

Forecasting constraints on scalar-induced gravitational waves with future pulsar timing array observations

Chiara Cecchini^{1,2,3,*}, Gabriele Franciolini^{3,†} and Mauro Pieroni^{3,‡}

¹*Department of Physics, University of Trento, Via Sommarive 14, 38122 Povo (Trento), Italy*

²*Trento Institute for Fundamental Physics and Applications-INFN,*

Via Sommarive 14, 38122 Povo (Trento), Italy

³*CERN, Theoretical Physics Department, Esplanade des Particules 1, Geneva 1211, Switzerland*



(Received 27 March 2025; accepted 19 May 2025; published 17 June 2025)

Pulsar timing arrays are playing a crucial role in the ongoing gravitational wave astronomy revolution. The current evidence for a stochastic gravitational wave background at nHz frequencies offers an opportunity to discover cosmological signals and threatens the observability of other subdominant gravitational waves (GWs). We explore prospects to constrain second-order scalar-induced GWs (SIGWs) associated with enhanced curvature perturbations in the primordial universe, forecasting realistic future pulsar timing array datasets. We assess how the currently observed signal could eventually limit future capabilities to search for GW relics of primordial phenomena and associated phenomenological consequences such as primordial black hole (PBH) formation. Given the sensitivity of PBH abundance to spectral parameters, measuring it remains a challenge for realistic signals. However, future observation could still rule out nearly subsolar mass PBHs formed through standard formation scenarios in some cases. Future progress in constraining PBH models is expected to stem from theoretical advancements in PBH computations, which should help resolve the tension between different computational methods. The analysis is based on and extends the Python code `fastPTA`.

DOI: [10.1103/nxx5-gx7d](https://doi.org/10.1103/nxx5-gx7d)

I. INTRODUCTION

Pulsar timing array (PTA) experiments represent an unprecedented opportunity to investigate the spectrum of gravitational waves (GWs) in the nHz frequency range. Recently, various PTA collaborations reported substantial evidence of a stochastic common spectrum with an angular correlation pattern consistent with the quadrupolar nature of GWs in general relativity [1]. As reported by NANOGrav (NG) [2], EPTA (in combination with InPTA) [3], PPTA [4], and CPTA [5], the evidence reaches $(2 \div 4)\sigma$ significance and is expected to grow in the not so distant future [6–8]. Most recently, the MeerKAT PTA (MPTA) reported similar results [9]. The excess is compatible with a stochastic GW background (SGWB) characterized by a blue-tilted spectrum of energy density $\Omega_{\text{GW}} \propto f^{n_T}$, with a best-fit value of $n_T \simeq 2$ and relatively large uncertainties.

In the astrophysical setting, this signal can be produced by the incoherent superposition of GWs emitted by a population of inspiralling supermassive black hole (SMBH) binaries [10,11]. While $\Omega_{\text{GW}} \propto f^{2/3}$ for circular orbits in the GW-driven regime, eccentricity and environmental effects can modify binary evolution. Additionally, the discrete number of sources and statistical fluctuation thereof in each frequency bin can lead to different predictions [12–14].

In a cosmological setting, plausible sources of nHz frequencies are extremely diverse. For example, first-order phase transitions, cosmic strings, domain walls, scalar-induced GWs (SIGWs), linear tensor perturbations generated during inflation, among others, could contribute to the observed signal, as addressed by an impressive body of literature (see Ref. [15] for a review). Inflationary models capable of amplifying scalar fluctuations are particularly promising candidates for generating an SGWB within the frequency range detectable by PTAs. These enhanced scalar fluctuations lead to the production of SIGWs at second order in perturbations [16–25], which could give rise to a significant SGWB. Amplified scalar perturbations can naturally emerge from single-field inflationary models with features in the potential, such as those inducing ultra-slow-roll phases. They may also arise in multifield inflationary scenarios or due to mechanisms like preheating and early matter-dominated periods.

* Contact author: chiara.cecchini@unitn.it

† Contact author: gabriele.franciolini@cern.ch

‡ Contact author: mauro.pieroni@cern.ch

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) 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 CERN.

Notably, the same scalar perturbations responsible for SIGWs are linked to the potential formation of primordial black holes (PBHs) [26–31]. When SIGWs fall in the nHz frequency band, they correspond to the stellar mass range of PBHs, which can contribute to a part of the dark matter and are complementarily searched for through lensing and ground-based GW searches looking for mergers (see Ref. [32] for a review). Through the detection of SIGWs or by establishing constraints on their amplitude, PTA data has the potential to provide crucial insights into the inflationary epoch, potentially strengthening bounds on dark matter candidates. On the other hand, the overproduction of dark matter in the form of PBHs in the stellar mass range implies a bound around $\mathcal{P}_\zeta \leq \mathcal{O}(10^{-2})$ [33–37] on the scalar curvature perturbations and thus also on the strength of the SIGW in the nHz frequency band. For all of these reasons, SIGWs are a compelling and high-priority observational target for PTAs. Thus, exploring PTA capabilities to constrain such a scenario is timely, and including theoretical information on SIGWs can help interpret current and future data.

Currently, many cosmological and astrophysical models compatible with the data predict GW spectra with similar shapes, and existing measurement uncertainties do not allow us to distinguish between them [10,38–41]. Therefore, it is possible that SIGWs are the dominant source, or that an astrophysical background dominates PTA observations, and SIGW searches would need to dig below the astrophysical foreground, as we will show. We do not consider the possibility of multiple cosmological signals, although the same logic would apply in that case too. We will show that future datasets will improve our current understanding, showing, however, how systematic uncertainties on the computation of the PBH abundance remain and could limit our ability to draw stringent conclusions (see also [42,43]).

This paper is organized as follows. In Sec. II we discuss the phenomenological parametrization we adopt of scenarios with enhanced curvature power spectra at small scales, including how to compute the spectrum of SIGW. In Sec. III we describe the essence of PTA searches for GWs, and the construction of the PTA likelihood that we adopt to estimate current and future sensitivity to SIGWs. In Sec. IV we forecast the future bound on both motivated astrophysical and SIGW signals and conclude the ability of future PTAs to constrain the spectrum of curvature perturbations. Moreover, in Sec. V we connect these results to the related abundance of PBH, commenting on the challenges associated with setting robust bounds due to the large model dependence and systematic uncertainties. We conclude in Sec. VI with some discussion and outlook.

II. SECOND-ORDER INDUCED GWS

The generation of GWs from scalar (curvature) perturbations is a direct consequence of general relativity. While scalar and tensor modes do not couple at first order in

perturbation theory, second-order tensor perturbations are sourced by first-order scalar modes. However, producing an observable SGWB requires sufficiently large scalar perturbations. Although the amplitude and shape of the curvature power spectrum are tightly constrained by cosmic microwave background (CMB) and large-scale structure observations on scales $k \lesssim \text{Mpc}^{-1}$, they remain unknown at smaller scales. The frequency range probed by PTAs corresponds to scales $k \sim (10^6 - 10^8) \text{Mpc}^{-1}$ which entered the Hubble sphere around the epoch of the QCD crossover [42,44–46]. Therefore, measuring SIGWs at PTA scales indirectly probes small-scale enhancements of the curvature power spectrum at such scales.

We note that SIGWs would also contribute to the effective number of neutrino species as $N_{\text{eff}} \equiv 3.044 + \Delta N_{\text{eff}}^{\text{GW}}$, with $\Delta N_{\text{eff}}^{\text{GW}} = \rho_{\text{GW}}/\rho_{\nu,1}$ and $\rho_{\nu,1}$ is the energy density of a single neutrino species. As a consequence of the measurements of the CMB [47] and baryon acoustic oscillations, which constrain $\Delta N_{\text{eff}} \leq 0.28$ at 95% CL, the total (integrated) GW abundance is forced to be [15]

$$h^2 \Omega_{\text{GW}} \simeq 1.6 \times 10^{-6} \left(\frac{\Delta N_{\text{eff}}^{\text{GW}}}{0.28} \right). \quad (1)$$

We will not include this bound in our analyses, as it is much weaker than the one coming from PBH overproduction, as we will see.

The current day abundance can be expressed as (see [15])

$$h^2 \Omega_{\text{GW}}(k) = \frac{h^2 \Omega_r}{24} \left(\frac{g_*}{g_*^0} \right) \left(\frac{g_{*s}}{g_{*s}^0} \right)^{-\frac{4}{3}} \mathcal{P}_h(k), \quad (2)$$

as a function of the tensor power spectrum $\mathcal{P}_h(k)$, where $g_{*s} \equiv g_{*s}(T_k)$ and $g_* \equiv g_*(T_k)$ represent the effective entropy and energy degrees of freedom at the time of Hubble crossing of mode k and at present (denoted by superscript 0), respectively. The current radiation abundance is $h^2 \Omega_r = 4.2 \times 10^{-5}$, while $g_*^0 = 3.38$ and $g_{*s}^0 = 3.93$.

Each mode k crosses the Hubble sphere at a temperature T_k given by

$$k = 1.5 \times 10^7 \text{Mpc}^{-1} \left(\frac{g_*}{106.75} \right)^{\frac{1}{2}} \left(\frac{g_{*s}}{106.75} \right)^{-\frac{1}{3}} \left(\frac{T_k}{\text{GeV}} \right), \quad (3)$$

and corresponds to a current GW frequency

$$f \equiv k/2\pi = 1.6 \text{ nHz} \left(\frac{k}{10^6 \text{Mpc}^{-1}} \right). \quad (4)$$

Finally, limiting ourselves to Gaussian curvature perturbations, the tensor mode power spectrum is expressed as [48,49]

$$\overline{\mathcal{P}_h(\eta, k)} = 4 \int_0^\infty dt \int_0^1 ds \left[\frac{t(2+t)(1-s^2)}{(1-s+t)(1+s+t)} \right]^2 \times \overline{I^2(t, s, k, \eta)} \times \mathcal{P}_\zeta(uk) \mathcal{P}_\zeta(vk), \quad (5)$$

where the transfer function assuming radiation domination (RD) can be written as

$$\overline{I_{\text{RD}}^2(t, s, k\eta \rightarrow \infty)} = \frac{1}{2(k\eta)^2} I_A^2(u, v) [I_B^2(u, v) + I_C^2(u, v)], \quad (6)$$

where

$$\begin{aligned} I_A(u, v) &= \frac{3(u^2 + v^2 - 3)}{4u^3 v^3}, \\ I_B(u, v) &= -4uv + (u^2 + v^2 - 3) \log \left| \frac{3 - (u+v)^2}{3 - (u-v)^2} \right|, \\ I_C(u, v) &= \pi(u^2 + v^2 - 3) \Theta(u + v - \sqrt{3}), \end{aligned} \quad (7)$$

with $u \equiv (1 + s + t)/2$, and $v \equiv (1 - s + t)/2$.

To simplify the analyses, we assume that the nHz SIGW background is emitted during an approximately radiation-dominated era in the early universe, ignoring equation of state and sound speed variation during the QCD era [45,50], which imprints specific features on the low-frequency tail of the SGWB [46].¹ Additionally, cosmic expansion may be influenced by unknown dark sector physics, potentially causing transient matter domination or kination [55–60]. Such nonstandard cosmologies can

significantly affect SIGW and PBH production [61–69]. We leave the exploration of these scenarios for future work.

A. SIGW template for a log-normal spectrum

The spectrum of SIGWs depends on the features of the curvature power spectrum $\mathcal{P}_\zeta(k)$, as described in Eq. (5). This is fully model dependent. In the simplest scenarios, for example, the enhanced curvature perturbations are induced by features in the potential of single-field inflation. See, e.g., Ref. [70] for a review on model building. One could perform the analysis by adopting different approaches. In particular, one could consider an explicit model of the early universe, and compute the curvature power spectrum from first principles, varying the model parameters. On the other hand, as often done in the literature, one can adopt a template-based method, where the shape of the curvature power spectrum is assumed to follow a given functional form, inspired by an explicit model, which is controlled by phenomenological parameters (see, e.g., [71] for a comparison of these approaches in the context of SIGWs in LISA). We will follow the latter approach and remain agnostic on the model producing the enhancement of curvature perturbations at small scales. Therefore, we consider a log-normal parametrization of the curvature power spectrum of the form²

$$\mathcal{P}_\zeta(k) = \frac{A_\zeta}{\sqrt{2\pi}\Delta} \exp \left[-\frac{\log^2(k/k_*)}{2\Delta^2} \right], \quad (8)$$

where k_* is the peak scale. This assumption allows us to speed up the computation of the double integral in Eq. (5) and express the spectrum of SIGWs analytically, as [72] (see also [42])

$$\begin{aligned} h^2 \Omega_{\text{GW}}(f) &\approx \frac{4}{5\sqrt{\pi}\Delta} c_g(f) A_\zeta^2 \kappa^3 \exp^{\frac{9\Delta^2}{4}} \left[\left(\mathcal{K}^2 + \frac{\Delta^2}{2} \right) \text{erfc} \left(\frac{\mathcal{K} + \frac{1}{2} \log \frac{3}{2}}{\Delta} \right) - \frac{\Delta}{\sqrt{\pi}} \exp \left(-\frac{\left(\mathcal{K} + \frac{1}{2} \log \frac{3}{2} \right)^2}{\Delta^2} \right) \left(\mathcal{K} - \frac{1}{2} \log \frac{3}{2} \right) \right] \\ &+ \frac{0.0659\kappa^2}{\Delta^2} e^{\Delta^2} \exp \left[-\frac{\left(\log \kappa + \Delta^2 - \frac{1}{2} \log \frac{4}{3} \right)^2}{\Delta^2} \right] + \frac{1}{3} \sqrt{\frac{2}{\pi}} \frac{e^{8\Delta^2}}{\kappa^4 \Delta} \exp \left(-\frac{\log^2 \kappa}{2\Delta^2} \right) \text{erfc} \left(\frac{4\Delta^2 - \log \kappa/4}{\sqrt{2}\Delta} \right), \end{aligned} \quad (9)$$

where $\kappa = f/f_*$, f_* being the peak frequency $f_* = k_*/(2\pi)$, $\mathcal{K} = \log(\kappa \exp(3\Delta^2/2))$, and we introduced

$$c_g(f) = h^2 \Omega_r \left(\frac{g_*}{g_*^0} \right) \left(\frac{g_{*S}}{g_{*S}^0} \right)^{-\frac{4}{3}}. \quad (10)$$

Note that Eq. (9) is valid in the broad peak regime of the log-normal parametrization in Eq. (8). We verified that it represents a sufficiently accurate template in the range of Δ we restrict our analyses.

¹The abundance of late-time universe SIGWs emitted during a RD phase is gauge independent [51–54].

²The choice of normalization of the amplitude $\propto A_\zeta/\Delta$ is standard in the literature, as it makes the variance of perturbations finite in the small width limit. However, this choice inevitably introduces degeneracies between the overall A_ζ amplitude and Δ .

III. FORECASTING FUTURE PTA CONSTRAINTS

In this section, we outline the methods adopted to estimate future PTA sensitivity, building on the framework first introduced in Ref. [8]. We assume the PTA data consists of residuals of time of arrivals (TOAs) $d_I(t)$, where $I = 1, \dots, N_p$ is the pulsar index. We further assume stationarity and introduce the Fourier domain data as

$$\tilde{d}_I^k = \int_{-T_{\text{obs}}/2}^{T_{\text{obs}}/2} d_I(t) e^{-i2\pi f_k t} dt, \quad (11)$$

where T_{obs} is the observation time and $f_k = k/T_{\text{obs}}$ are the Fourier basis components. We take the data as the sum of two zero mean, Gaussian processes: the stochastic GW signal and the noise part, $\tilde{d}_I^k = \tilde{s}_I^k + \tilde{n}_I^k$. Then, all of the information is contained in the covariance matrix,

$$C_{IJ} = C_{h,IJ} + C_{n,IJ}. \quad (12)$$

To model the SGWB, we take it as unpolarized and isotropic, which means that the Fourier modes $\tilde{h}_P(f, \hat{k})$ of the metric spin-2 perturbation follows³:

$$\langle \tilde{h}_P(f, \hat{k}) \tilde{h}_{P'}^*(f', \hat{k}') \rangle = \frac{1}{16\pi} S_h(f) \delta(f - f') \delta_{PP'} \delta^2(\hat{k}, \hat{k}'), \quad (13)$$

where $P = \{+, \times\}$ denotes the two polarizations and $S_h(f)$ is the (one-sided) strain power spectral density of the SGWB, related to the SGWB energy density as

$$\Omega_{\text{GW}} h^2 = \frac{h^2}{\rho_c} \frac{d\rho_{\text{GW}}}{d\log f} \equiv \frac{2\pi^2 f^3}{3H_0^2/h^2} S_h, \quad (14)$$

where $\rho_c/h^2 = 3(H_0/h)^2/(8\pi G)$ is the Universe critical energy density and $H_0/h = 1/(9.78 \text{ Gyr})$ is the Hubble parameter today.

Under the above assumptions, we can show that (see [8] for the details) the time delay covariance matrix is given by

$$\langle \tilde{\Delta t}_I \tilde{\Delta t}_J^* \rangle \equiv \frac{1}{2} \delta(f - f') C_{h,IJ} = \frac{1}{2} \delta(f - f') R_{IJ} S_h(f), \quad (15)$$

where we have defined

$$R_{IJ} = \chi_{IJ} \cdot \mathcal{R}(f) [T_I(f) T_J(f) T_{IJ}/T_{\text{obs}}]^{1/2}. \quad (16)$$

The first geometrical factor χ_{IJ} is the well-known Hellings-Downs correlation pattern as a function of the angular separation between any pair of pulsars [1]. The time-dependent part accounts for the pulsar's transmission function $\mathcal{T}(f)$, which models the low-frequency suppression of

the extracted signal [74], and the individual observation time T_I . The off-diagonal components are limited by the effective overlapping time between pulsars I and J , defined as $T_{IJ} = \min[T_I, T_J]$. We have further defined the sky-averaged quadratic response function $\mathcal{R}(f) \equiv 1/12\pi^2 f^2$.

We assume the noise contribution to the covariance matrix to be dominated by the diagonal terms, i.e., to be uncorrelated between pulsars. We thus approximate $C_{n,IJ} = \delta_{IJ} P_{n,I}$. We model the noise budget as a white noise (WN) with the addition of a red noise (RN) component, along with the temporal variations in dispersion measurements (DM) and scattering variations (SV),

$$P_{n,I} = P_I^{\text{WN}} + A_I^{\text{RN}} \left(\frac{f}{f_r} \right)^{\gamma_I^{\text{RN}}} + P_I^{\text{DM}} + P_I^{\text{SV}}, \quad (17)$$

where we arbitrarily fix the reference frequency at $f_r = f_{\text{yr}}$ and $\gamma_I^{\text{RN}} < 0$. See [8] for more details on the noise modeling adopted in fastPTA, along with its reliability in recovering the performance of current PTA datasets.

In what follows, we will simulate future PTA datasets by sampling random realizations of the noise components of the pulsars from the distribution of the *observed* properties of those contained in the current EPTA catalog. Also, we will not include noise parameters in the inference, but we will fix them at the best fit. While this may approximate some of the possible correlations between signal and noise, this choice was shown to have a negligible impact on the forecasts, as validated in [8]. Also, as we extrapolate to larger datasets, the observations enter the signal-dominated regime in the most relevant frequency range, which further justifies these approximations.

The log-likelihood of Gaussian and zero-mean data $\tilde{d}^k$ entirely described by their variance $C_{IJ}(f_k, \theta)$ is given by [75,76]

$$-\ln \mathcal{L}(\tilde{d}|\theta) = \text{const} + \sum_{k,IJ} [\ln [C_{IJ}(f_k, \theta)] + \tilde{d}_I^k C_{IJ}^{-1}(f_k, \theta) \tilde{d}_J^{k*}/T_{\text{obs}}], \quad (18)$$

also known as Whittle's likelihood, which represents the probability of the data $\tilde{d}$ given the model parameters θ . This is the main ingredient for the fast forecasting technique introduced in [8], which entails evaluating the Fisher information matrix (FIM). The FIM is defined as

$$F_{\alpha\beta} \equiv - \left. \frac{\partial^2 \ln \mathcal{L}}{\partial \theta^\alpha \partial \theta^\beta} \right|_{\theta=\theta_0} = \sum_k \text{Tr} \left[C^{-1} \frac{\partial C}{\partial \theta^\alpha} C^{-1} \frac{\partial C}{\partial \theta^\beta} \right]_{\theta=\theta_0}, \quad (19)$$

where θ_0 is the maximum likelihood estimator of the model parameters, obtained by imposing $(\partial \ln \mathcal{L} / \partial \theta^\alpha)|_{\theta_0} = 0$. For the likelihood in Eq. (18), the FIM reads

³See [73] for the inclusion of anisotropies.

$$F_{\alpha\beta} \equiv \sum_{f_k} C_{IJ}^{-1} C_{KL}^{-1} \frac{\partial(R_{JK}S_h)}{\partial\theta^\alpha} \frac{\partial(R_{LI}S_h)}{\partial\theta^\beta}. \quad (20)$$

The covariance matrix is then computed by inverting the FIM, and the uncertainties on the model parameters are easily obtained as $\sigma_\alpha \equiv \sqrt{F_{\alpha\alpha}^{-1}}$. This means that estimating the uncertainties boils down to inverting the FIM, with dramatically reduced computational times compared to a full-time-domain Bayesian analysis.

The FIM is considered to be a reliable estimator of the uncertainties in the limit of large signal-to-noise ratio, for which, typically, the posterior distribution can effectively be approximated by a Gaussian distribution. When this condition is not satisfied, we will compare the FIM results with full-fledged Bayesian parameter estimation methods. To this aim, we simulate frequency-domain data $\tilde{d}^k \equiv \tilde{d}(f_k)$, where k runs over the frequencies within detector sensitivity. We then generate zero-mean Gaussian realizations of signal and noise components, whose variance is defined by their spectral densities. The parameter estimation is performed via Markov chain Monte Carlo (MCMC) using EMCEE [77].

It is also instructive to define an effective sensitivity, which visually captures the reach of a given PTA dataset. Following Ref. [8], we define (see also [74])

$$S_{\text{eff}}(f) = (C_{IJ}^{-1} C_{KL}^{-1} R_{JK} R_{LI})^{-1/2}, \quad (21)$$

in such a way that the SNR computation takes the familiar form $\text{SNR}^2 = \sum_k (S_h/S_{\text{eff}})^2 = T_{\text{obs}} \int df (S_h/S_{\text{eff}})^2$.

IV. FUTURE BOUNDS ON SIGWS

In this section, we forecast the future reach of PTA datasets to constrain SIGWs and, indirectly, the primordial curvature power spectrum. We do so employing two working assumptions:

- (i) In the first stage (Sec. IV A), we assume that the current GW excess observed in the PTA dataset is due to a cosmological SIGW. Under this hypothesis, we show how the currently large uncertainties on the amplitude of the curvature perturbation will shrink significantly with longer observation time and larger datasets. The fact that we are observing this signal mostly in its low-frequency tail does not allow us to completely break the degeneracy between peak frequency and amplitude, leaving large uncertainties on the total amplitude of curvature perturbations.
- (ii) In a second stage (Sec. IV B), we assume that the current PTA signal is dominated by an astrophysical SGWB sourced by SMBH binaries, which we model with a power law (PL), while SIGWs are only a subdominant contribution. As in the first case, uncertainties will shrink significantly with larger datasets and the possibility of resolving the signal

peak below the dominant PL results in potentially stronger constraints on the overall SIGW amplitude.

A. SIGWs as the dominant signal

Under the assumption of a dominant SIGW signal in the nHz frequency range, our model parameters are

$$\theta = \{\log_{10} A_\zeta, \log_{10} \Delta, \log_{10}(f_*/\text{Hz})\}. \quad (22)$$

We simulate the observation of an SGWB dominated by SIGWs with

$$\begin{aligned} \log_{10} A_\zeta &= -1.64, \\ \log_{10} \Delta &= -0.15, \\ \log_{10}(f_*/\text{Hz}) &= -7.8. \end{aligned} \quad (23)$$

We choose this set of parameters as they are compatible with the current PTA observations (NG15, EPTA + InPTA, and PPTA) within 3σ [40,41] while complying with the PBH overproduction bounds assuming small non-Gaussianities, see below [78] for details. We will come back to the PBH abundance computation in Sec. V. We checked that assuming $N_p = 68$ and current noise configurations, the expected uncertainties on the SIGW noise parameters match current constraints based on available PTA data (see, e.g., [78]).

In Fig. 1, we report the forecasts obtained with a dataset assuming Square Kilometre Array (SKA)-like noise containing either $N_p = 100$ or $N_p = 400$ pulsars observed for $T_{\text{obs}} = 10$ yr. In the left panel, we show the effective sensitivity defined in Eq. (21), for both cases, alongside the SIGW signal with parameters in Eq. (23). In the right panel, we report the posterior distribution for the model parameters, inferred using the FIM approach. We observe a striking correlation between the amplitude A_ζ and the peak frequency f_* . This tight correlation arises because most of the constraining power comes from the observation of the low-frequency tail of the signal, where the signal is well approximated by a power law, while the signal peak is more loosely constrained. In such a regime, there is a leftover degeneracy between the amplitude and the pivot scale, so that larger pivot scales can be compensated by increasing the amplitude of the signal accordingly. Similar correlations, albeit slightly more pronounced, are found with the width Δ , due to the normalization condition in Eq. (8). Overall, we conclude that, given also the large degeneracies observed, the bounds on the model parameters remain rather weak. In terms of reconstructed Ω_{GW} , we expect it to be best constrained around $f \sim 10^{-8}$ Hz, with degeneracies maintaining fixed the best-observed region of the GW spectrum.

In Fig. 2, we show the posterior predictive distribution (PPD) of Ω_{GW} for the case with $N_p = 400$. We define the PPD as




FIG. 1. Left panel: SGWB abundance for the SIGW-only scenario with $\log_{10} A_\zeta = -1.64$, $\log_{10} \Delta = -0.15$, and $\log_{10}(f_*/\text{Hz}) = -7.8$. With red and blue lines we indicate the effective sensitivity for a future dataset with $N_p = 100$ and $N_p = 400$ pulsars with SKA-like noise budget and $T_{\text{obs}} = 10$ yr. Right panel: corner plot reporting the corresponding Fisher matrix uncertainty estimates.

$$p(\Omega_{\text{GW}}|d) = \int p(\Omega_{\text{GW}}|\theta)p(\theta|d)d\theta, \quad (24)$$

where we introduced $p(\theta|d)$ as the posterior distribution of the model parameters θ given the simulated data d , while $p(\Omega_{\text{GW}}|\theta)$ is the likelihood of new signal Ω_{GW} given a set of parameters θ . The integral marginalizes over all possible parameter values θ , weighted by their posterior probabilities. We obtain an upper bound on the signal amplitude at all scales, as the signal is observed with sufficient SNR across the whole range of frequencies. The high frequency tail is less constrained, as it fully depends on the width of the spectrum having $\Omega_{\text{GW}}(f \gg f_*) \propto \mathcal{P}_\zeta^2(f \gg f_*)$.




FIG. 2. Posterior predictive distribution of the SIGW measured by a simulated SKA-like dataset with $N_p = 400$, as shown in Fig. 1. Light red lines correspond to spectra obtained by random sampling points from the posterior distribution of model parameters. The red solid line indicates the 90% CL, while the red dashed line is the injected Ω_{GW} .

Narrow spectra have a characteristic double peak feature (see, e.g., discussion in [23,79]) that leads to the observed shape, a feature which is sufficiently well reconstructed. The reconstructed signal in the low-frequency tail is bounded to be very close to the injection, even extrapolating below the observed frequencies. This is because the causality tail does not allow for spectra narrower than $\Omega_{\text{GW}} \sim f^3$ [80–83] below the most constrained region of the spectrum.

We provide fitting formulas for the precision achieved as a function of the observation time and number of pulsars. The relative uncertainties scale as

$$\begin{aligned} \frac{\sigma_{A_\zeta}}{A_\zeta} &= 16\% \left(\frac{N_p}{70} \right)^{-1/2}, \\ \frac{\sigma_\Delta}{\Delta} &= 21\% \left(\frac{N_p}{70} \right)^{-1/2}, \\ \frac{\sigma_{f_*}}{f_*} &= 14\% \left(\frac{N_p}{70} \right)^{-1/2}. \end{aligned} \quad (25)$$

In the scaling considered above, we do not include the fact that, as a larger and larger set of pulsars becomes available, the SNR of the signal would become larger at higher frequencies, enhancing the number of effective frequencies, thus further reducing the uncertainties in the high-frequency portion of the signal.

Overall, we conclude that, given the relatively low amplitude of the SIGW injected in this scenario, Ω_{GW} would remain relatively weakly constrained also with realistic future datasets. To have tighter bounds, one would need signals with smaller f_* for which the peak appears in the observable band, while still complying with the

maximum $\mathcal{P}_\zeta$ imposed by the PBH abundance upper bound. Having the peak more clearly in the band would also reduce degeneracies. We will come back to the implication of such constraints for the production of PBHs.

B. SIGWs as a subdominant signal

In this section, we proceed with a complementary approach and assume the SIGW only amounts to a small portion of the nHz SGWB, which is instead already dominated by an astrophysical signal originating from the mergers of SMBH binaries.

1. Modeling the astrophysical foreground

In this work, we do not attempt to model the astrophysical signal from first principles but only adopt a simple phenomenological parametrization of its spectrum. We, therefore, assume the astrophysical SGWB to be described by a PL compatible with the recent observations (however, see [84])

$$h^2\Omega_{\text{GW}}^{\text{PL}}(f) = 10^{\alpha_{\text{PL}}} \left(\frac{f}{f_{\text{yr}}} \right)^{n_T}. \quad (26)$$

For definiteness, when performing the injections, the parameters are fixed to the maximum likelihood values obtained by the EPTA [3] collaboration, which are also in good agreement with the other observations [85], unless otherwise stated. Therefore, we fix

$$\begin{aligned} \alpha_{\text{PL}} &= -7.2, \\ n_T &= 2. \end{aligned} \quad (27)$$

Here, it is important to mention that relative uncertainties on the amplitude and tilt are of the order of 60% and 15%, respectively. We will explore how the result changes by varying the parameters within the 2σ range. To make contact with other quantities used in the literature, the spectrum in Eq. (26) corresponds to a strain amplitude ($h_c \equiv \sqrt{f S_h}$) at $f = f_{\text{yr}}$ of $A_{h_c} \simeq 10^{-14}$ and pulsar timing residuals spectral index $\gamma = 5 - n_T = 3$.

In the new physics search of the NANOGrav collaboration analysis [86], astrophysical population modeling was used to derive predictions for the amplitude and tilt of the SMBH SGWB spectrum, i.e., the parameters α_{PL} and n_T . These predictions, marginalized over the SMBH population model parameters, were then used as prior in the Bayesian inference analysis, on par with the cosmological bounds on the amplitude of curvature perturbation, for example. Thus, they worked with a bivariate normal distribution for amplitude and tilt of the spectrum [11,86]. When translated in terms of $(\alpha_{\text{PL}}, n_T)$, the mean μ_{SMBH} and covariance σ_{SMBH} are given by

$$\mu_{\text{SMBH}} = \begin{pmatrix} -10.4 \\ 0.30 \end{pmatrix}, \quad \sigma_{\text{SMBH}} = \begin{pmatrix} 1.1 & 0.0054 \\ 0.0054 & 0.12 \end{pmatrix}. \quad (28)$$

This prior distribution forces the amplitude of the SGWB to preferentially fall at smaller values than what is currently observed, thus downplaying the SMBH model in favor of new physics [86]. In this work, we assume a more agnostic approach and use uninformative priors with distant boundaries on both parameters.

To summarize, the total SGWB is assumed to be

$$h^2\Omega_{\text{GW}} = h^2\Omega_{\text{GW}}^{\text{PL}} + h^2\Omega_{\text{GW}}^{\text{SIGW}}, \quad (29)$$

and the model parameters in this case are

$$\theta = \{\alpha_{\text{PL}}, n_T, \log_{10} A_\zeta, \log_{10} \Delta, \log_{10}(f_*/\text{Hz})\}. \quad (30)$$

2. Multimodel analysis

We simulate a scenario where the dominant PL contribution (parameters of the PL are fixed to the reference values) is hiding a subdominant SIGW with parameters fixed to representative values:

$$\begin{aligned} \log_{10} A_\zeta &= -1.7, \\ \log_{10} \Delta &= -0.3, \\ \log_{10}(f_*/\text{Hz}) &= -7.8. \end{aligned} \quad (31)$$

In Fig. 3 (left panel), we plot the two separate contributions to the SGWB abundance and the future effective sensitivity achieved by the SKA 10 yr, assuming $N_p = 100, 400$ pulsars. Notice the injected SIGW signal is always subdominant with respect to the PL in the observed frequency range. First of all, the presence of the dominant PL contribution significantly degrades the effective sensitivity of the simulated dataset [8]. However, with a sufficient number of pulsars, the combined SGWB can still be measured with sufficient accuracy to allow constraining subdominant contribution with an abundance that is up to $\mathcal{O}(1)$ orders of magnitude below the astrophysical PL.

We show the corresponding posterior distribution of parameters in the right panel of Fig. 3. One obtains relative uncertainties of the order of $\mathcal{O}(20)\%$ on the curvature power spectrum parameters, roughly indicating SIGW detection at 5σ CL with more than $N_p = 400$ pulsars. Focusing on the SIGW parameters only, we see correlations similar to the ones observed in Fig. 1. However, since the SIGW peak appears in the sensitivity of the simulated dataset, these correlations are less pronounced. There exists a negative correlation between α_{PL} and A_ζ . This appears since PTA datasets are sensitive to the sum of the signals, and a stronger SIGW can be compensated by a fainter




FIG. 3. Same as Fig. 1, but assuming the presence of astrophysical SGWB with $\alpha_{\text{PL}} = -7.2$ and $n_T = 2$, while $\log_{10} A_\zeta = -1.7$, $\log_{10} \Delta = -0.3$, and $\log_{10}(f_*/\text{Hz}) = -7.8$. Right panel: corner plot reporting Fisher matrix uncertainty estimates for the same scenario shown in the left panel. The red (blue) posterior assumes a SKA-like with 100 (400) pulsars and $T_{\text{obs}} = 10$ yr.

astrophysical SGWB, and vice versa. Also, a similar positive correlation is observed with n_T , as a larger tilt removes power from the lower frequencies of relevance for the dataset we consider, due to our choice of pivot scale f_{yr} to normalize the PL. Finally, a negative correlation between f_* and α_{PL} appears. This is most likely a combination of multiple trends. As a SIGW pivot at lower frequencies correlates with a narrower tilt, the steeper low-frequency SIGW behavior is correlated with a shallower PL and a larger α_{PL} .

In Fig. 4, we report the PPD obtained from the results of the analysis, assuming $N_p = 400$. The PL signal is very tightly constrained up to frequencies of around




FIG. 4. Same as 2 for the SIGW + SMBH case. We show the posterior predictive distribution for the power law (green) and SIGW (red). We again assume a simulated SKA-like dataset with $N_p = 400$.

$f \gtrsim f \sim 10^{-7}$ Hz. The subdominant SIGW signal is best constrained in the low-frequency region of the spectrum, while it becomes unbounded from below above $f \gtrsim 2 \times 10^{-8}$ Hz. As the peak of the injected signal partly appears within the sensitivity band of the dataset, the width of the spectrum is more tightly constrained, compared to the previous case (see Fig. 2). As a consequence, the uncertainty on the signal amplitude mostly rescales the signal vertically around the injected spectrum in the low-frequency range.

While our results are subject to the assumption $\mathcal{P}_\zeta$ is described by a log-normal (LN) shape, we expect the qualitative conclusions to remain valid also for more complex shapes. It would be interesting to perform analyses with more flexible models (see, e.g., [71]). We leave this for future work.

It is interesting to explore how the uncertainties on the assumed subdominant SIGW background vary by modifying the assumptions. In the following, we explore how the uncertainties change as a function of

- (i) the number of pulsars in the dataset;
- (ii) the amplitude of the dominant astro-SGWB;
- (iii) the SIGW signal parameters.

3. Scaling with the number of pulsars

Still assuming a multimodel inference, in Fig. 5 we report how the uncertainties vary as a function of the number of pulsars in the dataset N_p . We derive these uncertainties assuming the reference values of the PL parameters, as well as the SIGW parameters, as defined in Eq. (31). The relative uncertainties retain an overall




FIG. 5. Scaling of the uncertainties on α_{PL} and n_T (first column), $\log_{10} A_\zeta$ (second column), $\log_{10} \Delta$ (third column), and $\log_{10}(f_*/\text{Hz})$ (fourth column), as a function of the number of pulsars N_p , assuming a $T_{\text{obs}} = 10.33$ yr. We assume an EPTA-like (top) or SKA-like (bottom) dataset. The dark (light) shaded area is obtained by varying α_{PL} within the 68% (95%) CL posterior distribution from the current EPTA + NANOGrav posterior, see [40]. Solid (dashed) lines correspond to larger (smaller) α_{PL} . The bars around the band edges are obtained from the standard deviation of ten realizations of the pulsar catalog (varying both location and noises).

scaling $\sim 1/\sqrt{N_p}$, as expected from the combination of N_p independent datasets with spectral parameters mostly constrained by autocorrelation information (see, e.g., [8,87]).

We analyze the scaling behavior of the uncertainties of the following parameters: α_{PL} and n_T (first column), $\log_{10} A_\zeta$ (second column), $\log_{10} \Delta$ (third column), and $\log_{10}(f_*/\text{Hz})$ (fourth column), as a function of the total number of pulsars, N_p . Our analysis assumes the observational period to be fixed to $T_{\text{obs}} = 10.33$ yr. We consider two distinct datasets: an EPTA-like dataset (top panel) and an SKA-like dataset (bottom panel). The shaded regions in the plots indicate the range of variation in the uncertainties due to different choices of the PL spectral amplitude α_{PL} . The dark-shaded region corresponds to the 68% CL, i.e., $\alpha_{\text{PL}} \in [-7.6, -7]$, while the light-shaded region extends to the 95% CL, i.e., $\alpha_{\text{PL}} \in [-7.9, -6.8]$, both derived from the posterior distribution obtained using current EPTA + NANOGrav data, see Sec. IV B 1.

Additionally, we differentiate between two limiting cases of α_{PL} : solid lines represent results obtained for larger values of α_{PL} , while dashed lines correspond to smaller values. As expected, in the leftmost panel, the uncertainties are larger for smaller values of α_{PL} , as it corresponds to weaker injected signals. On the other hand, in the remaining columns, the uncertainties are larger for larger values of α_{PL} , as expected for a larger dominant astrophysical foreground.

To account for statistical variations in the pulsar population, we generate multiple realizations of the pulsar catalog. Specifically, the error bars surrounding the edges of the shaded bands reflect the standard deviation computed over ten independent realizations, where both the spatial

distribution of pulsars and their intrinsic noise properties are varied. This variability is very small for large N_p . Furthermore, we see that it reduces significantly when passing from EPTA to an SKA-like dataset, as the observations enter the signal-dominated regime.

We provide fitting formulas for the precision achieved as a function of the observation time and number of pulsars. The relative uncertainties of the PL parameters scale as

$$\begin{aligned}\sigma_{\alpha_{\text{PL}}} &= [10, 24]\% \left(\frac{N_p}{70}\right)^{-1/2}, \\ \sigma_{n_T} &= [3.2, 5.5]\% \left(\frac{N_p}{70}\right)^{-1/2},\end{aligned}\quad (32)$$

while the SIGW parameter uncertainties are

$$\begin{aligned}\frac{\sigma_{A_\zeta}}{A_\zeta} &= [52, 150]\% \left(\frac{N_p}{70}\right)^{-1/2}, \\ \frac{\sigma_\Delta}{\Delta} &= [48, 150]\% \left(\frac{N_p}{70}\right)^{-1/2}, \\ \frac{\sigma_{f_*}}{f_*} &= [32, 96]\% \left(\frac{N_p}{70}\right)^{-1/2}.\end{aligned}\quad (33)$$

In the above equations, we bracket the results obtained by varying α_{PL} within the 68% CL.

As already hinted in the previous section, the relatively large uncertainties found for this injected SIGW signal, which results in a large uncertainty on the amplitude of the curvature power spectrum, imply relatively weak




FIG. 6. Fisher matrix relative uncertainties on A_ζ for EPTA-like configuration with $T_{\text{obs}} = 16.03$ yr and $N_p = 120$ (solid lines) and SKA-like configuration with $T_{\text{obs}} = 10.33$ yr and $N_p = 200$ (shaded), as a function of $\log_{10} \Delta$ (left panel) and $\log_{10}(f_*/\text{Hz})$ (right panel). When fixed, the parameters are set to $\log_{10}(f_*/\text{Hz}) = -8.45$ (left panel) and $\Delta = 0.5$ (right panel). The injection values for the MCMC analysis are marked with a red cross (star) for present (future) configuration. They coincide in the left panel.

constraining power on the parameter space of relevance for PBH formation. In Sec. V, we will discuss the implications of such constraints for PBH production.

4. Scaling with SIGW model parameters

We further derive future bounds on the curvature power spectrum by varying the SIGW model parameters. In Fig. 6, we show the relative uncertainties on A_ζ varying the injected SIGW parameters, Δ (left panel) and f_* (right panel). We keep the PL parameters fixed to the reference values chosen in Eq. (27). The results for an EPTA-like dataset with $T_{\text{obs}} = 16.03$ yr and $N_p = 120$, similar to the upcoming IPTA data release, are shown with a dashed line. The future SKA-like configuration with $T_{\text{obs}} = 10.33$ yr and $N_p = 200$ is shown with colored shading.

We can identify simple trends in the left plot of Fig. 6. We see that relative uncertainties on A_ζ decrease both for larger amplitude and smaller Δ , up to $\log_{10} \Delta \simeq -0.5$. Because of our choice of f_* , narrower spectra exit from the observable window and thus correspond to larger uncertainties. The EPTA-like case, due to the longer assumed observation time, can reach a larger parameter space, as it can access slightly lower frequencies $\sim 1/T_{\text{obs}}$. The secondary island is induced by the large bump in the signal obtained in the small Δ limit. Interestingly, in this specific setting, the uncertainties are smaller for the EPTA-like configuration. This is simply due to our choice of f_* .

Focusing on the right panel, we see that the best reconstruction is obtained for a peak frequency around $f_* \sim \text{few}/T_{\text{obs}}$, depending on the observation time we consider. For the current NANOGrav duration of 16.03 yr, the peak sensitivity falls at $f_* \simeq 10^{-8.5}$ Hz. The double peak feature reflects the shape of the SIGW spectrum for the assumed narrow width of $\mathcal{P}_\zeta$. The uncertainty quickly grows going away from the sweet spot values either because the signal falls outside the observable band (at low pivot frequencies) or it can only be observed above the pulsar

noise in its tail (at high pivot frequencies). As we already saw in the examples of the previous sections, a subdominant SIGW can be measured with a precision better than $\mathcal{O}(30)\%$ only in a limited region of parameter space.

We validate the results of the FIM estimates by comparing them with MCMC Bayesian inference of simulated frequency domain datasets. When running MCMC analysis, we impose the following priors on the model parameters:

$$\theta \in [\theta_0 - 5\sigma_\theta, \theta_0 + 5\sigma_\theta], \quad (34)$$

where the subscript 0 denotes the injection values, while σ_θ is the uncertainty estimated using the FIM analysis. The SIGW priors are further limited by the following ranges:

$$\begin{aligned} \log_{10} A_\zeta &\in [-3.5, 0], \\ \log_{10} \Delta &\in [-1.5, 0.9], \\ \log_{10}(f_*/\text{Hz}) &\in [-10, -6]. \end{aligned} \quad (35)$$

Moreover, in the MCMC run we also impose a prior that avoids PBH overproduction (see next section), i.e., we impose $f_{\text{PBH}} < 1$, which partially cuts a portion of the parameter space with large A_ζ .

In Fig. 7, we show the posterior distribution on the PL + SIGW analyses for some benchmark points identified in the plane of Fig. 6 by crosses. Notice that, since we are interested in comparing the measurement uncertainties from the MCMC procedure and with the FIM, in this case, we do not generate data by sampling from a Gaussian distribution with zero mean and standard deviation given by the corresponding spectral density. Rather, we fixed the data to match with the ensemble average. Effectively, this corresponds to anchoring the mean of the realization (and, correspondingly, the mean values of all the parameters) to its expectation value, which allows us to focus on the uncertainties. In general, we find very good agreement between the results of the MCMC analysis and the




FIG. 7. Comparison between Fisher matrix (red) and Bayesian MCMC parameter estimation (green) obtained from EPTA-like configuration with $T_{\text{obs}} = 16.03$ yr and $N_p = 120$ (left panel) and SKA-like configuration with $T_{\text{obs}} = 10.33$ yr and $N_p = 200$ (right panel). The signal is composed of a PL as in Eq. (27), and an SIGW with $\log_{10} A_\zeta = -1.7$, $\Delta = 0.5$, $\log_{10}(f_*/\text{Hz}) = -8.45(-8.25)$ for the left (right) plot. The Fisher matrix results excluding the points with $f_{\text{PBH}} > 1$ are plotted in blue. In the 2-d panels, 68%, 95%, and 99% confidence interval are shown. The prior ranges coincide with the boundaries.

uncertainties estimated through the FIM. In the left panel, a slight non-Gaussianity of the green MCMC posterior is found, with small tails of the distribution appearing at low values of α_{PL} , A_ζ , and Δ . This is due to the relatively small SNR of the injected signal. Contrarily, in the right panel, we notice that the green contours appear to be slightly more constraining than the red ones. The origin of this small difference comes from the PBH overproduction prior, which is imposed in the MCMC, but not in the FIM. For this reason, we also show, in blue, the impact of including the same prior in the FIM analysis.⁴ After including the same prior, while this does not introduce major corrections to the posterior, we see the constraints to match nearly perfectly.

We repeat the exercise in Fig. 8, by varying the amplitude of the PL foreground within the $2\text{-}\sigma$ CL. As expected, the uncertainties on A_ζ are anticorrelated with the choice of injected α_{PL} . Overall, we conclude that varying the amplitude of the PL foreground within the 95% CL results in a shift of around half an order of magnitude in the reach of future PTA datasets on A_ζ . This would be significant in terms of constraining power on PBH formation scenarios, as we motivate in the following section.

V. IMPLICATIONS FOR PBHS

In this section, we explore the implications of a detection or upper bound on curvature perturbations for PBH models. Large primordial density perturbations can lead to the

formation of a significant abundance of PBHs (see [88] for a recent review). Let us connect the scale of the perturbations, or equivalently the comoving wave number k , with the Hubble sphere mass M_H as [89]

$$M_H \simeq 10M_\odot \left(\frac{g_*}{106.75} \right)^{1/2} \left(\frac{106.75}{g_{*s}} \right)^{2/3} \left(\frac{10^6 \text{ Mpc}^{-1}}{k/\kappa} \right)^2. \quad (36)$$

The resulting PBH mass after the collapse is related to M_H by an order-unity factor described by the theory of critical collapse and validated through dedicated numerical relativistic simulations [90] (see Ref. [91] for a recent review). Notice that we include the prefactor $\kappa \equiv kr_m$ that relates k to the characteristic perturbation size at Hubble crossing $r_m = 1/aH$ [92–95], which also corresponds to the peak of the compaction function. For definiteness, we will fix it to the value $\kappa = 2$ that corresponds to the case of a LN curvature spectrum with width $\Delta \simeq \mathcal{O}(0.5)$ [96,97]. Directly inspecting the right panel of Fig. 6, we see that PTA observations with $T_{\text{obs}} = \mathcal{O}(10)$ yr have a peak sensitivity for masses around $M \simeq \mathcal{O}(\text{few})M_\odot$, with an extended tail in the subsolar mass range.

The PBH abundance can be computed by evaluating the probability that the compaction function, a measure of the mass excess over the areal radius, is larger than the threshold for collapse, which is defined through dedicated numerical simulations. We follow the general prescription presented in Ref. [98] (see also Refs. [99,100]), where more details can be found. The approach is based on threshold statistics on the compaction function and can include NG

⁴In practice, this is achieved by sampling data from the FIM and rejecting points that are not compatible with $f_{\text{PBH}} < 1$.




FIG. 8. Same as Fig. 6 but assuming α_{PL} to be at the current best fit $+2\sigma$ (top row) and -2σ (bottom row).

corrections to the statistics, although we only assume Gaussian perturbations in this work. We review these computations in the Appendix for completeness. This is also built in an independent module of the *fastPTA* code and can be used to evaluate overproduction bounds when running MCMC Bayesian analysis both on simulated and real GW data.

In this section, we derive the upper bound on the spectral amplitude and compare it to the values required to produce a significant population of PBHs. Assuming no SIGW detection, an agnostic upper bound is obtained by injecting a PL signal with best-fit parameters and performing a multimodel fit which could include a subdominant SIGW signal. The priors we impose on the SIGW parameters are

$$\begin{aligned} \log_{10} A_\zeta &\in [-3.5, 0], \\ \log_{10} \Delta &\in [-1.5, 0.9], \\ \log_{10}(f_*/\text{Hz}) &\in [-10, -6]. \end{aligned} \quad (37)$$

We do not include additional priors from PBH overproduction in this first stage.

In Fig. 9, we derive the upper bound on the power spectral parameters A_ζ and k_* , marginalized over the power spectral width Δ , assuming no SIGW is detected using the MCMC analyses. In the left panel, we show the full MCMC posterior distribution. We see that the PL parameters are well reconstructed, with bounds on α_{PL} and n_T which are competitive with analyses assuming a single PL. A slight

degeneracy between the PL and SIGWs is observed, and a subdominant tail of the posterior explains the injected signal as coming from SIGWs, with $A_\zeta \simeq 0.1$ and $\log_{10}(f_*/\text{Hz}) \simeq -7.2$. This is also observed as a darker contour in the (A_ζ, f_*) plane. For pivot scales away from the sweet-spot frequency, the shape of the SIGW becomes unconstrained, as can be seen by the flat posterior distribution of Δ .

The plots depict a clear upper bound on A_ζ as a function of the pivot scale, marginalized over the spectral width. We focus on the (A_ζ, f_*) plane again in the right panel, where we show marginalized upper bounds at 95% CL for three different datasets, assuming $N_p = 68$ (green), $N_p = 136$ (orange), and $N_p = 272$ (blue) of EPTA-like catalogs with $T_{\text{obs}} = 16.03$ yr. The green line, corresponding to $N_p = 68$, is compatible with the current most stringent upper bound derived from existing NANOGrav data. This line serves as a benchmark for comparison, illustrating the present constraints on parameter uncertainties. The other lines correspond to projections based on future IPTA analyses or subsequent datasets, such as those anticipated from enhanced observations with larger PTA datasets. Slight oscillatory features observed in the plot are only due to the finite sampling of the posterior. These scenarios demonstrate a progressive tightening of the uncertainties, leading to stronger upper bounds. Notice there exists a difference with the result reported in Fig. 2 (left panel) of [43], where a much tighter upper bound was derived from the currently available PTA dataset (to be compared with




FIG. 9. Left panel: posterior distribution obtained performing a Bayesian inference of the simulated power law signal as in Eq. (27) with a negligible SIGW contribution. The dataset simulates a EPTA-like observation campaign with $T_{\text{obs}} = 16.03$ yr, and $N_p = 272$. Right panel: enlargement of the (k_*, A_c) panel, reporting the upper bound on A_c at the 95% CL, marginalizing over the spectral width and assuming $N_p = 68$ (green), $N_p = 136$ (orange), and $N_p = 272$ (blue).

our green line $N_p = 68$). The difference comes from our marginalization over Δ , which was instead fixed to the representative value of $\Delta = 0.5$ in [43].

To facilitate the comparison of the forecasted bounds with PBH abundance, as well as previous literature, we further restrict the width of the PBH spectrum to be $\Delta = 0.5$. In Fig. 10, we show the corresponding upper bound on the spectral amplitude as a function of the pivot




FIG. 10. Same as Fig. 9 (right panel), but fixing $\Delta = 0.5$. Additionally, the bound obtained from SKA-like data with $T_{\text{obs}} = 10.33$ yr and $N_p = 272$ is reported in red. $T_{\text{obs}} = 16.03$ yr is assumed for EPTA-like catalogs with different numbers of pulsars. The band indicates the upper bound at 95% and 99.7% CL. The black solid and dashed lines indicates $f_{\text{PBH}} = 1$ using threshold statistics or peak theory, respectively. In both cases, we assume Gaussian perturbations.

scale. The color code is the same adopted in Fig. 9 (right panel), with the colored band corresponding to 95% and 99.7% CL. We find that, by fixing the width to the choice made in [43], our estimate of the current PTA bound agrees with their derivation from the real NANOGrav data when assuming nearly Gaussian perturbations.⁵ We also estimate an SKA-like scenario with $N_p = 272$ and $T_{\text{obs}} = 10.33$ yr. The bound shifts to a higher pivot frequency due to the shorter observation time we assumed. It is interesting to stress, though, that the bound does not improve significantly moving from EPTA-like to SKA-like noise. This is a consequence of the fact that we enter in the signal-dominated regime with the PL dominating the SGWB in the observed frequency range.

We also superimpose a black solid line to show the amplitude corresponding to the representative value $f_{\text{PBH}} = 1$ computed assuming threshold statistics. In the same plot, the black dashed line indicates the value of the amplitude computed assuming peak theory [101]. The slight reduction observed in both lines around $k \simeq 6 \times 10^6 \text{ Mpc}^{-1}$ is due to the effect of the QCD era in reducing the threshold for collapse [89,102–105]. The difference between these two lines indicates the systematic theoretical uncertainties associated with this computation [106–109].

⁵A minor difference is observed around $k_* \simeq 10^7 \text{ Mpc}^{-1}$, where the actual NANOGrav 15 yr bound is slightly more stringent. As we work with simulated datasets, we are not able to reproduce real data statistical fluctuations or unmodeled systematics. Furthermore, it should be noted that those frequency bins only provide an upper bound on the SGWB amplitude, and the 95% CL upper bound is thus still sensitive to the choice of the lower edge of the prior range for the SGWB amplitude.

While threshold statistics predicts a larger amplitude for fixed abundance and reduces the tension with the SIGW explanation of the current PTA signal, it would also lead to tighter bounds on the PBH production in the range $k \in [4 \times 10^5, 10^7] \text{ Mpc}^{-1}$, which corresponds to masses in the range $M_{\text{PBH}} \sim 0.6 M_H \in [0.28, 200] M_\odot$, right in the range of masses probed by the LVK collaboration [110] and by future next-generation ground-based experiments [111–114]. Conversely, assuming peak theory, one obtains a smaller required amplitude, which exacerbates the tension with a possible SIGW interpretation of current PTA data, but at the same time corresponds to very weak bounds on PBH production, only currently scratching the 3σ significance around $M_{\text{PBH}} \sim 0.6 M_H \in [10, 56] M_\odot$. This theoretical uncertainty should be addressed elsewhere in dedicated studies.

It is striking that, although the dataset is assumed to grow by a factor of 4 in terms of the number of pulsars, the upper bound from SIGW only improves by less than a factor of 2, due to the dominant role of the PL signal in the nHz band. We, therefore, foresee much of the progress in constraining PBH models coming from theoretical advancements in reducing *uncertainties* in the PBH computations, which should resolve the tension between the different methods. Additionally, we do not address the additional *model dependence* introduced by possible non-Gaussianity of curvature perturbations, and refer to [43] for a thorough discussion.

It should be stressed that the conclusions above are subject to the assumption that the dominant PL signal is described by the current best-fit PL parameters (27). The actual future observational bound may also change




FIG. 11. 1σ and 3σ detectability level for A_ζ for a SKA-like configuration with fixed $T_{\text{obs}} = 10.33 \text{ yr}$ and $N_p = 200$. The black horizontal line shows the PBH overproduction ($f_{\text{PBH}} = 1$) amplitude assuming threshold statistics. The colored vertical bands are the 95% CL for α_{PL} as computed from EPTA data (gray), NANOGrav data (green), and EPTA + NANOGrav data (orange), see [40]. The other parameters are fixed as in Fig. 3. The injection values for the MCMC analysis shown in Fig. 7 are marked with a red cross.

depending on the assumptions we make on the true value of α_{PL} . We explore this dependence in Fig. 11, where we compare the FIM uncertainties on the power spectral amplitude to the upper bounds of the PBH abundance while also varying the assumed amplitude of the foreground α_{PL} . In blue, we indicate the upper bound at 1σ CL, i.e., $\Delta A_\zeta / A_\zeta < 1$, and 3σ CL, i.e., $\Delta A_\zeta / A_\zeta < 1/3$. As expected, the upper bound on A_ζ scales with the assumed value of α_{PL} . At the high end, we find $\text{Max}(A_\zeta) \sim \alpha_{\text{PL}}^{1/3}$. On the other hand, the upper bound flattens for sufficiently low values of α_{PL} , where the GW foreground becomes irrelevant and the bound is driven by the pulsar noises. The SIGW parameters not shown are fixed to $\log_{10} \Delta = -0.3$ and $\log_{10}(f_*/\text{Hz}) = -7.8$, as in Fig. 3. Most notably, varying α_{PL} within the current most stringent 95% CL reported by the NANOGrav collaboration (green band) results in a shift of the upper bound on A_ζ by around half an order of magnitude. This range is compatible with the current systematic uncertainties on the computation of the PBH abundance. Therefore, further reduction of the uncertainties expected with future datasets will have to be accompanied by a drastic reduction of the theoretical uncertainties associated with the PBH abundance computation.

VI. CONCLUSIONS

In this work, we have explored the future constraints that PTA observations could impose on SIGWs and the associated primordial curvature power spectrum. We mostly employed an FIM approach, which was validated through specific full Bayesian MCMC runs, and we evaluated the precision with which upcoming PTA datasets—characterized by increased numbers of pulsars, extended observation times, and lower white noise—will be able to probe SIGWs.

Our analysis indicates that, while current PTA data is still plagued by significant uncertainties on the SIGW properties, future PTA observations with larger pulsar networks will drastically improve sensitivity. This will lead to reduced degeneracies between key parameters, such as the amplitude, width, and peak frequency of the SIGW spectrum, ultimately providing stronger constraints on the primordial power spectrum. In case SIGW is the sole responsible for the PTA excess, future PTA datasets will slowly increase precision in the determination of the spectral amplitude and tilt, with uncertainties remaining relatively large due to the necessarily low signal amplitude forced by the PBH overproduction bound. Efforts should be directed towards testing this hypothesis using information beyond the frequency spectrum. Moreover, in the scenario in which SIGW only accounts for a subdominant contribution to current data, we have investigated how future datasets could improve our constraints on PBH formation. While current uncertainties limit precise determinations of the PBH abundance, the enhanced sensitivity of the next

PTA datasets will either impose stringent upper bounds on PBH populations or detect signatures of primordial density fluctuations responsible for their formation. We stressed how as measurement uncertainties narrow down, the theoretical uncertainties on the determination of the PBH abundance could remain a limiting factor, potentially preventing us from drawing strong conclusions.

The potential existence of an astrophysical foreground, i.e., sourced by SMBH mergers, remains a key factor in SIGW detectability. A reduction of the uncertainties on the possible foreground amplitude, i.e., leading to a more precise determination of the GW foreground to be expected, would allow a more precise determination of the future upper bound on SIGWs.

We conclude by mentioning a few caveats underlying our analyses, which could be independently addressed in future work. The present analysis relies solely on the information provided by the spectral shape of different SGWB contributions without identifying any smoking-gun feature distinguishing astrophysical from cosmological sources. While this Bayesian model comparison approach is standard in GW data interpretation, the conclusions drawn remain inherently dependent on the reliability of the assumed models. In this respect, the template employed for the curvature power spectrum is relatively flexible, yet the second-order nature of the induced signal imposes constraints on its shape. A more robust modeling of the astrophysical sources would refine these constraints and enhance their significance. Moreover, an explicit model comparison incorporating a first-principles computation of $\mathcal{P}_\zeta$ within an early-universe framework would yield even more stringent bounds, as discussed in [71].

In our study, we assumed that the dominant astrophysical foreground follows a power-law spectrum. Future observations with greater sensitivity might reveal deviations from this assumption, manifesting as spectral oscillations absent from the SIGW signal. Such deviations could improve the differentiation between astrophysical and cosmological components. Additionally, we conservatively considered the astrophysical SGWB to be irreducible through the subtraction of individually resolved sources. However, in practice, improved source subtraction techniques might reduce its contribution, thereby further clarifying the potential SIGW signature.

Another distinguishing aspect is that the astrophysical foreground exhibits significantly stronger anisotropies, particularly at high frequencies, which are instead not present in the cosmological SIGW at large scales [115]. Once detected, incorporating these anisotropies into the analysis could improve the separation of the subdominant cosmological SGWB component. The study of these anisotropic effects is left for future research.

To extend the scope of our analysis, we envision the use of our forecast code in conjunction with SIGWAY. This

combination would enable the exploration of more model-independent spectral parametrizations and facilitate the direct modeling of the inflationary dynamics responsible for large curvature perturbations. Additionally, one should explore more in detail deviations from the perfect radiation assumption adopted here, as would QCD thermal history unavoidably affect the PTA scale [44,46,116], as well as possible new physics beyond the standard model. Finally, it would be interesting to investigate the effect of accounting for a more realistic damping of curvature modes, which we neglect in this work, as discussed in [117], to test whether it may be measured with future PTA datasets. This would particularly affect the tail of the signal and the resonant peak. We leave this for future work.

It would also be interesting to complement the constraints on the SIGW scenarios with information derived through additional observations beyond the PTA dataset. One possibility is to include bounds from CMB spectral distortions (see, e.g., [118,119]). It should be stressed that the scales of relevance for PTA observations fall in the range between $k \sim 10^6 \div 10^8 \text{ Mpc}^{-1}$, while CMB spectral distortions can only reach up to around $k \lesssim 10^4 \text{ Mpc}^{-1}$. Therefore, *a priori*, there is no direct connection between the two observations. However, within an explicit scenario where the power spectrum of curvature perturbations is enhanced at small scales, the spectral growth required to deviate from the small CMB amplitude would cross the scales relevant for distortions. Thus, models featuring mild growth would lead to signatures at both scales.

ACKNOWLEDGMENTS

We thank Valerie Domcke, Jonas El Gammal, Antonio Iovino, Gabriele Perna, Nataliya Poryako, Antonio Riotto, and the EPTA GW interpretation group for interesting discussions. C. C. thanks the TH department at CERN for the kind hospitality during the beginning of this project. M. P. acknowledges the hospitality of Imperial College London, which provided office space during some parts of this project. C. C. acknowledges support from the Istituto Nazionale di Fisica Nucleare (INFN) through the Commissione Scientifica Nazionale 4 (CSN4) Iniziativa Specifica “Quantum Fields in Gravity, Cosmology and Black Holes” (FLAG).

DATA AVAILABILITY

No data were created or analyzed in this study.

APPENDIX: COMPUTING THE PBH ABUNDANCE

One can compute the PBH abundance by integrating the mass fraction $\beta(M_H)$ produced when perturbations re-enter the Hubble sphere at all relevant scales, including the appropriate redshift scaling. This can be computed as [120]

$$f_{\text{PBH}} = \frac{1}{\Omega_{\text{DM}}} \int_{M_{H,\min}}^{M_{H,\max}} \frac{dM_H}{M_H} \left(\frac{1}{M_H} \right)^{\frac{1}{2}} \times \left(\frac{g_*}{106.75} \right)^{\frac{3}{4}} \left(\frac{g_{*S}}{106.75} \right)^{-1} \left(\frac{\beta(M_H)}{7.9 \times 10^{-10}} \right), \quad (\text{A1})$$

where we introduced the fraction of PBHs compared to the DM budget $f_{\text{PBH}} \equiv \Omega_{\text{PBH}}/\Omega_{\text{DM}}$. Furthermore, the effective degrees of freedom should be evaluated at the corresponding epoch, identifying the Hubble mass with the temperature using

$$M_H(T_k) = 4.8 \times 10^{-2} M_{\odot} \left(\frac{g_*}{106.75} \right)^{-\frac{1}{2}} \left(\frac{T_k}{\text{GeV}} \right)^{-2}. \quad (\text{A2})$$

Given Eq. (A1), computing the abundance boils down to evaluating the probability a given Hubble patch collapses to form a PBH. We will follow the general prescription presented in Ref. [98] (see also Refs. [99,100]), where more details can be found. The approach is based on threshold statistics on the compaction function, a fundamental variable defined as twice the local mass $\delta M(r, t)$ excess over the areal radius $R(r, t)$,

$$\mathcal{C}(r, t) \equiv \frac{2\delta M(r, t)}{R(r, t)} = \frac{2}{R(r, t)} \int_{S_R^2} d^3\vec{x} \rho_b(t) \delta(\vec{x}, t). \quad (\text{A3})$$

Assuming spherical symmetry and adopting the gradient expansion approximation, on the super-Hubble scale the relation between the density contrast δ and the curvature perturbation ζ is nonlinear [93,121]:

$$\delta(r, t) = -\frac{2\Phi}{3} \left(\frac{1}{aH} \right)^2 e^{-2\zeta(r)} \left[\zeta''(r) + \frac{2}{r} \zeta'(r) + \frac{1}{2} \zeta'(r)^2 \right]. \quad (\text{A4})$$

The parameter Φ is introduced to keep track of how the equation of state of the Universe changes due to the thermal history, which is particularly relevant for the formation of stellar mass PBH across the QCD phase transition.

As the relation between the density contrast and curvature perturbations is nonlinear, the above equation unavoidably introduces a certain amount of NGs in the probability density function (PDF) of δ , and hence on that of $\mathcal{C}$. Integrating Eq. (A4) over the radial coordinate, one finds

$$\mathcal{C}(r) = -2\Phi r \zeta'(r) \left[1 + \frac{r}{2} \zeta'(r) \right] = \mathcal{C}_1(r) - \frac{1}{4\Phi} \mathcal{C}_1(r)^2, \quad (\text{A5})$$

where $\mathcal{C}_1(r) \equiv -2\Phi r \zeta'(r)$ is defined as the linear component of the compaction function. In full generality, the latter can be recast in terms of its Gaussian counterpart as follows:

$$\mathcal{C}_1(r) = \mathcal{C}_G(r) \frac{dF}{d\zeta_G}, \quad \mathcal{C}_G(r) := -2\Phi r \zeta'_G(r), \quad (\text{A6})$$

where F encodes the relation between ζ and the Gaussian component ζ_G .

We stress that $\mathcal{C}(r)$ depends both on ζ_G and $\mathcal{C}_G$ which are, by definition, Gaussian distributed. Therefore, their joint PDF can be computed as

$$P_G(\mathcal{C}_G, \zeta_G) = \frac{1}{2\pi\sigma_c\sigma_r\sqrt{1-\gamma_{cr}^2}} \exp\left(-\frac{\zeta_G^2}{2\sigma_r^2}\right) \times \exp\left[-\frac{1}{2(1-\gamma_{cr}^2)} \left(\frac{\mathcal{C}_G}{\sigma_c} - \frac{\gamma_{cr}\zeta_G}{\sigma_r}\right)^2\right], \quad (\text{A7})$$

where $\gamma_{cr} \equiv \sigma_{cr}^2/\sigma_c\sigma_r$ and the correlators are given by

$$\begin{aligned} \sigma_c^2 &= \frac{4\Phi^2}{9} \int_0^\infty \frac{dk}{k} (kr_m)^4 W^2(k, r_m) P_\zeta^T(k), \\ \sigma_{cr}^2 &= \frac{2\Phi}{3} \int_0^\infty \frac{dk}{k} (kr_m)^2 W(k, r_m) W_s(k, r_m) P_\zeta^T(k), \\ \sigma_r^2 &= \int_0^\infty \frac{dk}{k} W_s^2(k, r_m) P_\zeta^T(k), \end{aligned} \quad (\text{A8})$$

with $P_\zeta^T(k) \equiv T^2(k, r_m) P_\zeta(k)$, $W_s(k, r) = \sin(kr)/kr$ while $W(k, R)$ and $T(k, \tau)$ are given by

$$\begin{aligned} T(k, \tau) &= 3 \left[\frac{\sin(k\tau/\sqrt{3}) - (k\tau/\sqrt{3}) \cos(k\tau/\sqrt{3})}{(k\tau/\sqrt{3})^3} \right], \\ W(k, R) &= 3 \left[\frac{\sin(kR) - kR \cos(kR)}{(kR)^3} \right]. \end{aligned} \quad (\text{A9})$$

The transfer function $T(k, \tau)$ we adopt in this section is derived assuming a perfect radiation fluid. Furthermore, we have selected the Hubble crossing epoch by fixing the conformal time $\tau = 1/(aH) = r_m$, using the Hubble crossing condition $aHr_m = 1$.

Finally, $\beta(M_H)$ entering in Eq. (A1), which is identified with the mass fraction derived from the collapse of a single mode is

$$\beta(M_H) = \int_{\mathcal{D}} \mathcal{K}(\mathcal{C} - \mathcal{C}_{\text{th}})^{\gamma} P_G(\mathcal{C}_G, \zeta_G) d\mathcal{C}_G d\zeta_G, \quad (\text{A10})$$

with the collapse condition being expressed as

$$\mathcal{D} = \{\mathcal{C}(\mathcal{C}_G, \zeta_G) > \mathcal{C}_{\text{th}} \wedge \mathcal{C}_1(\mathcal{C}_G, \zeta_G) < 2\Phi\}. \quad (\text{A11})$$

The first condition requires overthreshold perturbations, while the second allows us only to consider type I PBH formation.

Throughout this work, we have assumed Gaussian curvature perturbations. In this case, $F \equiv 1$ and one can marginalize over ζ_G in Eq. (A7), being the compaction function independent of ζ_G . Therefore, (A7) simply

becomes a Gaussian distribution with variance σ_c^2 . Notice that nonlinearities induced by general relativity are still accounted for through the full nonlinear compaction function entering the threshold condition in Eq. (A10).

-
- [1] R. w. Hellings and G. s. Downs, *Astrophys. J. Lett.* **265**, L39 (1983).
 - [2] G. Agazie *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **951**, L8 (2023).
 - [3] J. Antoniadis *et al.* (EPTA Collaboration), *Astron. Astrophys.* **678**, A50 (2023).
 - [4] D. Reardon *et al.* (PPTA Collaboration), *Astrophys. J. Lett.* **951**, L6 (2023).
 - [5] H. Xu *et al.* (CPTA Collaboration), *Res. Astron. Astrophys.* **23**, 075024 (2023).
 - [6] X. Siemens, J. Ellis, F. Jenet, and J. D. Romano, *Classical Quantum Gravity* **30**, 224015 (2013).
 - [7] N. S. Pol *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **911**, L34 (2021).
 - [8] S. Babak, M. Falxa, G. Franciolini, and M. Pieroni, *Phys. Rev. D* **110**, 063022 (2024).
 - [9] M. T. Miles *et al.*, *Mon. Not. R. Astron. Soc.* **536**, 1489 (2024).
 - [10] J. Antoniadis *et al.* (EPTA, InPTA Collaborations), *Astron. Astrophys.* **685**, A94 (2024).
 - [11] G. Agazie *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **952**, L37 (2023).
 - [12] A. Sesana, A. Vecchio, and C. N. Colacino, *Mon. Not. R. Astron. Soc.* **390**, 192 (2008).
 - [13] B. Kocsis and A. Sesana, *Mon. Not. R. Astron. Soc.* **411**, 1467 (2011).
 - [14] J. Ellis, M. Fairbairn, G. Hütsi, M. Raidal, J. Urrutia, V. Vaskonen, and H. Veermäe, *Astron. Astrophys.* **676**, A38 (2023).
 - [15] C. Caprini and D. G. Figueroa, *Classical Quantum Gravity* **35**, 163001 (2018).
 - [16] K. Tomita, *Prog. Theor. Phys.* **54**, 730 (1975).
 - [17] S. Matarrese, O. Pantano, and D. Saez, *Phys. Rev. D* **47**, 1311 (1993).
 - [18] S. Matarrese, O. Pantano, and D. Saez, *Phys. Rev. Lett.* **72**, 320 (1994).
 - [19] S. Matarrese, S. Mollerach, and M. Bruni, *Phys. Rev. D* **58**, 043504 (1998).
 - [20] V. Acquaviva, N. Bartolo, S. Matarrese, and A. Riotto, *Nucl. Phys. B* **667**, 119 (2003).
 - [21] S. Mollerach, D. Harari, and S. Matarrese, *Phys. Rev. D* **69**, 063002 (2004).
 - [22] C. Carbone and S. Matarrese, *Phys. Rev. D* **71**, 043508 (2005).
 - [23] K. N. Ananda, C. Clarkson, and D. Wands, *Phys. Rev. D* **75**, 123518 (2007).
 - [24] D. Baumann, P. J. Steinhardt, K. Takahashi, and K. Ichiki, *Phys. Rev. D* **76**, 084019 (2007).
 - [25] G. Domènech, *Universe* **7**, 398 (2021).
 - [26] Y. B. Zel'dovich and I. D. Novikov, *Sov. Astron. AJ (Engl. Transl.)*, **10**, 602 (1967).
 - [27] S. W. Hawking, *Mon. Not. R. Astron. Soc.* **152**, 75 (1971).
 - [28] B. J. Carr and S. W. Hawking, *Mon. Not. R. Astron. Soc.* **168**, 399 (1974).
 - [29] B. J. Carr, *Astrophys. J.* **201**, 1 (1975).
 - [30] G. F. Chapline, *Nature (London)* **253**, 251 (1975).
 - [31] *Primordial Black Holes*, edited by C. Byrnes, G. Franciolini, T. Harada, P. Pani, and M. Sasaki, Springer Series in Astrophysics and Cosmology (Springer, New York, 2025).
 - [32] B. Carr, K. Kohri, Y. Sendouda, and J. Yokoyama, *Rep. Prog. Phys.* **84**, 116902 (2021).
 - [33] R. Saito and J. Yokoyama, *Phys. Rev. Lett.* **102**, 161101 (2009); **107**, 069901(E) (2011).
 - [34] E. V. Bugaev and P. A. Klimai, *JETP Lett.* **91**, 1 (2010).
 - [35] E. Bugaev and P. Klimai, *Phys. Rev. D* **83**, 083521 (2011).
 - [36] C. T. Byrnes, P. S. Cole, and S. P. Patil, *J. Cosmol. Astropart. Phys.* **06** (2019) 028.
 - [37] A. D. Gow, C. T. Byrnes, P. S. Cole, and S. Young, *J. Cosmol. Astropart. Phys.* **02** (2021) 002.
 - [38] A. Afzal *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **951**, L11 (2023); **971**, L27(E) (2024).
 - [39] E. Madge, E. Morgante, C. Puchades-Ibáñez, N. Ramberg, W. Ratzinger, S. Schenk, and P. Schwaller, *J. High Energy Phys.* **10** (2023) 171.
 - [40] D. G. Figueroa, M. Pieroni, A. Ricciardone, and P. Simakachorn, *Phys. Rev. Lett.* **132**, 171002 (2024).
 - [41] J. Ellis, M. Fairbairn, G. Franciolini, G. Hütsi, A. Iovino, M. Lewicki, M. Raidal, J. Urrutia, V. Vaskonen, and H. Veermäe, *Phys. Rev. D* **109**, 023522 (2024).
 - [42] V. Dandoy, V. Domcke, and F. Rompineve, *SciPost Phys. Core* **6**, 060 (2023).
 - [43] A. J. Iovino, G. Perna, A. Riotto, and H. Veermäe, *J. Cosmol. Astropart. Phys.* **10** (2024) 050.
 - [44] D. J. Schwarz, *Mod. Phys. Lett. A* **13**, 2771 (1998).
 - [45] K. T. Abe, Y. Tada, and I. Ueda, *J. Cosmol. Astropart. Phys.* **06** (2021) 048.
 - [46] G. Franciolini, D. Racco, and F. Rompineve, *Phys. Rev. Lett.* **132**, 081001 (2024).
 - [47] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020); **652**, C4(E) (2021).
 - [48] K. Kohri and T. Terada, *Phys. Rev. D* **97**, 123532 (2018).
 - [49] J. R. Espinosa, D. Racco, and A. Riotto, *J. Cosmol. Astropart. Phys.* **09** (2018) 012.
 - [50] F. Hajkarim and J. Schaffner-Bielich, *Phys. Rev. D* **101**, 043522 (2020).
 - [51] K. Inomata and T. Terada, *Phys. Rev. D* **101**, 023523 (2020).

- [52] V. De Luca, G. Franciolini, A. Kehagias, and A. Riotto, *J. Cosmol. Astropart. Phys.* **03** (2020) 014.
- [53] C. Yuan, Z.-C. Chen, and Q.-G. Huang, *Phys. Rev. D* **101**, 063018 (2020).
- [54] G. Domènech and M. Sasaki, *Phys. Rev. D* **103**, 063531 (2021).
- [55] P. G. Ferreira and M. Joyce, *Phys. Rev. D* **58**, 023503 (1998).
- [56] C. Pallis, *Nucl. Phys.* **B751**, 129 (2006).
- [57] K. Redmond, A. Trezza, and A. L. Erickcek, *Phys. Rev. D* **98**, 063504 (2018).
- [58] R. T. Co, D. Dunsy, N. Fernandez, A. Ghalsasi, L. J. Hall, K. Harigaya, and J. Shelton, *J. High Energy Phys.* **09** (2022) 116.
- [59] Y. Gouttenoire, G. Servant, and P. Simakachorn, *arXiv*: 2111.01150.
- [60] C.-F. Chang and Y. Cui, *J. High Energy Phys.* **03** (2022) 114.
- [61] I. Dalianis and G. Tringas, *Phys. Rev. D* **100**, 083512 (2019).
- [62] S. Bhattacharya, S. Mohanty, and P. Parashari, *Phys. Rev. D* **102**, 043522 (2020).
- [63] S. Bhattacharya, S. Mohanty, and P. Parashari, *Phys. Rev. D* **103**, 063532 (2021).
- [64] A. Ireland, S. Profumo, and J. Scharnhorst, *Phys. Rev. D* **107**, 104021 (2023).
- [65] S. Bhattacharya, *Galaxies* **11**, 35 (2023).
- [66] A. Ghoshal, Y. Gouttenoire, L. Heurtier, and P. Simakachorn, *J. High Energy Phys.* **08** (2023) 196.
- [67] G. Domènech, *Int. J. Mod. Phys. D* **29**, 2050028 (2020).
- [68] G. Domènech and A. Ganz, *J. Cosmol. Astropart. Phys.* **01** (2025) 020.
- [69] U. Kumar, U. Thattarampilly, and P. Chaturvedi, *arXiv*: 2410.12931.
- [70] O. Özsoy and G. Tasinato, *Universe* **9**, 203 (2023).
- [71] J. E. Gammal *et al.* (LISA Cosmology Working Group), *arXiv*:2501.11320.
- [72] S. Pi and M. Sasaki, *J. Cosmol. Astropart. Phys.* **09** (2020) 037.
- [73] P. F. Depta, V. Domcke, G. Franciolini, and M. Pieroni, *Phys. Rev. D* **111**, 083039 (2025).
- [74] J. S. Hazboun, J. D. Romano, and T. L. Smith, *Phys. Rev. D* **100**, 104028 (2019).
- [75] J. R. Bond, A. H. Jaffe, and L. Knox, *Phys. Rev. D* **57**, 2117 (1998).
- [76] C. R. Contaldi, M. Pieroni, A. I. Renzini, G. Cusin, N. Karnesis, M. Peloso, A. Ricciardone, and G. Tasinato, *Phys. Rev. D* **102**, 043502 (2020).
- [77] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [78] G. Franciolini, A. Iovino, Jr., V. Vaskonen, and H. Veermäe, *Phys. Rev. Lett.* **131**, 201401 (2023).
- [79] M. Braglia *et al.* (LISA Cosmology Working Group), *J. Cosmol. Astropart. Phys.* **11** (2024) 032.
- [80] C. Caprini, R. Durrer, T. Konstandin, and G. Servant, *Phys. Rev. D* **79**, 083519 (2009).
- [81] R.-G. Cai, S. Pi, and M. Sasaki, *Phys. Rev. D* **102**, 083528 (2020).
- [82] R. Allahverdi *et al.*, *Open J. Astrophys.* **4**, 52 (2021).
- [83] A. Hook, G. Marques-Tavares, and D. Racco, *J. High Energy Phys.* **02** (2021) 117.
- [84] G. Agazie *et al.*, *Astrophys. J.* **978**, 31 (2025).
- [85] G. Agazie *et al.* (International Pulsar Timing Array Collaboration), *Astrophys. J.* **966**, 105 (2024).
- [86] A. Afzal *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **951**, L11 (2023).
- [87] J. Samsing, M. MacLeod, and E. Ramirez-Ruiz, *Astrophys. J.* **784**, 71 (2014).
- [88] E. Bagui *et al.* (LISA Cosmology Working Group), *Living Rev. Relativity* **28**, 1 (2025).
- [89] G. Franciolini, I. Musco, P. Pani, and A. Urbano, *Phys. Rev. D* **106**, 123526 (2022).
- [90] I. Musco, J. C. Miller, and A. G. Polnarev, *Classical Quantum Gravity* **26**, 235001 (2009).
- [91] A. Escrivà, *Universe* **8**, 66 (2022).
- [92] C. Germani and I. Musco, *Phys. Rev. Lett.* **122**, 141302 (2019).
- [93] I. Musco, *Phys. Rev. D* **100**, 123524 (2019).
- [94] A. Escrivà, C. Germani, and R. K. Sheth, *Phys. Rev. D* **101**, 044022 (2020).
- [95] S. Young, *Int. J. Mod. Phys. D* **29**, 2030002 (2019).
- [96] I. Musco, V. De Luca, G. Franciolini, and A. Riotto, *Phys. Rev. D* **103**, 063538 (2021).
- [97] M. S. Delos and G. Franciolini, *Phys. Rev. D* **107**, 083505 (2023).
- [98] G. Ferrante, G. Franciolini, A. Iovino, Jr., and A. Urbano, *Phys. Rev. D* **107**, 043520 (2023).
- [99] S. Young, *J. Cosmol. Astropart. Phys.* **05** (2022) 037.
- [100] A. D. Gow, H. Assadullahi, J. H. P. Jackson, K. Koyama, V. Vennin, and D. Wands, *Europhys. Lett.* **142**, 49001 (2023).
- [101] J. M. Bardeen, J. R. Bond, N. Kaiser, and A. S. Szalay, *Astrophys. J.* **304**, 15 (1986).
- [102] K. Jedamzik, *Phys. Rev. D* **55**, R5871 (1997).
- [103] C. T. Byrnes, M. Hindmarsh, S. Young, and M. R. S. Hawkins, *J. Cosmol. Astropart. Phys.* **08** (2018) 041.
- [104] A. Escrivà, E. Bagui, and S. Clesse, *J. Cosmol. Astropart. Phys.* **05** (2023) 004.
- [105] I. Musco, K. Jedamzik, and S. Young, *Phys. Rev. D* **109**, 083506 (2024).
- [106] A. M. Green, A. R. Liddle, K. A. Malik, and M. Sasaki, *Phys. Rev. D* **70**, 041502 (2004).
- [107] S. Young, C. T. Byrnes, and M. Sasaki, *J. Cosmol. Astropart. Phys.* **07** (2014) 045.
- [108] V. De Luca, G. Franciolini, A. Kehagias, M. Peloso, A. Riotto, and C. Ünal, *J. Cosmol. Astropart. Phys.* **07** (2019) 048.
- [109] L. Frosina and A. Urbano, *Phys. Rev. D* **108**, 103544 (2023).
- [110] G. Franciolini, R. Cotesta, N. Loutrel, E. Berti, P. Pani, and A. Riotto, *Phys. Rev. D* **105**, 063510 (2022).
- [111] V. De Luca, G. Franciolini, P. Pani, and A. Riotto, *J. Cosmol. Astropart. Phys.* **11** (2021) 039.
- [112] O. Pujolas, V. Vaskonen, and H. Veermäe, *Phys. Rev. D* **104**, 083521 (2021).
- [113] G. Franciolini, F. Iacovelli, M. Mancarella, M. Maggiore, P. Pani, and A. Riotto, *Phys. Rev. D* **108**, 043506 (2023).

- [114] M. Branchesi *et al.*, *J. Cosmol. Astropart. Phys.* **07** (2023) 068.
- [115] N. Bartolo, D. Bertacca, V. De Luca, G. Franciolini, S. Matarrese, M. Peloso, A. Ricciardone, A. Riotto, and G. Tasinato, *J. Cosmol. Astropart. Phys.* **02** (2020) 028.
- [116] Y. Watanabe and E. Komatsu, *Phys. Rev. D* **73**, 123515 (2006).
- [117] G. Domènech and J. Chluba, [arXiv:2503.13670](#).
- [118] B. Cyr, T. Kite, J. Chluba, J. C. Hill, D. Jeong, S. K. Acharya, B. Bolliet, and S. P. Patil, *Mon. Not. R. Astron. Soc.* **528**, 883 (2024).
- [119] M. Tagliazucchi, M. Braglia, F. Finelli, and M. Pieroni, *Phys. Rev. D* **111**, L021305 (2025).
- [120] M. Sasaki, *Prog. Theor. Phys.* **76**, 1036 (1986).
- [121] T. Harada, C.-M. Yoo, T. Nakama, and Y. Koga, *Phys. Rev. D* **91**, 084057 (2015).