model noise and signal components at the same time?
(ii) Could a flexible instrumental noise model account for the totality of the data, mistakenly absorbing also stochastic contributions from GW origin? In this case, what would be our chances of confidently claiming an SGWB detection?

We describe the properties of the instrumental noise and signal considered in this work in Sec. II. We introduce the details of our inference pipeline in Sec. III, while in Sec. IV we present our results. Finally, we draw our conclusions and explore possible future research directions in Sec. V.

II. GRAVITATIONAL WAVE SIGNALS AND NOISE MODELS

In this section, we report the properties of the instrumental noise and SGWBs considered in this work. We then combine them to describe the datasets simulated for the analysis. In particular, we describe in Secs. II B and II C the properties of the parametrized noise and SGWB templates, respectively, that we use as baselines in our flexible model.

A. Time-delay interferometry

The LISA constellation will be composed of three identical spacecraft forming a quasi-equilateral triangle. The arm length between the three satellites is 2.5 million kilometers, and they are interconnected by six active laser links. Because of the motion of the satellites, the constellation presents unequal and time-varying arm lengths [51], which makes the suppression of laser frequency noise in the measurement channels highly nontrivial [52,53]. To mitigate this, a postprocessing method called time-delay interferometry (TDI) is employed [48]. This technique works by combining raw phasemeter data on the ground and applying suitable time delays to simulate an interferometer with equal arm lengths, thereby effectively suppressing by eight orders of magnitude the laser frequency noise. TDI combinations can be of different generations. First-generation TDI variables can suppress laser frequency noise for a constellation with static unequal arm lengths [54], while second-generation variables can also be used for time-varying arm lengths [55,56]. Other dominant noise contributions that have to be subtracted from the data are the clock and the tilt-to-length (TTL) noises [57–59].

In this work, we assume that the aforementioned noises have already been suppressed [60–62]. We also adopt the idealized configuration of LISA forming a perfect equilateral triangle, which allows us to assume static and equal arm lengths [63]. This simplification is valid in the context of analyzing secondary noise contributions, i.e., those not eliminated by the initial noise reduction pipelines [64–67], and for modeling the instrument's response to GWs under the assumptions introduced in this work. We still employ second-generation TDI variables to lay the ground for future, more realistic investigations. We assume a noise model made of two components: the test mass (TM) acceleration noise and the optical metrology system (OMS) noise [68].

While many different TDI combinations can be (and have been) constructed [53,65,69], the Michelson variables X, Y, and Z are widely employed in the community [70–73]. These can be further combined to form the three noise-orthogonal A, E, T channels [63,74]. In principle, this TDI basis does not present cross-correlations between the different channels, significantly simplifying GW likelihoods and speeding up their evaluation. However, this orthogonality holds only under the assumptions of equal arm lengths, uncorrelated noise contributions, and identical OMS and TM noises in all the links. This TDI basis is constructed as a linear combination of the Michelson channels,

$$A = \frac{Z - X}{\sqrt{2}}, \qquad E = \frac{X - 2Y + Z}{\sqrt{6}}, \qquad T = \frac{X + Y + Z}{\sqrt{3}}.$$

On the other hand, the second-generation TDI variable X_2 is defined as

$$\begin{aligned} X_2 = {} & (1 - D_{12}D_{21} - D_{12}D_{21}D_{13}D_{31} \\ & + D_{13}D_{31}D_{12}D_{21}D_{12}D_{21}) \times (\eta_{13} + D_{13}\eta_{31}) \\ & - (1 - D_{13}D_{31} - D_{13}D_{31}D_{12}D_{21} \\ & + D_{12}D_{21}D_{13}D_{31}D_{13}D_{31}) \times (\eta_{12} + D_{12}\eta_{21}). \end{aligned} \tag{1}$$

In the above, the quantities $\eta_{ij}(t)$ represent the single-link measurements between spacecrafts [65], describing the relative phase difference between the laser beam emitted from spacecraft j at time $t - \tau_{ij}$ and the laser local to spacecraft i at time t. Here, $\tau_{ij} = L_{ij}/c$ represents the light travel time between spacecrafts j (the emitter) and i (the receiver). D_{ij} is, instead, an operator that applies a constant delay equal to τ_{ij}, with

$$D_{ij}x(t) = x(t - \tau_{ij}). \tag{2}$$

In the frequency domain, its effect can be expressed as a multiplication by the phase factor $\mathcal{F}(D_{ij}) = e^{-i2\pi f\tau_{ij}}$. In this work, we assume equal and constant arm lengths, so that $L_{ij} = L = 2.5 \times 10^9$ m for all the i, j indices. The other two TDI variables, Y_2 and Z_2, are defined similarly to X_2 by cyclic permutations of the three satellites.

B. Parametrized noise model

Denoting the time series of the TM and OMS noises in the single link η_{ij} by $x^g_{ij}(t)$ and $x^m_{ij}(t)$, respectively, the noise in a single-link measurement can be expressed as

$$\tilde{\eta}^N_{ij}(f) = \tilde{x}^g_{ji}(f)e^{-if\tau_{ij}} + \tilde{x}^g_{ij}(f) + \tilde{x}^m_{ij}(f), \tag{3}$$

where $\tilde{\eta}^N_{ij}(f)$ denotes the total noise in $\eta_{ij}(t)$. This expression shows that each single-link measurement includes TM noise contributions from both spacecraft involved, leading to correlations between measurements taken at either end of the same arm. Specifically, the TM noise introduces a nonzero cross-correlation between the two corresponding link measurements [42].

In this work, we assume a noise curve [70,75] such that the TM acceleration component is described by

$$\begin{aligned} \langle \tilde{x}^g_{ij}(f)\tilde{x}^{g*}_{lm}(f')\rangle &= \frac{1}{2}\delta_{il}\delta_{jm}\delta(f - f')S_{\mathrm{TM}}(f) \\ S_{\mathrm{TM}}(f) &= A^2_{\mathrm{TM}}\left[1 + \left(\frac{0.4\,\mathrm{mHz}}{f}\right)^2\right]\left[1 + \left(\frac{f}{8\,\mathrm{mHz}}\right)^4\right] \\ &\quad \times \left(\frac{1}{2\pi f c}\right)^2. \end{aligned} \tag{4}$$

We set the test mass noise amplitude to $A_{\mathrm{TM}} = 2.4 \times 10^{-15}$ ms^{-2} Hz$^{-1/2}$ [70]. Conversely, the OMS component reads

$$\langle \tilde{x}^m_{ij}(f)\tilde{x}^{m*}_{lm}(f')\rangle = \frac{1}{2}\delta_{il}\delta_{jm}\delta(f-f')S_{\rm OMS}(f)$$
$$S_{\rm OMS}(f) = A^2_{\rm OMS}\left[1+\left(\frac{2\ \rm mHz}{f}\right)^4\right]\left(\frac{2\pi f}{c}\right)^2. \quad (5)$$

We set the OMS noise amplitude to $A_{\rm OMS} = 7.9\times 10^{-12}\ \rm m\,Hz^{-1/2}$ [70]. Asterisks represent complex conjugates, and we use the notation $\langle\cdot\rangle$ to indicate expectation values. $S_{\rm TM}$ and $S_{\rm OMS}$ represent the one-sided power spectral densities (PSDs) of the two components. We assume that the PSDs of all six test mass noise terms and, separately, all six OMS noise terms are the same, and the noise components are independent. These assumptions mean that the noises in the A, E, and T channels will be uncorrelated.

The A, E, and T one-sided power spectral densities are built applying the TDI transfer functions to the noise model,

$$\begin{aligned}S_{n,AA}(f) &= S_{n,EE}(f)\\ &= 32\sin^2\left(\frac{2\pi fL}{c}\right)\sin^4\left(\frac{2\pi fL}{c}\right)\\ &\times\left\{\left[2+\cos\left(\frac{2\pi fL}{c}\right)\right]S_{\rm OMS}(f)\right.\\ &\left.+4\left[1+\cos\left(\frac{2\pi fL}{c}\right)+\cos^2\left(\frac{2\pi fL}{c}\right)\right]S_{\rm TM}(f)\right\},\end{aligned} \quad (6a)$$

$$\begin{aligned}S_{n,TT}(f) &= 64\sin^2\left(\frac{2\pi fL}{c}\right)\sin^4\left(\frac{2\pi fL}{c}\right)\\ &\times\left\{\left[1-\cos\left(\frac{2\pi fL}{c}\right)\right]S_{\rm OMS}(f)\right.\\ &\left.+2\left[1-\cos\left(\frac{2\pi fL}{c}\right)\right]^2 S_{\rm TM}(f)\right\},\end{aligned} \quad (6b)$$

$$\begin{aligned}S_{n,AE}(f) &= S_{n,AT}(f) = S_{n,EA}(f)\\ &= S_{n,ET}(f) = S_{n,TA}(f) = S_{n,TE}(f) = 0.\end{aligned} \quad (6c)$$

In the following sections, we refer to Eq. (6) as the template for the noise.

C. Parametrized SGWB model

Stochastic backgrounds in the LISA band can have a large variety of spectral shapes, depending on the physical processes that generate them [76]. For the purposes of this work, we only consider SGWB templates whose energy density can be described by a power law,

$$h^2\Omega(f) = \frac{1}{\rho_c}\frac{{\rm d}\rho_{\rm GW}(f)}{{\rm d}\log f} = A\left(\frac{f}{f_{\rm knee}}\right)^n. \quad (7)$$

Here, $\rho_{\rm GW}(f)$ is the GW energy density at frequency f, h is the dimensionless Hubble parameter [77] $H_0 = h\times 100\,\rm Kms^{-1}Mpc^{-1} \simeq h\times 3.24\times 10^{-18}\,Hz$, $\rho_c = 3H_0^2/(8\pi G h^2)$ is the critical energy density, $f_{\rm knee}$ is a reference frequency, and A and n are the amplitude and spectral slope of the background, respectively. The values of A and n are determined by the physical processes generating the background [18,78]. As an example of an SGWB that can be described by Eq. (7) in the LISA frequency band, we consider here the concrete case of the SGWB produced by cosmic strings loops [23,27]. For the purposes of the present work, we can rewrite Eq. (7) in the cosmic string case as

$$h^2\Omega_{\rm cs} = 0.55\times 10^{-11}\left(\frac{G\mu}{10^{-13}}\right). \quad (8)$$

Here, $G\mu \sim 10^{-6}(\eta/10^{16}\ \rm GeV)^2$ is the string tension, and η is a characteristic energy scale. For the cosmic string background considered here, we assume a flat plateau characterized by $n = 0$. The overall spectral shape of the background is expected to be more complex, being a function also of the string loop size and the string reconnection probability [23,27,79]. While the location and amplitude of the flat part of the spectrum is affected by the aforementioned parameters, the simple description given in Eq. (8) is sufficient for the purposes of this work, as we are interested in assessing our pipeline's response to the introduction of flexible models when trying to infer the presence of a SGWB in the data. The analysis of a more comprehensive set of signal models is left for future work.

Finally, we express the SGWB energy density in terms of the one-sided signal power spectral density in the TDI channels as

$$\boldsymbol{S}_h(f) = \boldsymbol{\mathcal{R}}(f)h^2\Omega(f)\frac{3(H_0/h)^2}{4\pi^2 f^3}, \quad (9)$$

where $\boldsymbol{\mathcal{R}}(f)$ denotes the sky-averaged LISA response function. Here, bold symbols represent matrices in the chosen TDI basis. Derivations of the LISA response to stochastic signals can be found in Refs. [42,45,46]. In this work, we follow the derivation and implementation described in Ref. [42] and use the associated *Mathematica* code [80]. We first numerically evaluate $\boldsymbol{\mathcal{R}}(f)$ on a grid of frequencies logarithmically spaced in the range $[5\times 10^{-6}, 0.1]$ Hz. We then interpolate these data to evaluate the response function at the frequencies of interest. The aforementioned code can compute the response function for different LISA constellation configurations (e.g., for both equal and unequal arm lengths). Therefore, adopting this setup allows our inference pipeline to be easily adapted to different scenarios by switching the grids used to build the interpolants.

From the one-sided PSDs of noise and signal in the chosen TDI basis, it is straightforward to define a signal-to-noise ratio (SNR) of a SGWB as

$$\mathrm{SNR} = \sqrt{T_{\mathrm{obs}} \sum_{i,j} \int_0^\infty \mathrm{d}f \frac{S_{h,ij}^2}{S_{n,ij}^2}}, \tag{10}$$

with T_{obs} the LISA observation time [81]. This is not strictly equivalent to the definition of signal-to-noise ratio for deterministic sources, as it does not correspond to the output of a search statistic. However, it is commonly used to compare the detectability of different SGWBs with LISA.

In the remainder of this paper, we refer to Eq. (9) as the template for the signals.

D. Data generation

We test our pipeline on two different setups to study its performance under different conditions. In the first case, we generate data consistent with a parametrized stochastic background from cosmic strings just outside the reach of square kilometre array [82], choosing $G\mu = 10^{-13}$ in Eq. (8). We inject the signal in a realization of LISA instrumental noise generated according to the parametrized LISA PSD introduced in Eq. (6). Conversely, in the second case, we analyze a noise realization that is produced from a PSD whose spectral shape differs from what we expect from the instrument. This is described in Sec. IV B.

For all the analyzed datasets, we generate the noise-orthogonal *AET* TDI variables directly in the frequency domain under the assumptions of stationarity and Gaussianity. LISA instrumental noise will neither be stationary nor Gaussian (see, e.g., Ref. [83]), but a more realistic description of it is out of the scope of this work. We generate data according to the models described in the previous sections, working with second-generation TDI variables and assuming constant, equal arm lengths $L = 2.5 \times 10^9$ m. This assumption, despite being highly simplifying, allows us to streamline the description and generation of the simulated datasets and, therefore, to focus this work on presenting the details of our inference methodology.

We generate data assuming an observation time of $T_{\mathrm{obs}} = 1$ yr and a sampling rate of 2 seconds, or equivalently a sampling frequency $f_s = 0.5$ Hz. We then restrict the analyzed frequency range in the interval $f \in [10^{-4}, 2.9 \times 10^{-2}]$ Hz. Although this narrower range limits our constraining power, we can avoid artifacts due to the zeros introduced in the LISA noise transfer function [84]. For these parameters, the aforementioned SGWB has SNR equal to 262.

III. METHODOLOGY

A. Power spectral density breakdown

In this work, we are interested only in stochastic contributions to the overall LISA data streams. Following [45,33], for a single data stream, we can separate its deterministic and stochastic signal components and express it as

$$\begin{aligned}\tilde{d}(f) &= \sum_{i,j} \tilde{h}_{ij}(f) + \sum_i \tilde{y}_i^{\mathrm{GW}}(f) + \tilde{n}(f) \\ &= \tilde{H}(f) + \tilde{y}(f),\end{aligned} \tag{11}$$

where $\tilde{d}(f)$ represents the frequency domain data, $\tilde{h}_{ij}(f)$ the resolvable sources, $\tilde{y}_i^{\mathrm{GW}}(f)$ the stochastic contributions due to gravitational waves, and $\tilde{n}(f)$ the instrumental noise. We use the index i to refer to the different source types (MBHBs, EMRIs, GWDs, SBHBs) and the index j to label the individually resolved sources belonging to the same type. In this work, we do not include any noise artifacts (gaps, glitches), and we assume the perfect subtraction of all the resolvable sources $\tilde{H}(f) = \sum_{i,j} \tilde{h}_{ij}(f)$, so that we can consider only the total stochastic contribution $\tilde{y}(f) = \sum_i \tilde{y}_i^{\mathrm{GW}}(f) + \tilde{n}(f)$.

We perform our analysis using multiple channels. Thus, Eq. (11) applies to each of them. We can write the data vector of the Fourier transformed TDI variables as $\tilde{\mathbf{d}} \equiv (\tilde{A}, \tilde{E}, \tilde{T})^T$, and then, at each frequency, we can define the 3×3 periodogram matrix $\mathbf{P}(f) \equiv \tilde{\mathbf{d}}(f)\tilde{\mathbf{d}}(f)^\dagger$. Daggers represent conjugate transposes. Following [45], we include in the Fourier variables the normalization factor $\sqrt{2/(Nf_s)}$, with $N = T_{\mathrm{obs}} f_s$ the length of the time series. In this way, the periodogram is directly comparable to the one-sided PSD.

Since signal and noise are independent processes, the data covariance matrix $\mathbf{C}_d(f) = \langle \tilde{\mathbf{d}}\tilde{\mathbf{d}}^\dagger \rangle(f)$ can be rewritten as the sum of the instrumental noise and SGWBs covariances. On the TDI level, this reads

$$\mathbf{C}_{AET}(f) = \frac{1}{2} \begin{bmatrix} S_{AA}(f) & 0 & 0 \\ 0 & S_{EE}(f) & 0 \\ 0 & 0 & S_{TT}(f) \end{bmatrix}, \tag{12}$$

with

$$S_{ii}(f) = S_{n,ii}(f) + S_{h,ii}(f), \qquad i \in \{A, E, T\}. \tag{13}$$

Here, $S_{n,ii}(f)$ and $S_{h,ii}(f)$ represent, respectively, the noise and background contribution to the overall one-sided PSD in each channel. We account for the nonperfect knowledge of their spectral shape, factoring the overall PSD as the product of a parametrized template $S_{n/h}^{\mathrm{template}}(f|\boldsymbol{\lambda}_{n/h})$ and a perturbation $\mathcal{P}_{n/h}(f|\boldsymbol{\gamma}_{n/h})$. We perform our inference algorithm on the parameters of both the templates ($\boldsymbol{\lambda} = \{\boldsymbol{\lambda}_n, \boldsymbol{\lambda}_h\}$) and the perturbations ($\boldsymbol{\gamma} = \{\boldsymbol{\gamma}_n, \boldsymbol{\gamma}_h\}$). We represent the entire parameter space with $\boldsymbol{\theta}$: $\boldsymbol{\theta} = \{\boldsymbol{\theta}_n, \boldsymbol{\theta}_h\} = \{\boldsymbol{\lambda}, \boldsymbol{\gamma}\} = \{\boldsymbol{\lambda}_n, \boldsymbol{\lambda}_h, \boldsymbol{\gamma}_{n,ii}, \boldsymbol{\gamma}_h\}$. Therefore, each component reads

$$S_{n,ii}(f|\boldsymbol{\theta}_n) = S^{\text{template}}_{n,ii}(f|\boldsymbol{\lambda}_n) \times 10^{\mathcal{P}_{n,ii}(f|\boldsymbol{\gamma}_{n,ii})}, \tag{14a}$$

$$S_{h,ii}(f|\boldsymbol{\theta}_h) = S^{\text{template}}_{h,ii}(f|\boldsymbol{\lambda}_h) \times 10^{\mathcal{P}_h(f|\boldsymbol{\gamma}_h)}. \tag{14b}$$

This factorization results in two correlated terms, but we choose it to allow the templates to catch the overall frequency evolution of the PSD and let eventual deviations from it be absorbed into $\mathcal{P}_{n/h}(f)$. In this work, we assume three independent perturbations for the noise contributions across the three TDI channels, whereas the background spectral shape is perturbed in a consistent (i.e., shared) way. Hence, the presence of the indices ii only in the exponent of Eq. (14). This assumption, which translates into using four different perturbations, is based on the hypothesis that the transfer function of the signal will be well known. While the same holds for any noise source entering the measurement chain, we cannot know *a priori* the full set of noises that will be present in the data stream. In addition, we will not have access to all the physical parameters that constitute the noises we expect to be there (see, e.g., Refs. [1,85]). Thus, we decided to be fully generic only in the noise contribution to the overall PSD in each channel. This approach has a caveat: the *AET* basis is noise orthogonal only under the assumption of a perfect equilateral constellation where all the secondary noise sources are exactly equal in the three spacecrafts [86], thus using independent splines in the three channels may seem a violation of this assumption. For a complete analysis, we should also include perturbations in the off-diagonal components of the noise matrix, which would increase the number of perturbations for the noise contributions by another three terms. The off-diagonal terms would also account for cross-correlations between the noise and SGWB components at high frequencies [38,42]. Cross-correlations can be as large as $\sim$10% of the diagonal terms [38]. These can have an impact on the distinguishability of instrumental noise and SGWBs, but from the results of Ref. [42] we do not expect our conclusions to be heavily affected. For completeness, this should be validated against a full analysis in the Michelson TDI basis, where the (Hermitian, positive-definite) covariance matrix $\mathbf{C}_{XYZ}(f)$ is dense with six independent components, since it is possible to construct a positive definite matrix whose elements are functions of six independent, flexible terms [42]. This requests the evaluation of six splines at every likelihood call only for the noise contribution, significantly increasing the computational cost. Here, we neglect off-diagonal terms for computational simplicity. Having generated the frequency domain data directly in the *AET* basis, we know the noise-orthogonal assumption holds. We leave a full covariance treatment of the three noise channels for future analysis.

We use as perturbations $\mathcal{P}_{n/h}(f)$ realizations of Akima splines [87], which although only being C^1 differentiable, were introduced to produce smoother fits to data in the presence of rapidly varying second derivatives with respect to other widely used C^2 Cubic splines implementations [88–90]. A description of the Akima splines algorithm can be found in the Appendix. To be as generic as possible, we do not fix the number or the position of the spline knots beforehand. Instead, we use a reversible jump Markov chain Monte Carlo (RJ-MCMC) scheme [91] to dynamically infer the most probable number of knots. We use the parallel tempered ensemble sampler Eryn [92–94] for this. Our knot placement scheme can be described as follows:

(i) We fix the frequency position of the first and last spline knots to the extrema of the frequency range considered. We label these two points as edges, and we only fit for their amplitudes. This is crucial to always make the splines cover the whole frequency range.

(ii) For the in-between knots, we sample both the position f_k and the amplitude w_k. The sampler proposes a change in their number at each step. Following Eryn's jargon, this is done through a *between-models* move, proposing the addition of a new knot to the active ones (i.e., the ones currently used to evaluate the splines) or removing one. When the addition of a new knot is proposed, its parameters are drawn from the prior. We refer the reader to Refs. [33,45,93,95] for additional details about the Eryn reversible jump implementation and its recent applications in LISA data analysis.

We set the minimum and maximum number of in-between knots to 0 and 10, respectively. Since our Akima splines implementation uses boundary conditions that require at least four points to be computed, we rely on a linear interpolant for cases with less than two in-between knots.

With the nomenclature introduced in this section and the templates presented in Eqs. (6) and (9), we can summarize the parameters of our model as follows:

(i) $\boldsymbol{\lambda}_n = \{A_{\text{TM}}, A_{\text{OMS}}\}$: the two noise amplitudes

(ii) $\boldsymbol{\lambda}_h = \{A, n\}$: the amplitude and the spectral index of the signal power law

(iii) $\boldsymbol{\gamma}_\alpha = \{w_0^\alpha, \{f_k, w_k\}^\alpha_{k=1,\dots,n-1}, w_n^\alpha\}$: the amplitudes at the edges and the position and amplitude of each in-between knot added through reversible jump, for each independent spline component α

Given the properties presented in this section, it is clear that our inference model includes components without a "physical" origin (the splines) and does not leverage the different noise and signal TDI transfer functions to discriminate between the two contributions. Although these may sound like disadvantages, we believe they represent the core features of our work. In fact, the generality we introduced enables analyses in which the totality of the stochastic contributions are fitted as a whole in a global analysis that also includes deterministic sources. The total stochastic power found in this way could be broken down into

individual contributions in a postprocessing step. We leave the assessment of the feasibility of this approach to future investigations.

B. Hardware acceleration

In our setup, we must evaluate many potentially different splines at each MCMC step. All the walkers of our ensemble can have different numbers of spline knots with different positions and amplitudes (f_k, w_k). Therefore, we implement a version of Akima splines [96], which leverages parallelization and GPU (graphics processing units) acceleration to speed up their evaluation substantially. While the Cupy library [97] already offers a GPU implementation of Akima splines, this is mostly suitable for scenarios where the number and the position of the interpolation knots are fixed. We tailor our code to handle the variability caused by the reversible jump setup.

Additionally, we employ the JAX library [98] extensively in our code base, leveraging its built-in GPU support, vectorization, and just-in-time compilation.

C. MCMC setup

We run Eryn using 20 different temperatures with 20 walkers each. We initialize our walkers, drawing their position in the parameter space randomly from the prior. The only exception is for the spline knots, whose initial amplitudes are drawn from a multivariate normal centred on zero $\mathcal{N}(\boldsymbol{\mu} = \mathbf{0}, \boldsymbol{\Sigma} = \mathbf{I} \times 10^{-8})$. This choice follows our interpretation of the splines as "small" perturbations of parametrized templates. For each of our parameter estimation (PE) runs, we monitor the convergence of the samples by making sure that the chains are longer than $200 \times \max(\hat{\tau})$, with $\hat{\tau}$ being the average autocorrelation time across all the model parameters [99]. We discard (at least) the first $5 \times \max(\hat{\tau})$ samples to account for the burn-in phase, and we thin the chains by a factor of $0.5 \times \min(\hat{\tau})$.

The large number of parallel temperatures helps to increase the acceptance rate of our runs and allows us to estimate the evidence Z through the stepping-stone method [100]. In the standard Bayesian framework, the evidence, or marginal likelihood, is a normalization constant equal to the integral over the parameter space of the unnormalized posterior distribution,

$$p(\boldsymbol{\theta}|d) = \frac{\mathcal{L}(d|\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{Z}, \qquad Z = \int \mathrm{d}\boldsymbol{\theta}\mathcal{L}(d|\boldsymbol{\theta})\pi(\boldsymbol{\theta}), \tag{15}$$

where d represents the observed data, $\boldsymbol{\theta}$ the model parameters, $\pi(\boldsymbol{\theta})$ their prior, and $\mathcal{L}$ the assumed likelihood. We report our prior distributions in Table I. The evidence is usually expensive to compute, and being a normalization constant, it can safely be disregarded in parameter estimation studies. However, it is required to perform Bayesian model selection. The ratio of the evidence computed for two competing models, M_1 and M_2, is called the Bayes factor, $\mathcal{B}_{12} = Z_1/Z_2$. A large Bayes factor indicates that model M_1 describes the observed data better than M_2. Examples of the use of Bayes factors for model selection in the field of GW astronomy can be found in Refs. [101–103], while applications of the stepping-stone method in this field can be found in Refs. [104,105]. Following [100], the stepping-stone approximation for the evidence can be expressed as

TABLE I. Prior distributions used for the MCMC runs presented in this work. A and n represent the amplitude and slope of the power-law in Eq. (7). We use uniform distributions for most of the sampling parameters. The prior distribution assumed for the weight of each spline knot and edge is a normal distribution centered on $\mu = 0$. In each channel and for each contribution, we assume the same priors on f_k and w_k.

Parameter	Prior distribution
$A_{\mathrm{TM}}/10^{-15}[\mathrm{ms}^{-2}\ \mathrm{Hz}^{-1/2}]$	$U[10^{-2}, 10^{2}]$
$A_{\mathrm{OMS}}/10^{-12}[\mathrm{m}\ \mathrm{Hz}^{-1/2}]$	$U[10^{-2}, 10^{2}]$
$\log_{10} A$	$U[-20, -10]$
n	$U[-5, 5]$
$\log_{10} f_k$	$U[\log_{10} f_{\min}, \log_{10} f_{\max}]$
w_k	$\mathcal{N}(\mu = 0, \sigma = 0.1)$

$$\log \hat{Z} = \sum_{k=1}^{K} \log \sum_{i=1}^{n} \mathcal{L}(d|\boldsymbol{\theta}^{i}_{\beta_{k-1}}, M)^{\beta_k - \beta_{k-1}} - K \log n, \tag{16}$$

where K is the number of temperature T_k in the ladder, $\beta_k = 1/T_k$ are the inverse temperatures, M is the model the evidence is computed for, and $\boldsymbol{\theta}^{i}_{\beta_{k-1}}$ are n draws from the annealed posterior at temperature T_{k-1}. For low β values, the likelihood is suppressed by the tempering. When β tends to 1, the tempered posterior approaches the target distribution. To cover these behaviors and to resolve the turnover point between the two, we place five chains with β logarithmically equispaced between 10^{-20} and 10^{-12}, 10 between 10^{-12} and 10^{-5}, and 5 between 10^{-5} and 1.

To further reduce the computational cost of our runs, we do not sample using the totality of the frequency domain data and the Whittle likelihood [106]. This represents the standard in the context of GW astronomy, and it has already been widely implemented in the literature (see, e.g., Ref. [46], and Ref. [107] for a more comprehensive SGWB likelihoods review). Here, we average the periodograms over N_{bins} consecutive and nonoverlapping frequency bins, and we sample using a complex Wishart likelihood, as described in Ref. [45]. We define the averaged periodogram in a segment i as

$$\mathbf{P}(f_i^{\mathrm{bin}}) \equiv \frac{1}{n_i} \sum_{k=i-\frac{n_i}{2}}^{i+\frac{n_i}{2}} \tilde{\mathbf{d}}(f_k)\tilde{\mathbf{d}}(f_k)^{\dagger}, \tag{17}$$

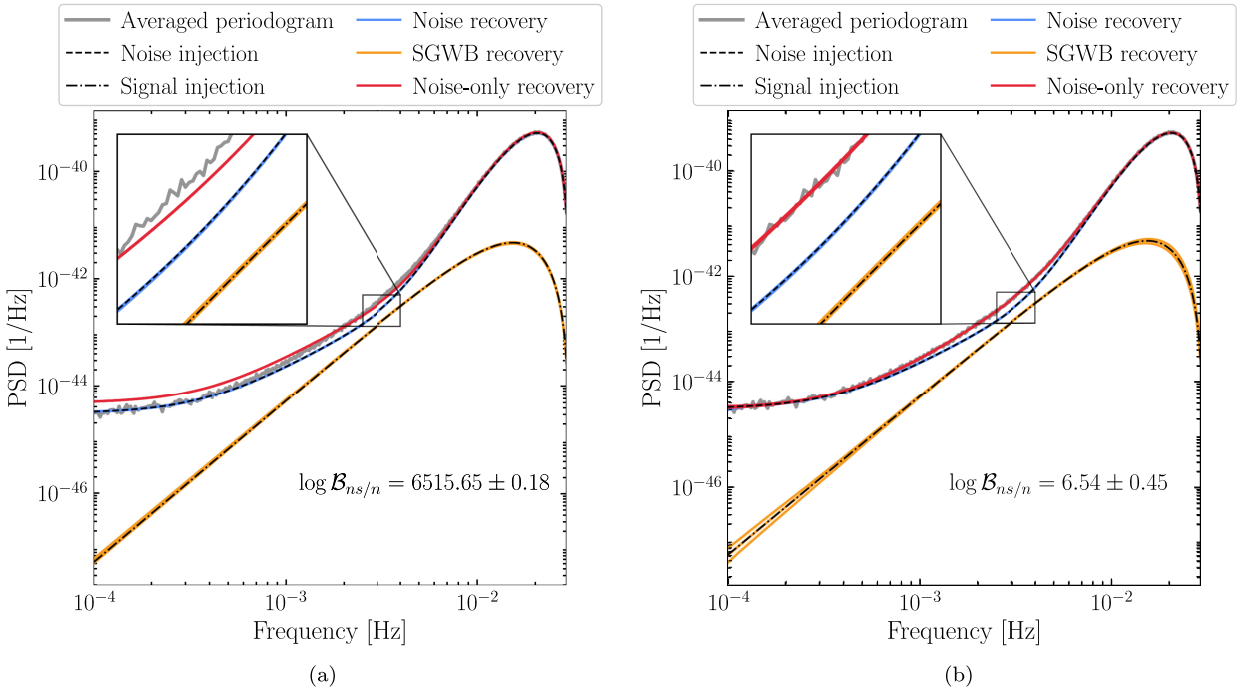

FIG. 1. Reconstructed posterior distribution of the different contributions to the total PSD in the A channel. The black dashed (dash-dotted) line represents the PSD used to generate the noise (SGWB from cosmic strings) dataset. The averaged data periodogram is shown as a gray line. Blue (orange) lines and filled contours represent medians and 90% credible regions for the reconstructed noise (signal) PSD when both contributions are present in the model. Red lines and contours refer to the case where only the noise has been used to fit the totality of the data. Insets show a zoom in on the region $f \in [2.5, 4]$ mHz. Left: result of the analysis performed using only the baseline templates. Right: result of the analysis performed introducing splines in the model.

where n_i is the number of frequency samples in a bin centred on $f_i^{\rm bin}$. Introducing the number of degrees of freedom (d.o.f.'s) in each bin $\nu(f_i^{\rm bin}) = n_i$, the matrix $\mathbf{Y}(f) = \nu(f)\mathbf{P}(f)$ would follow the aforementioned complex Wishart distribution if $\tilde{\mathbf{d}}(f_k)$ were uncorrelated between frequency bins. When working with time domain data, we apply a window function to the time series before converting it into the frequency domain. This would introduce correlations between frequency bins, effectively reducing the number of d.o.f.'s $\nu(f)$. Since we generate our data directly in the frequency domain, this step is not necessary. Keeping only the terms that depend on the model parameters, for each bin we can write

$$\log p(\mathbf{Y}(f)|\boldsymbol{\theta}) = -\mathrm{tr}(\mathbf{C}_d^{-1}\mathbf{Y}(f)) - \nu(f)\log|\mathbf{C}_d(f)|, \quad (18)$$

where $\mathrm{tr}(\cdot)$ denotes the trace operator, and $|\cdot|$ represents the determinant of a matrix. Combining all the bins together, our likelihood reads out,

$$\log \mathcal{L}(\mathbf{Y}|\boldsymbol{\theta}) = \sum_{i=0}^{N_{\rm bins}-1} \log p(\mathbf{Y}(f_i^{\rm bin})|\boldsymbol{\theta}). \quad (19)$$

In the limit of $n_i = 1$ for all i, having only one frequency sample per bin, the complex Wishart likelihood converges to the Whittle likelihood [45]. For completeness, we include in our code the possibility of using directly the latter with the totality of the frequency domain data. We have checked that the two methods yield consistent results.

D. Two-step inference

Our last acceleration step consists of splitting the inference into two consecutive steps. There is a degeneracy in the model in that changes to the model template parameters can be mimicked by changes to the spline parameters. We address this degeneracy by first running an initial MCMC with only the modeled noise parameters and the background templates. This corresponds to enforcing all the spline knots to have $w_k = 0$ in Eq. (14). The resulting posterior distributions are used to produce the results labeled as "template recovery" in Figs. 1 and 4 and to fix the template parameters in a second MCMC run. In the second run, we let the splines free to explore the parameter space, while the template parameters $\boldsymbol{\lambda}$ are fixed to the medians of the aforementioned posteriors. If compared to running one single MCMC where all the parameters $\boldsymbol{\theta}$ are updated together, this separation dramatically improves the convergence of the chains, as we remove the degeneracy between the template parameters and the splines. We can do this because in our model the splines are not used to describe the overall spectral shape of the

PSD, but only to account for (small) deviations from the templates. We can interpret the recovered template model parameters as the best fit model and the spline parameters as deviations from that best fit. This applies to both the noise and the signal analyzed. In general, if the spline model was used to describe the totality of the stochastic contribution as one individual component, the priors would be increased to be as large as needed. Instead, we interpret the splines as deviations from the best-fit model; therefore, we use tight zero-mean Gaussians as priors on w_k.

We have checked that when fixing the template parameters to different values sampled from the posterior inferred in step 1, the splines parameters change accordingly, resulting in consistent recovered PSDs.

This two-step approach, combined with the hardware acceleration and averaging procedure described in Secs. III B and III C, allows our runs to converge and accumulate enough samples within a day on an NVIDIA A100-SXM4 GPU despite the large number of parallel walkers and the dimensionality of the parameter space.

IV. RESULTS

In this section, we present the results of our pipeline on the different datasets described in Sec. II. We compare the results obtained using only the templates with those obtained using splines to account for deviations from them. The analysis in Sec. IV B of a data stream generated from a different PSD with respect to the template used in the recovery allows us to test the robustness of our pipeline in scenarios where the analyzed data do not match the baseline parametrized templates. We leave the analysis of SGWBs whose spectral shape differs from our baseline power-law template for future work.

A. Cosmic strings

First, we present results for the cosmic string-like SGWB. Our reconstructed posterior distributions on the PSDs used for the data generation are shown in Fig. 1, together with the respective injections. As anticipated in Sec. III D, we report in the two panels the results obtained with only the templates (left) and with the inclusion of the splines (right). In each panel, we report the results obtained under two separate hypotheses: first, we include both the noise and signal components in the model fitted to the data, shown as the blue and orange curves in the plot. Next, we perform the inference using only the noise component, shown in red. The two models read, respectively, like the totality and the first term of Eq. (13).

While this is not the case for the analysis performed only with templates, when we introduce splines in the model, both approaches can successfully account for the totality of the data. Figure 1(a) clearly shows that the red curve, representing the recovery performed using only the noise template in Eq. (6), fails to model the totality of the data. This is particularly evident at low frequencies, where the reconstruction significantly overestimates the data content. The main driver for this is the inclusion of the T channel in the analysis. T is mostly insensitive to GWs (especially at low frequency) [108], so it can, in principle, be used as a noise monitor. This holds especially in the idealized equal arm length configuration assumed here. Its noise PSD is dominated by the OMS component, which controls the A and E PSD only at high frequency. Since T is largely unaffected by the presence of a signal, when the latter is present in the data but not included in the model, $A_{\rm OMS}$ is better constrained than $A_{\rm TM}$. While the median of the $A_{\rm TM}$ posterior is more than 167 standard deviations (σ) away from the corresponding injected value, the bias for $A_{\rm OMS}$ is of "only" 8σ. The poor estimate of $A_{\rm TM}$ dominates the error at low frequency.

The situation is different for the splines analysis in Fig. 1(b). Under the first hypothesis, the reconstructed 90% credible regions, represented as filled contours, entirely encompass the injected PSDs. In the second, the reconstruction perfectly overlays with the averaged data periodogram shown in gray. In both panels of Fig. 1, the 90% credible regions for the noise recovery under both assumptions are tight enough to be mostly indistinguishable from the median of the posterior distributions. We can compute the Bayes factor between the two hypotheses to assess which is better supported by the data. This is a pivotal question posed by resorting to flexible methods, as lifting the assumption of knowing the spectral shape of the different contributions inevitably reduces our constraining power. Flexible noise models can accommodate the excess power present in the data because of a stochastic signal, making it harder to detect an SGWB and constrain its parameters. Labeling the first hypothesis as "noise+signal" (ns) and the second as "noise" (n), we compute the respective evidences using Eq. (16) and the natural logarithm of the Bayes factor $\log \mathcal{B}_{ns/n} = \log Z_{ns} - \log Z_n$. We find a log-Bayes factor of $\log \mathcal{B}_{ns/n} = 6.54 \pm 0.44$, which, according to the criterion outlined in Ref. [109], translates to "very strong" evidence for the presence of a stochastic signal in the data. Thus, in this case, our flexible parametrization allows us to confidently discriminate between the presence or absence of an SGWB. Although the splines can absorb the SGWB, they do so by varying considerably more parameters, so the more parsimonious description of an SGWB is still preferred. Nevertheless, we want to stress that the introduction of the splines reduced the log-Bayes factor by three orders of magnitude, as shown in Fig. 1. Therefore, the increased flexibility introduced in the noise (and signal) model(s) could make detecting SGWBs with SNR lower than the one considered much more challenging, as also shown in Ref. [45]. This is particularly relevant for SGWBs that are quieter than our example, but that would still be detectable, in principle, in a templates-only analyses.

The interpretation of the absolute value of the Bayes factor in the context of SGWB analyses is an active area of research with the authors of Refs. [43–46]. Moreover, it is always influenced by the prior distributions assumed during the analysis. Nevertheless, here it still provides a valuable metric to assess the impact of model-agnostic analyses on the detectability of these signals. We also refer the reader to Ref. [110], where the impact of the inclusion of model uncertainties in the Bayes factor computation was investigated. For comparison, we also apply the deviance information criterion (DIC) [111,112], an alternative model selection tool that does not require the computation of the evidence. We define the deviance as $D(\boldsymbol{\theta}) = -2\log\mathcal{L}(d|\boldsymbol{\theta}) + C$, where C is a constant that cancels out in all model comparisons. Using the notation $\bar{x}$ to indicate the mean over the posterior distribution of the quantity x, we can define the effective number of model parameters as $p_D = \overline{D(\boldsymbol{\theta})} - D(\bar{\boldsymbol{\theta}})$. From these quantities, the DIC is computed as

$$\mathrm{DIC} = \overline{D(\boldsymbol{\theta})} + p_D. \tag{20}$$

Models with smaller DIC values should be preferred. In particular, Ref. [109] suggests that the absolute value of the difference in DIC between two models, $|\Delta\mathrm{DIC}|$, should roughly follow the same scale of twice the natural logarithm of the Bayes. For the template analysis, we find $\Delta\mathrm{DIC} = \mathrm{DIC}_{ns} - \mathrm{DIC}_{n} = -1.3\times10^{4}$, while in the splines analysis we obtain $\Delta\mathrm{DIC} = -30$. This agrees with our Bayes factor results, favouring the "noise+signal" hypothesis in both cases. Once again, the inclusion of splines significantly reduces the strength of the support for the presence of a signal. However, when analyzing the "noise +signal" data, we find DIC_{ns} to be (slightly) smaller in the templates-only analysis, while the splines case shows a (slightly) larger evidence than the templates-only case. Therefore, under our inference setup, the DIC seems (slightly) more penalizing toward increased model complexity than the Bayes factor.

Given the stochastic nature of the datasets, we would expect that flexible models are capable of achieving larger likelihoods than the analytic templates, even if the latter ones match the spectral shape used during the data generation. This is a consequence of the flexible models being able to better accommodate the random fluctuations inherent to stochastic processes. Both the noise and signal datasets simulated here are random realizations whose PSD is described by Eqs. (6) and (9) only on average. Therefore, we expect the presence of local fluctuations that deviate from the templates. Given the local nature of the splines, they can account for these fluctuations, while the templates cannot. In principle, without an upper bound on the number of knots, one could reach larger and larger likelihood values by adding knots until the recovered PSD exactly matches the data periodogram everywhere. This effect is balanced by the increased prior volume associated with the additional parameters. The fact that only a small number of interior knots are added when the data are generated from a PSD consistent with the template family gives us confidence that our priors on the spline parameters have been chosen well.

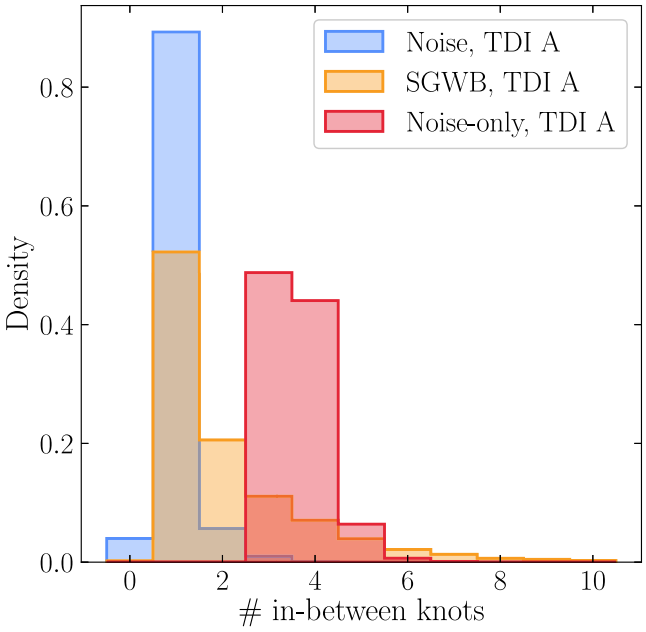

FIG. 2. Posteriors on the number of in-between spline knots for the data stream containing noise and an SGWB due to cosmic strings. Blue (orange) bars refer to the noise in the A channel (signal) contribution. Red bars refer to the analysis performed only with the noise component.

We report the posterior on the number of in-between spline knots in the A channel in Fig. 2. The posteriors in A and E always look the same, since the PSD in the two channels has the same shape. Conversely, given its low sensitivity to GWs, the T channel posteriors always converge to one single knot. Colors reflect the labels already described for Fig. 1. Under the "noise+signal" hypothesis, all the spline perturbations converge to having only one in-between knot. On the other hand, given the loud nature of the signal, the recovery under the noise-only hypothesis requires between a total of five and six knots to account for the totality of the power present in the data stream. This behavior reflects our expectations. Since the simulated data are consistent with the templates used during the inference stage, we do not have to rely on the spline component to catch deviations from the parametrized baseline. Nevertheless, since one extra knot is still added by the reversible jump algorithm, we believe that its role is to account for the aforementioned random fluctuations. The different posteriors help interpret the Bayes factors previously presented. In fact, even if our flexible

model under the "noise-only" hypothesis can provide a good description of the data, it requires a larger number of parameters than the model including both noise and signal. Usually, we would like our analyses to implement the concept of "Occam's razor", i.e., if two models describe the data equally well, the more complex one is disfavored. It is often said that the Bayes always automatically includes it, but this holds only in cases where suitable priors are chosen (see, e.g., Ref [113]). The Bayes factors found here demonstrate that our priors are chosen appropriately, since we prefer the simplest description of the data in this case, which is the noise model template plus an SGWB.

Figure 1 shows the 90% credible regions getting wider when splines are included in the model. The effect is more pronounced for the SGWB contribution, while it is barely visible for the remaining components. To better highlight this behavior, we show in Fig. 3 the posterior distribution of the reconstructed total PSD in the three TDI channels at a reference frequency of $\bar{f} = 10^{-3}$ Hz. We note that the color scheme in this figure is different to the one used in Figs. 1 and 2. This figure allows us to compare how the location and width of the PSD posteriors change with the assumed model. The tightest contours are obtained by the two template only analysis. The noise-only recovery in purple is clearly biased in all three channels. Including the splines in the model has the effect of broadening the posteriors for both the "noise+signal" and the "noise" recoveries and to shift the latter enough to be in agreement with the injection. Both here and in Fig. 1(b), the credible intervals for the splines analysis account only for the uncertainty in the spline parameters, while the template parameters remain fixed. The same applies to the evidence and DIC calculations in the splines analyses. While being a caveat, the posteriors in Fig. 3 suggest that the uncertainty in the spline coefficients translates in larger contours than the uncertainty in the template parameters alone. Therefore, we do not expect our conclusions to change significantly.

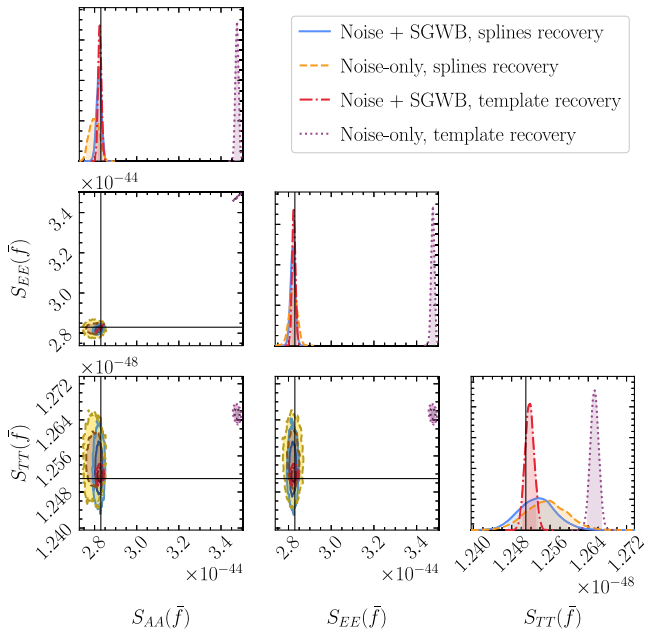

FIG. 3. Posterior distribution of reconstructed total PSD in the three TDI channels at a reference frequency of $\bar{f} = 10^{-3}$ Hz, for the data stream containing noise and an SGWB due to cosmic strings. Black solid lines represent the values evaluated at the injected parameters. 2D contours represent the $1-\sigma$ and $2-\sigma$ regions, while the 1D shaded areas represent 95% credible regions.

B. Perturbed noise spectrum

Finally, we present our results for the "perturbed" noise PSD. In this case, we have generated the data according to a PSD of the same shape as Eq. (14) with $\mathcal{P}_{n,ii}(f)$ being splines with knot positions $\log_{10} f_k = \{-7.5, -4, -2.01, -1.71, -1.54, -0.6\}$ and amplitudes $w_k = \{0.0, 0.2, -0.1, 0.3, 0.2, 0.0\}$. We assume the same knot positions and amplitudes for all the channels, and we include the leftmost and rightmost points to cover the whole frequency range implied by our assumed observation period and cadence. The base, unperturbed noise PSD and the one used for the injection are shown as the black dash-dotted and dashed lines in Fig. 4. Red and blue lines and

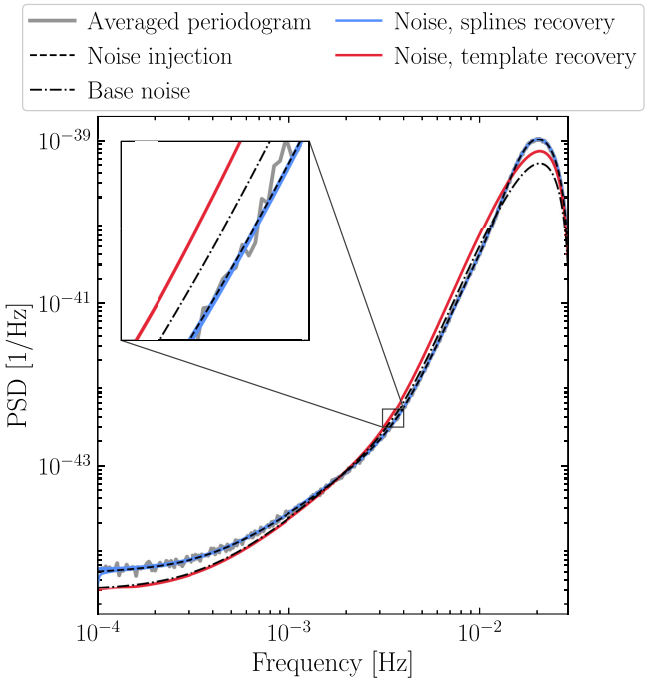

FIG. 4. Reconstructed posterior distribution of the reconstruction of the total PSD in the A channel. The black dashed (dash-dotted) line represents the "perturbed" PSD used to generate the noise dataset (the base noise PSD before applying the spline perturbation). The averaged data periodogram is shown as a gray line. Blue lines and filled contours represent the median and 90% credible interval for the reconstructed noise PSD. The inset shows a zoom in on the region $f \in [3.1, 4]$ mHz.

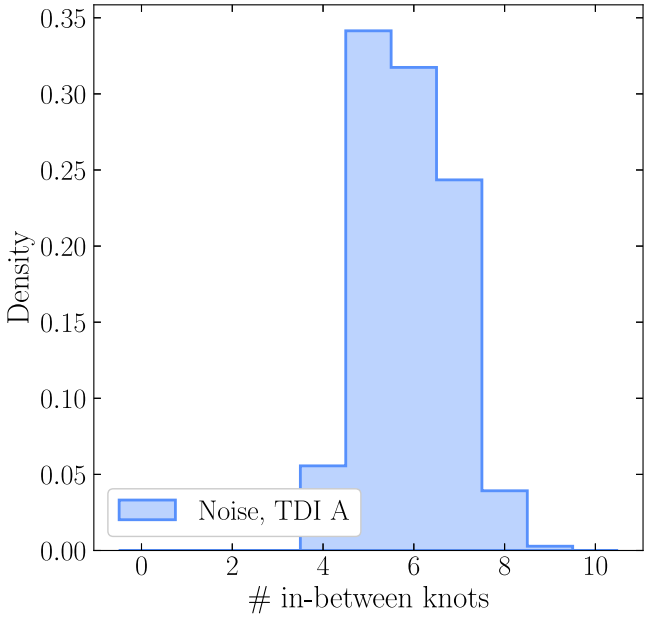

FIG. 5. Posteriors on the number of in-between spline knots for the data stream containing only noise generated from the "perturbed" PSD.

contours represent the median and 90% credible interval for the reconstructed PSD for the template and spline recoveries, respectively. While the first approach fails to successfully represent the power present in the data, the spline analysis is able to recover the injected PSD despite it being different from the assumed baseline. This gives us confidence that the results of the previous section do not strongly depend on having the right template as the baseline of the flexible models, as long as the differences in the spectral shape are small enough to be compatible with our priors.

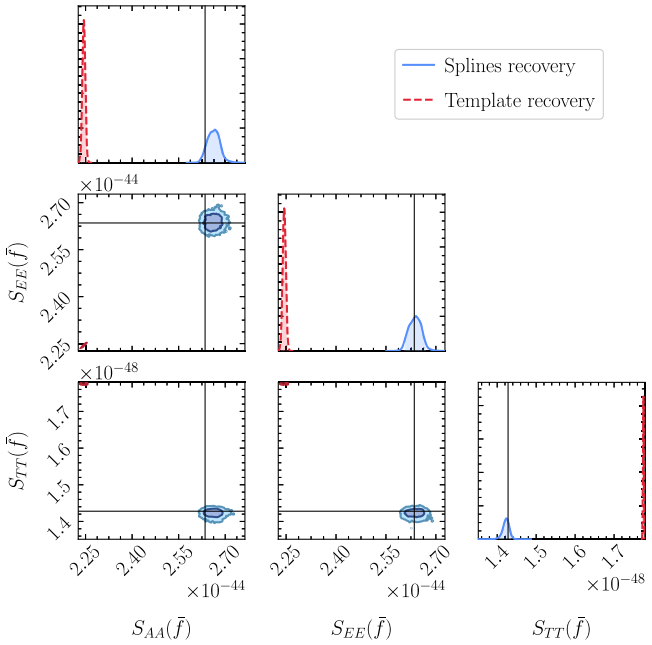

FIG. 6. Posterior distribution of reconstructed total PSD in the three TDI channels at a reference frequency of $\bar{f} = 10^{-3}$ Hz, for the data stream containing only "perturbed" noise. Black solid lines represent the values evaluated at the injected parameters. 2D contours represent the $1-\sigma$ and $2-\sigma$ regions, while the 1D shaded areas represent 95% credible regions.

We report the posterior on the number of spline knots required to account for the totality of the data in Fig. 5. In this case, we need a larger number of knots to model the simulated data, a result that is consistent with our expectations. Notably, the posterior does not converge to the number of knots used when generating the data, which was four. This is a consequence of the intrinsic correlation between the baseline templates and the spline components of the model. Therefore, the recovered parameters of the two components are free to be different from those used during the data generation as long as the total PSD remains the same.

Finally, we show in Fig. 6 the posterior distribution of the reconstructed total PSD in the three TDI channels at the reference frequency $\bar{f} = 10^{-3}$ Hz. Once again, the inclusion of splines in the model has the effect of broadening the posterior distributions, and shifting them enough to encompass the injected values.

V. CONCLUSIONS

In this paper, we have presented a novel code designed to simultaneously extract the properties of the instrumental noise and stochastic backgrounds from LISA data under model-agnostic assumptions. For both components, we provide flexible models that do not necessarily have to follow rigid spectral shapes. The LPF mission has already shown that analytical models could not account for the totality of the observed instrumental noise [40,41], and we expect the same to hold for LISA. Conversely, we ultimately want SGWB models that are ready to capture unexpected features in the signal spectral shape. Hence, the requirement of flexible models for both stochastic contributions, which have to be able to account for the unmodeled contribution that will enter our measurement chain.

Our fit is based on the inclusion of Akima splines. In each TDI channel, we apply them to the state-of-the-art noise and signal models [1] and run our inference scheme on the parameters of both. We introduce the splines directly at the TDI level, after the LISA response and transfer functions, to account for any arbitrary contribution. This level of flexibility comes at a price, as it reduces our capability of discriminating stochastic signals from pure noise. We investigate the behavior of our pipeline in two different scenarios.

We first test our setup against a simulated dataset containing a realization of instrumental noise and a stochastic background compatible with cosmic strings. In this case, both the noise and the signal injections fall inside

the parameter space of our baseline templates alone. Relying on a detection criterion based on the Bayes factor, we find that our inference scheme can successfully disentangle the signal and the noise, reconstructing both. We find a log-Bayes factor of $\log \mathcal{B}_{ns/n} = 6.54 \pm 0.44$ in favor of the presence of a signal in the data. Nevertheless, the extra flexibility introduced comes at the price of losing constraining power. Running our pipeline with only rigid templates produces a much larger Bayes factor ($\log \mathcal{B}_{ns/n} = 6515.65 \pm 0.18$) and tighter posteriors on the total PSDs with respect to the model with the splines, as shown in Figs. 1 and 3. For the power-law signal considered, having an SNR of 262, we find a decrease in the log-Bayes factor of 3 orders of magnitude. This is a consequence of our flexible noise model's ability to account for the totality of power contained in the data stream, at the price of a larger parameter space dimensionality.

Next, we assess the performance of our model when the analyzed noise data do not match the assumed baseline template. Even in this case, our flexible noise model successfully recovers the spectral shape of the analyzed data, but it needs a large number of spline knots to do so.

While we work under a set of simplifying assumptions, this paper represents a further step forward toward an efficient, general, and agnostic setup able to address the challenges posed by stochastic signals in LISA. From here, natural future research directions include (i) extending our approach to the full covariance matrix, (ii) applying our methodology to datasets of increasing realism, including the future LISA data challenges [70], (iii) analyzing an SGWB whose PSD spectral shape does not match the analytic template assumed as baseline during the inference, and (iv) leveraging our agnostic setup to fit the totality of the data as a whole and discriminate between noise and signals *a posteriori*.

VI. CODE AVAILABILITY

The library used to perform the analysis presented in this paper is available at [114]. This package is still under early development, and it is likely to undergo significant changes in the future. We also release our parallel, GPU-accelerated implementation of Akima splines at [96]. Besides the few packages already cited, this work made extensive use of NumPy [115], Matplotlib [116], Numba [117], and Pysco [118].

ACKNOWLEDGMENTS

We thank Alexandre Toubiana, Lorenzo Speri, and the LISA Cosmology Working Group members for fruitful discussions. In particular, the authors acknowledge the collaboration of Germano Nardini, Chiara Caprini, Jean-Baptiste Bayle, Nikolaos Karnesis, Quentin Baghi, Riccardo Buscicchio, Federico Pozzoli, Mauro Pieroni, and Dam Quang Nam. A. S., M. M., and O. H. gratefully acknowledge the support of the German Space Agency, DLR. The work is supported by the Federal Ministry for Economic Affairs and Climate Action based on a resolution of the German Bundestag (Project Ref. No. FKZ 50 OQ 2301).

DATA AVAILABILITY

The data that support the findings of this article are openly available [119].

APPENDIX: AKIMA SPLINES IMPLEMENTATION

Given a set of knot points with coordinates $(x_i, y_i)_{i=1,\ldots,n}$, Akima splines construct a piecewise function $S(x)$ such that $S(x_i) = y_i$ for all i. In those points, its slope s_i is a function of (x_j, y_j) with $j \in \{i-2, \ldots, i+2\}$. We can define the slope of the line connecting two consecutive knots (x_i, y_i) and (x_{i+1}, y_{i+1}) as

$$m_i = \frac{y_{i+1} - y_i}{x_{i+1} - x_i}, \quad i = 0, \ldots, n-1. \tag{A1}$$

Defining the quantity

$$w_i = |m_{i+1} - m_i|, \tag{A2}$$

we can construct the spline slopes as

$$s_i = \begin{cases} \frac{m_{i+1}+m_i}{2} & \text{if } w_i + w_{i-2} = 0, \\ \frac{w_i m_{i-1} + w_{i-2} m_i}{w_i + w_{i-2}} & \text{otherwise.} \end{cases} \tag{A3}$$

For each knot (x_i, y_i) the coefficients s_i require four extra points to be computed; therefore, we need to introduce boundary conditions for the first and last two knots. In our implementation [96], we follow Ref. [88],

$$s_1 = \frac{3m_1 - m_2}{2}, \tag{A4a}$$

$$s_2 = \begin{cases} \frac{m_3+m_2}{2} & \text{if } w_2 + w_1 = 0 \\ \frac{w_2 m_1 + w_1 m_2}{w_2 + w_1} & \text{otherwise,} \end{cases} \tag{A4b}$$

$$s_{n-1} = \begin{cases} \frac{m_{n-3}+m_{n-2}}{2} & \text{if } w_{n-3} + w_{n-2} = 0 \\ \frac{w_{n-3} m_{n-2} + w_{n-2} m_{n-3}}{w_{n-3} + w_{n-2}} & \text{otherwise,} \end{cases} \tag{A4c}$$

$$s_n = \frac{3m_{n-1} - m_{n-2}}{2}. \tag{A4d}$$

Since the spline is defined as a piecewise cubic polynomial, in each interval $x \in [x_i, x_{i+1})$ we have

$$S(x) = a_i + b_i(x - x_i) + c_i(x - x_i)^2 + d_i(x - x_i)^3. \tag{A5}$$

The coefficients are determined by imposing the continuity of $S(x)$ and its first derivative $S'(x)$ at each knot,

$$S(x_i) = y_i, \qquad S(x_{i+1}) = y_{i+1}, \tag{A6a}$$

$$S'(x_i) = s_i, \qquad S'(x_{i+1}) = s_{i+1}. \tag{A6b}$$

Therefore, the polynomial coefficients read

$$a_i = y_i, \tag{A7a}$$

$$b_i = s_i, \tag{A7b}$$

$$c_i = \frac{3m_i - 2s_i - s_{i+1}}{x_{i+1} - x_i}, \tag{A7c}$$

$$d_i = \frac{s_i + s_{i+1} - 2m_i}{(x_{i+1} - x_i)^2}. \tag{A7d}$$

Since there are no conditions on the second derivative of $S(x)$, this is not necessarily continuous.

[1] M. Colpi *et al.*, arXiv:2402.07571.
[2] B. P. Abbott *et al.*, Phys. Rev. X **9**, 031040 (2019).
[3] R. Abbott *et al.*, Phys. Rev. X **11**, 021053 (2021).
[4] R. Abbott *et al.*, Phys. Rev. D **109**, 022001 (2024).
[5] R. Abbott *et al.*, Phys. Rev. X **13**, 041039 (2023).
[6] J. Antoniadis *et al.* (EPTA Collaboration and InPTA Collaboration), Astron. Astrophys. **685**, A94 (2024).
[7] A. Afzal *et al.* (NANOGrav Collaboration), Astrophys. J. Lett. **951**, L11 (2023); **971**, L27(E) (2024).
[8] D. J. Reardon *et al.*, Astrophys. J. Lett. **951**, L6 (2023).
[9] H. Xu *et al.*, Res. Astron. Astrophys. **23**, 075024 (2023).
[10] A. Klein *et al.*, Phys. Rev. D **93**, 024003 (2016).
[11] V. Korol *et al.*, Astron. Astrophys. **638**, A153 (2020).
[12] 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).
[13] D. Gerosa, S. Ma, K. W. K. Wong, E. Berti, R. O'Shaughnessy, Y. Chen, and K. Belczynski, Phys. Rev. D **99**, 103004 (2019).
[14] R. Buscicchio, A. Klein, E. Roebber, C. J. Moore, D. Gerosa, E. Finch, and A. Vecchio, Phys. Rev. D **104**, 044065 (2021).
[15] A. Toubiana, S. Babak, S. Marsat, and S. Ossokine, Phys. Rev. D **106**, 104034 (2022).
[16] N. Karnesis, S. Babak, M. Pieroni, N. Cornish, and T. Littenberg, Phys. Rev. D **104**, 043019 (2021).
[17] A. J. Farmer and E. S. Phinney, Mon. Not. R. Astron. Soc. **346**, 1197 (2003).
[18] 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.
[19] L. Barack and C. Cutler, Phys. Rev. D **70**, 122002 (2004).
[20] M. Bonetti and A. Sesana, Phys. Rev. D **102**, 103023 (2020).
[21] F. Pozzoli, S. Babak, A. Sesana, M. Bonetti, and N. Karnesis, Phys. Rev. D **108**, 103039 (2023).
[22] L. S. S. Team, Report No. ESA-L3-EST-SCI-RS-001 **5**, 65 (2018).
[23] C. Caprini, J. Phys. Conf. Ser. **610**, 012004 (2015).
[24] C. Caprini *et al.*, J. Cosmol. Astropart. Phys. 03 (2020) 024.
[25] 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.
[26] C. Caprini *et al.*, J. Cosmol. Astropart. Phys. 04 (2016) 001.
[27] P. Auclair, J. J. Blanco-Pillado, D. G. Figueroa, A. C. Jenkins, M. Lewicki, M. Sakellariadou, S. Sanidas, L. Sousa, D. A. Steer, J. M. Wachter, and S. Kuroyanagi (LISA Cosmology Working Group), J. Cosmol. Astropart. Phys. 04 (2020) 034.
[28] A. Ricciardone, J. Phys. Conf. Ser. **840**, 012030 (2017).
[29] R.-g. Cai, S. Pi, and M. Sasaki, Phys. Rev. Lett. **122**, 201101 (2019).
[30] N. Bartolo, V. De Luca, G. Franciolini, A. Lewis, M. Peloso, and A. Riotto, Phys. Rev. Lett. **122**, 211301 (2019).
[31] P. Auclair *et al.* (LISA Cosmology Working Group), Living Rev. Relativity **26**, 5 (2023).
[32] T. B. Littenberg and N. J. Cornish, Phys. Rev. D **107**, 063004 (2023).
[33] M. L. Katz, N. Karnesis, N. Korsakova, J. R. Gair, and N. Stergioulas, Phys. Rev. D **111**, 024060 (2025).
[34] S. Deng, S. Babak, M. Le Jeune, S. Marsat, E. Plagnol, and A. Sartirana, Phys. Rev. D **111**, 103014 (2025).
[35] S. H. Strub, L. Ferraioli, C. Schmelzbach, S. C. Stähler, and D. Giardini, Phys. Rev. D **110**, 024005 (2024).
[36] G. Boileau, N. Christensen, C. Gowling, M. Hindmarsh, and R. Meyer, J. Cosmol. Astropart. Phys. 02 (2023) 056.
[37] M. R. Adams and N. J. Cornish, Phys. Rev. D **89**, 022001 (2014).
[38] O. Hartwig, M. Lilley, M. Muratore, and M. Pieroni, Phys. Rev. D **107**, 123531 (2023).
[39] A. W. Criswell, S. Rieck, and V. Mandic, Phys. Rev. D **111**, 023025 (2025).
[40] M. Armano *et al.*, Phys. Rev. Lett. **116**, 231101 (2016).
[41] M. Armano *et al.* (LISA Pathfinder Collaboration), Phys. Rev. D **110**, 042004 (2024).
[42] M. Muratore, J. Gair, and L. Speri, Phys. Rev. D **109**, 042001 (2024).
[43] 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.

[44] R. Flauger, N. Karnesis, G. Nardini, M. Pieroni, A. Ricciardone, and J. Torrado, J. Cosmol. Astropart. Phys. 01 (2021) 059.
[45] Q. Baghi, N. Karnesis, J.-B. Bayle, M. Besançon, and H. Inchauspé, J. Cosmol. Astropart. Phys. 04 (2023) 066.
[46] F. Pozzoli, R. Buscicchio, C. J. Moore, F. Haardt, and A. Sesana, Phys. Rev. D **109**, 083029 (2024).
[47] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, MA, 2005).
[48] M. Tinto and S. V. Dhurandhar, Living Rev. Relativity **8**, 4 (2005).
[49] C. Cutler and E. E. Flanagan, Phys. Rev. D **49**, 2658 (1994).
[50] M. Vallisneri, Phys. Rev. D **77**, 042001 (2008).
[51] P. Amaro-Seoane *et al.*, arXiv:1702.00786.
[52] M. Tinto, G. Giampieri, R. W. Hellings, P. L. Bender, and J. E. Faller, in *Proceedings of the 7th Marcel Grossmann Meeting on General Relativity (MG 7)* (1994), pp. 1668–1670.
[53] J. W. Armstrong, F. B. Estabrook, and M. Tinto, Astrophys. J. **527**, 814 (1999).
[54] D. A. Shaddock, B. Ware, R. E. Spero, and M. Vallisneri, Phys. Rev. D **70**, 081101 (2004).
[55] D. A. Shaddock, M. Tinto, F. B. Estabrook, and J. W. Armstrong, Phys. Rev. D **68**, 061303 (2003).
[56] M. Vallisneri, Phys. Rev. D **72**, 042003 (2005); **76**, 109903 (E) (2007).
[57] S. Paczkowski, R. Giusteri, M. Hewitson, N. Karnesis, E. D. Fitzsimons, G. Wanner, and G. Heinzel, Phys. Rev. D **106**, 042005 (2022).
[58] M.-S. Hartig and G. Wanner, Phys. Rev. D **108**, 022008 (2023).
[59] G. Wanner, S. Shah, M. Staab, H. Wegener, and S. Paczkowski, Phys. Rev. D **110**, 022003 (2024).
[60] M.-S. Hartig, S. Schuster, and G. Wanner, J. Opt. **24**, 065601 (2022).
[61] M. Chwalla, K. Danzmann, M. D. Álvarez, J. E. Delgado, G. Fernández Barranco, E. Fitzsimons, O. Gerberding, G. Heinzel, C. Killow, M. Lieser, M. Perreur-Lloyd, D. Robertson, J. Rohr, S. Schuster, T. Schwarze, M. Tröbs, G. Wanner, and H. Ward, Phys. Rev. Appl. **14**, 014030 (2020).
[62] O. Hartwig and J.-B. Bayle, Phys. Rev. D **103**, 123027 (2021).
[63] O. Hartwig and M. Muratore, Phys. Rev. D **105**, 062006 (2022).
[64] M. Muratore, Time delay interferometry for LISA science and instrument characterization, Ph.D. thesis, Trento University, 2021.
[65] O. Hartwig, Instrumental modelling and noise reduction algorithms for the Laser Interferometer Space Antenna, Ph.D. thesis, Leibniz University, Hannover, 2021.
[66] M. B. Staab, Time-delay interferometric ranging for LISA: Statistical analysis of bias-free ranging using laser noise minimization, Ph.D. thesis, Leibniz University, Hannover, 2023.
[67] J. N. Reinhardt, Intersatellite clock synchronization and absolute ranging for space-based gravitational-wave detectors, Ph.D. thesis, Leibniz University, Hannover, 2025.
[68] M. Muratore, O. Hartwig, D. Vetrugno, S. Vitale, and W. J. Weber, Phys. Rev. D **107**, 082004 (2023).
[69] M. Muratore, D. Vetrugno, and S. Vitale, Classical Quantum Gravity **37**, 185019 (2020).
[70] LISA Data Challenges Working Group, LISA Data Challenges (2022).
[71] LISA Data Challenges Working Group, Radler dataset description.
[72] LISA Data Challenges Working Group, Sangria dataset description.
[73] LISA Data Challenges Working Group, Spritz dataset description.
[74] T. A. Prince, M. Tinto, S. L. Larson, and J. W. Armstrong, Phys. Rev. D **66**, 122002 (2002).
[75] S. Babak, M. Hewitson, and A. Petiteau, arXiv:2108.01167.
[76] C. Caprini and D. G. Figueroa, Classical Quantum Gravity **35**, 163001 (2018).
[77] N. Aghanim *et al.*, Astron. Astrophys. **641**, A6 (2020).
[78] L. Lehoucq, I. Dvorkin, R. Srinivasan, C. Pellouin, and A. Lamberts, Mon. Not. R. Astron. Soc. **526**, 4378 (2023).
[79] P. Binetruy, A. Bohe, C. Caprini, and J.-F. Dufaux, J. Cosmol. Astropart. Phys. 06 (2012) 027.
[80] https://github.com/martinaAEI/noise_knowledge_uncertainty.
[81] E. Thrane and J. D. Romano, Phys. Rev. D **88**, 124032 (2013).
[82] G. Janssen, G. Hobbs, M. McLaughlin, C. Bassa, A. Deller, M. Kramer, K. Lee, C. Mingarelli, P. Rosado, S. Sanidas, A. Sesana, L. Shao, I. Stairs, B. Stappers, and J. P. W. Verbiest, Proc. Sci., AASKA14 (**2015**) 037 [arXiv:1501.00127].
[83] M. C. Edwards, P. Maturana-Russel, R. Meyer, J. Gair, N. Korsakova, and N. Christensen, Phys. Rev. D **102**, 084062 (2020).
[84] D. Quang Nam, Y. Lemière, A. Petiteau, J.-B. Bayle, O. Hartwig, J. Martino, and M. Staab, Phys. Rev. D **108**, 082004 (2023).
[85] D. Quang Nam, J. Martino, Y. Lemière, A. Petiteau, J.-B. Bayle, O. Hartwig, and M. Staab, Phys. Rev. D **108**, 082004 (2023).
[86] O. Hartwig, M. Lilley, M. Muratore, and M. Pieroni, Phys. Rev. D **107**, 123531 (2023).
[87] H. Akima, J. ACM **17**, 589 (1970).
[88] Virtanen, Pauli, *et al.*, Nat. Methods **17**, 261 (2020).
[89] M. L. Katz, A. J. K. Chua, L. Speri, N. Warburton, and S. A. Hughes, Phys. Rev. D **104**, 064047 (2021).
[90] C. E. A. Chapman-Bird *et al.*, arXiv:2506.09470.
[91] P. J. Green, Biometrika **82**, 711 (1995).
[92] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, Publ. Astron. Soc. Pac. **125**, 306 (2013).
[93] N. Karnesis, M. L. Katz, N. Korsakova, J. R. Gair, and N. Stergioulas, Mon. Not. R. Astron. Soc. **526**, 4814 (2023).
[94] M. Katz, N. Karnesis, and N. Korsakova, mikekatz04/eryn: first full release (2023).
[95] M. Muratore, J. Gair, O. Hartwig, M. L. Katz, and A. Toubiana, Phys. Rev. D **112**, 063041 (2025).
[96] A. Santini, asantini29/cudakima: First official release (2024), https://github.com/asantini29/CudAkima.

[97] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis, in *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)* (2017).
[98] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: composable transformations of Python+NumPy programs (2018).
[99] J. Goodman and J. Weare, Commun. Appl. Math. Comput. Sci. **5**, 65 (2010).
[100] W. Xie, P. O. Lewis, Y. Fan, L. Kuo, and M.-H. Chen, Syst. Biol. **60**, 150 (2010).
[101] W. Del Pozzo, J. Veitch, and A. Vecchio, Phys. Rev. D **83**, 082002 (2011).
[102] A. Toubiana, K. W. K. Wong, S. Babak, E. Barausse, E. Berti, J. R. Gair, S. Marsat, and S. R. Taylor, Phys. Rev. D **104**, 083027 (2021).
[103] A. Toubiana, L. Pompili, A. Buonanno, J. R. Gair, and M. L. Katz, Phys. Rev. D **109**, 104019 (2024).
[104] P. Maturana-Russel, R. Meyer, J. Veitch, and N. Christensen, Phys. Rev. D **99**, 084006 (2019).
[105] E. M. Zahraoui, P. Maturana-Russel, W. van Straten, R. Meyer, and S. Gulyaev, Mon. Not. R. Astron. Soc. **540**, 3818 (2025).
[106] P. Whittle, J. R. Stat. Soc. Ser. B **15**, 125 (1953).
[107] G. Franciolini, M. Pieroni, A. Ricciardone, and J. D. Romano, arXiv:2505.24695.
[108] M. Muratore, D. Vetrugno, S. Vitale, and O. Hartwig, Phys. Rev. D **105**, 023009 (2022).
[109] R. E. Kass and A. E. Raftery, J. Am. Stat. Assoc. **90**, 773 (1995).
[110] F. Pozzoli, J. Gair, R. Buscicchio, and L. Speri, Phys. Rev. D **112**, 064035 (2025).
[111] D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. Van Der Linde, J. R. Stat. Soc. Ser. B **64**, 583 (2002).
[112] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*, Texts in Statistical Science Series (Chapman & Hall/CRC, Boca Raton, FL, 2004), 2nd ed., pp. xxvi–668.
[113] D. J. C. MacKay, *Information Theory, Inference & Learning Algorithms* (Cambridge University Press, Cambridge, England, 2002).
[114] A. Santini, asantini29/lisa-ps: First release (2025), https://github.com/asantini29/lisa-ps.
[115] C. R. Harris *et al.*, Nature (London) **585**, 357 (2020).
[116] J. D. Hunter, Comput. Sci. Eng. **9**, 90 (2007).
[117] S. K. Lam, A. Pitrou, and S. Seibert, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC* (2015), pp. 1–6.
[118] A. Santini, asantini29/pysco: First release (2024).
[119] The data used for this publication are available at: https://zenodo.org/records/17203220. The main code base used for the analysis is available at https://github.com/asantini29/lisa-ps.