

Unbiased estimation of gravitational-wave anisotropies from noisy data

Nikolaos Kouvatsos^{1,*} Alexander C. Jenkins^{2,†} Arianna I. Renzini^{3,4,5,‡}
Joseph D. Romano^{6,§} and Mairi Sakellariadou^{1,||}

¹Theoretical Particle Physics and Cosmology Group, Physics Department, King's College London, University of London, Strand, London WC2R 2LS, United Kingdom

²Department of Physics and Astronomy, University College London, London WC1E 6BT, United Kingdom

³LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA

⁴Dipartimento di Fisica "G. Occhialini", Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, 20126 Milano, Italy

⁵INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, 20126 Milano, Italy

⁶Department of Physics and Astronomy, University of Texas Rio Grande Valley, One West University Boulevard, Brownsville, Texas 78520, USA

 (Received 18 December 2023; accepted 23 April 2024; published 22 May 2024)

One of the most exciting targets of current and future gravitational-wave observations is the angular power spectrum of the astrophysical GW background. This cumulative signal encodes information about the large-scale structure of the Universe, as well as the formation and evolution of compact binaries throughout cosmic time. However, the finite rate of compact binary mergers gives rise to *temporal shot noise*, which introduces a significant bias in measurements of the angular power spectrum if not explicitly accounted for. Previous work showed that this bias can be removed by cross-correlating GW sky maps constructed from different observing times. However, this work considered an idealized measurement scenario, ignoring detector specifics and in particular noise contributions. Here we extend this temporal cross-correlation method to account for these difficulties, allowing us to implement the first unbiased anisotropic search pipeline for LIGO-Virgo-KAGRA data. In doing so, we show that the existing pipeline is biased *even in the absence of shot noise*, due to previously neglected subleading contributions to the noise covariance. We apply our pipeline to mock LIGO data, and find that our improved analysis will be crucial for stochastic searches from the current observing run (O4) onwards.

DOI: [10.1103/PhysRevD.109.103535](https://doi.org/10.1103/PhysRevD.109.103535)

I. INTRODUCTION

During its first three observing runs, the LIGO-Virgo-KAGRA Collaboration detected dozens of gravitational-wave (GW) signals from binary black hole coalescences [1–3], with the fourth run currently ongoing. This number is expected to increase to several thousand by the end of the fifth observing run [4]. As a result, interest is now shifting to other, undetected observables, like the GW background (GWB)—a persistent signal made up of numerous superimposed GW sources throughout the history of the

Universe [5,6]. With the recent strong evidence for a GWB in the nanohertz band provided by pulsar timing arrays [7–9], we can expect exciting developments in the field within the next few years.

While the GWB can originate from many different early-universe cosmological mechanisms [10], the astrophysical GW background (AGWB) [11] arising from the combined emission of compact binary coalescences (CBCs) is expected to be the dominant component. In the past years, significant research has been devoted to probing the anisotropies in the AGWB, as these could provide us with valuable information with regards to the large-scale structure (LSS) of the Universe and, in particular, galaxy clustering [12–18].

Much like other cosmological observables such as the cosmic microwave background (CMB), the anisotropies in the AGWB are typically probed by means of their angular power spectrum, C_ℓ . A significant challenge in the case of the AGWB is that the finite rate of CBCs leads to temporal shot noise, which can significantly bias inferences of the angular power spectrum of the AGWB [19–21]. It is essential that this bias be properly accounted for in

*nikolaos.kouvatsos@kcl.ac.uk

†alex.jenkins@ucl.ac.uk

‡arenzini@caltech.edu

§joseph.romano@utrgv.edu

||mairi.sakellariadou@kcl.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

order to calculate the true angular power spectrum. Reference [22] introduced a technique to remove this shot noise bias by performing a cross-correlation of GW sky maps produced from nonoverlapping epochs within the full dataset. However, this work did not consider limitations related to GW detector noise and interferometer orientation. In this paper, we show both in theory and in practice how the optimal C_ℓ estimator first introduced in Ref. [22] does indeed compute the correct angular power spectrum even after one accounts for the aforementioned effects, unlike the estimator currently used by both the LVK Collaboration and other authors [23–26] to probe the angular power spectrum of the AGWB [27].

As a byproduct of this analysis, we show that the existing search pipeline is biased *even in the absence of shot noise*, due to subleading contributions to the covariance matrix. These contributions are sourced by the GWB signal itself, which is much weaker than the noise power in the detectors, and is thus usually neglected. However, as we show explicitly below, these terms can be significantly larger than the angular power spectrum one is searching for, and therefore bias the C_ℓ estimator by more than 100%. The method we describe here is robust to these errors, and is thus the first unbiased search pipeline for GWB anisotropies with LVK data.

The paper is organized as follows. In Sec. II, we show how the current C_ℓ estimator is biased with contributions from detector noise, shot noise and the signal itself, while the optimal estimator is by construction unbiased. Section III describes the anisotropic analysis we perform on mock data to validate our theoretical predictions, for detector noise sampled from O3 and LIGO A + sensitivity. We also make predictions for the currently ongoing O4 observing run. We conclude in Sec. IV by discussing the importance of our findings for current and increased detector sensitivity, and providing possible future extensions of our work.

II. UNBIASED SEARCH METHOD

Our data vector \hat{P} (referred to as the cross-spectral density (CSD) in relevant literature [28]) is formed from cross-correlations between the strain measured in the different detectors in our network,

$$\hat{P}_{IJ}(f, t) = \frac{2}{\tau} \hat{s}_I(f, t) \hat{s}_J^*(f, t), \quad (1)$$

with $I \neq J$ labeling these detectors, and τ the length of the time segments used to measure this cross-power (typically 192 s in anisotropic LVK analyses, with 50% of the data overlapping [26]). Each cross-correlation is formed from data at a particular frequency f (which can be positive or negative) and time segment t . We assume that the strain data, \hat{s} , are Gaussian, so that the first and second moments of \hat{P} are

$$\langle \hat{P}_{IJ}(f, t) \rangle_N = \sum_{\ell m} \tilde{\gamma}_{IJ}^{\ell m}(f, t) \Omega_{\ell m}(t), \quad (2)$$

$$\begin{aligned} \text{Cov}[\hat{P}_{IJ}(f, t), \hat{P}_{I'J'}(f', t')]_N \\ = \delta_{ff'} \delta_{tt'} \left(\delta_{II'} N_I(f, t) + \sum_{\ell m} \tilde{\gamma}_{II'}^{\ell m}(f, t) \Omega_{\ell m}(t) \right) \\ \times \left(\delta_{JJ'} N_J(f, t) + \sum_{\ell' m'} \tilde{\gamma}_{JJ'}^{\ell' m'*}(f, t) \Omega_{\ell' m'}^*(t) \right) \\ + \delta_{f, -f'} \delta_{t, t'} \left(\delta_{II'} N_I(f, t) + \sum_{\ell m} \tilde{\gamma}_{II'}^{\ell m}(f, t) \Omega_{\ell m}(t) \right) \\ \times \left(\delta_{I'J'} N_{J'}(f, t) + \sum_{\ell' m'} \tilde{\gamma}_{I'J'}^{\ell' m'*}(f, t) \Omega_{\ell' m'}^*(t) \right), \quad (3) \end{aligned}$$

where $\langle \dots \rangle_N$ denotes an average over noise realizations, $N_I(f, t)$ is the noise power spectral density PSD in detector I , and $\tilde{\gamma}_{IJ}^{\ell m}(f, t)$ are the spherical harmonic components of the overlap reduction function (ORF) of detector baseline IJ . We have absorbed a frequency-dependent factor into the definition of these compared to the definition usually found in the literature, as this simplifies many of the expressions below,

$$\tilde{\gamma}_{IJ}^{\ell m} \equiv \frac{3H_0^2}{2\pi^2 f_{\text{ref}}^3} (f/f_{\text{ref}})^{\alpha-3} \gamma_{IJ}^{\ell m}, \quad (4)$$

where α is the assumed spectral index of the astrophysical/cosmological signal and $f_{\text{ref}} = 25$ Hz. This definition factors out the frequency dependence of the signal, so that the GWB spherical harmonics $\Omega_{\ell m}$ depend only on time, due to different realizations of time-dependent shot noise. (Here we assume for simplicity that the signal is a power law in frequency, as expected for an astrophysical background from quasi-circular binary inspirals well before merger; incorporating more detailed modeling of the frequency dependence of the signal may allow for more effective removal of the shot noise [29].)

We estimate the spherical harmonic components of the GWB as [27]

$$\hat{\Omega}_{\ell m} = \sum_{\ell' m'} (\Gamma^{-1})_{\ell m, \ell' m'} \hat{X}_{\ell' m'}, \quad (5)$$

where the *dirty map* $\hat{X}_{\ell m}$ and *Fisher matrix* $\Gamma_{\ell m, \ell' m'}$ are defined as

$$\begin{aligned} \hat{X}_{\ell m} &= (\tilde{\gamma}^{\ell m} | \hat{P}), \\ \Gamma_{\ell m, \ell' m'} &= (\tilde{\gamma}^{\ell m} | \tilde{\gamma}^{\ell' m'}). \end{aligned} \quad (6)$$

Here we have introduced a noise-weighted inner product over time-, frequency-, and baseline-dependent complex functions,

$$(A|B) = \sum_{J>I} \sum_{f,t} \frac{A_{IJ}^*(f,t)B_{IJ}(f,t)}{\bar{P}_I(f,t)\bar{P}_J(f,t)}. \quad (7)$$

Note that there is a subtlety associated with the quantity we are using to do the noise weighting, \bar{P}_I . To leading order this is just the noise PSD N_I , but since it is measured from the strain data, it will inevitably include some contribution from the GWB signal. In principle, since this is a data-dependent quantity, we should treat it as a random variable. However, in practice it is estimated independently of the data segment under consideration (e.g. using neighboring segments), so that it can be effectively considered as a deterministic quantity. It is still unclear exactly what value this deterministic noise estimate will take, particularly when we consider corrections of order $\Omega_{\ell m}$. As we show below, this is a deficiency of the existing search method, which can be rectified in a way that also minimizes the impact of shot noise.

To see this, it is necessary first to write down a form of Eqs. (2) and (3) which averages over shot noise and cosmological fluctuations, as well as detector noise. The first and second moments of the GWB spherical harmonics under these averages are [22]

$$\langle \Omega_{\ell m}(t) \rangle_{S,\Omega} = \delta_{\ell 0} \sqrt{4\pi\bar{\Omega}}, \quad (8)$$

$$\text{Cov}[\Omega_{\ell m}(t), \Omega_{\ell' m'}(t')]_{S,\Omega} = \delta_{\ell\ell'} \delta_{mm'} (C_\ell + \delta_{t't'} W_\tau), \quad (9)$$

where C_ℓ is the true GWB angular power spectrum that we are aiming to measure, and $W_\tau \gg C_\ell$ is the shot noise power in a segment of length τ , which can be calculated for a given model of the AGWB using expressions in Ref. [22]. We see that the shot noise gives an excess contribution to the variance of GWB spherical harmonic measurements made at the same time $t = t'$, and that this contribution is spectrally white (i.e. independent of angular multipole ℓ). Combining Eqs. (2), (3), and (9), we find

$$\begin{aligned} & \text{Cov}[\hat{P}_{IJ}(f,t), \hat{P}_{I'J'}(f',t')]_{N,S,\Omega} \\ &= \sum_{\ell m} \tilde{\gamma}_{IJ}^{\ell m}(f,t) \tilde{\gamma}_{I'J'}^{\ell m*}(f',t') (C_\ell + \delta_{t't'} W_\tau) \\ &+ \delta_{ff'} \delta_{t't'} \left[(\delta_{II'} N_I(f,t) + \tilde{\gamma}_{II'}^{00}(f,t) \sqrt{4\pi\bar{\Omega}}) (\delta_{JJ'} N_J(f,t) + \tilde{\gamma}_{JJ'}^{00}(f,t) \sqrt{4\pi\bar{\Omega}}) + \sum_{\ell m} \tilde{\gamma}_{II'}^{\ell m}(f,t) \tilde{\gamma}_{JJ'}^{\ell m*}(f,t) (C_\ell + W_\tau) \right] \\ &+ \delta_{f,-f'} \delta_{t't'} \left[(\delta_{II'} N_I(f,t) + \tilde{\gamma}_{II'}^{00}(f,t) \sqrt{4\pi\bar{\Omega}}) (\delta_{I'J'} N_{J'}(f',t') + \tilde{\gamma}_{I'J'}^{00}(f',t') \sqrt{4\pi\bar{\Omega}}) + \sum_{\ell m} \tilde{\gamma}_{II'}^{\ell m}(f,t) \tilde{\gamma}_{I'J'}^{\ell m*}(f',t') (C_\ell + W_\tau) \right]. \quad (10) \end{aligned}$$

A naive first guess for an appropriate C_ℓ estimator is to compute the average power in the $(2\ell + 1)$ spherical harmonics corresponding to a given multipole ℓ ,

$$\hat{C}_\ell^{(\text{raw})} = \frac{1}{2\ell + 1} \sum_m |\hat{\Omega}_{\ell m}|^2. \quad (11)$$

Using Eqs. (6) and (10) we can show that the expected value of this quantity is

$$\langle \hat{C}_\ell^{(\text{raw})} \rangle_{N,S,\Omega} = C_\ell + \delta_{\ell 0} 4\pi\bar{\Omega}^2 + \mathcal{N}_\ell. \quad (12)$$

The second term here is to be expected, and simply reflects the fact that the monopole ($\ell = 0$) is not zero-mean like the other spherical harmonics. The third term is a *noise-induced bias*, which is obtained by expanding out the expressions in Eq. (6),

$$\begin{aligned} \mathcal{N}_\ell &= \frac{1}{2\ell + 1} \sum_m \sum_{\ell' m'} \sum_{\ell'' m''} (\Gamma^{-1})_{\ell m, \ell' m'} (\Gamma^{-1})_{\ell'' m'', \ell m} \sum_{J>I} \sum_{J'>I'} \sum_{f,t} \sum_{f',t'} \frac{\tilde{\gamma}_{IJ}^{\ell' m'}(f,t) \tilde{\gamma}_{I'J'}^{\ell'' m''}(f',t')}{\bar{P}_I(f,t) \bar{P}_J(f,t) \bar{P}_{I'}(f',t') \bar{P}_{J'}(f',t')} \\ &\times \delta_{t't'} \left\{ \sum_{LM} \tilde{\gamma}_{IJ}^{LM}(f,t) \tilde{\gamma}_{I'J'}^{LM*}(f',t') W_\tau \right. \\ &+ \delta_{ff'} \left[(\delta_{II'} N_I(f,t) + \tilde{\gamma}_{II'}^{00} \sqrt{4\pi\bar{\Omega}}) (\delta_{JJ'} N_J(f,t) + \tilde{\gamma}_{JJ'}^{00} \sqrt{4\pi\bar{\Omega}}) + \sum_{LM} \tilde{\gamma}_{II'}^{LM} \tilde{\gamma}_{JJ'}^{LM*} (C_L + W_\tau) \right] \\ &\left. + \delta_{f,-f'} \left[(\delta_{II'} N_I(f,t) + \tilde{\gamma}_{II'}^{00} \sqrt{4\pi\bar{\Omega}}) (\delta_{I'J'} N_{J'}(f',t') + \tilde{\gamma}_{I'J'}^{00} \sqrt{4\pi\bar{\Omega}}) + \sum_{LM} \tilde{\gamma}_{II'}^{LM} \tilde{\gamma}_{I'J'}^{LM*} (C_L + W_\tau) \right] \right\}. \quad (13) \end{aligned}$$

This expression involves contributions from pure detector noise, shot noise, and from the signal itself. In the limiting case where there is no signal and no shot noise, this simplifies to

$$\mathcal{N}_\ell^{(\text{lim})} = \frac{1}{2\ell + 1} \sum_m (\Gamma^{-1})_{\ell m, \ell m}, \quad (14)$$

which is the standard expression first derived in Ref. [27]. (We have assumed that $\bar{P}_I = N_I$ in this case.)

In existing LVK analyses [26], this leading-order contribution to the bias is explicitly subtracted from the angular power spectrum estimates, so that the corresponding estimator is given by

$$\begin{aligned} \hat{C}_\ell^{(\text{curr})} &= \hat{C}_\ell^{(\text{raw})} - \mathcal{N}_\ell^{(\text{lim})} \\ &= \frac{1}{2\ell + 1} \sum_m [|\hat{\Omega}_{\ell m}|^2 - (\Gamma^{-1})_{\ell m, \ell m}]. \end{aligned} \quad (15)$$

However, it is clearly not possible to set the GWB signal and shot noise to zero in real data, so Eq. (13) tells us that, even if we subtract the usual bias term (14), the current estimator $\hat{C}_\ell^{(\text{curr})}$ is still biased by an amount much larger than the signal we are searching for, since $N_I \bar{\Omega} \gg C_\ell$ and $W_\tau \gg C_\ell$. The exact size of this bias depends on the properties of the noise estimates $\bar{P}_I(f, t)$, which themselves include further contributions from the GWB. It is possible in principle that one could define a procedure for generating the noise estimates such that the bias *is* on average given by Eq. (14); however, it is not at all clear how to do this in practice.

A better alternative is to notice that all of the contributions to the bias (13) can be traced back to terms proportional to $\delta_{tt'}$ in Eq. (10). In other words, *the bias comes from autocorrelations of each data segment with itself*. One can thus remove the bias *exactly*, by explicitly excluding such autocorrelations when constructing the C_ℓ estimate. This is the same method that was first suggested in Ref. [22] for removing the bias induced by shot noise; however, we see here that the same procedure removes all contributions to the bias, including those from detector noise.

Practically, one proceeds by constructing multiple dirty maps $\hat{X}_{\ell m}^{(i)}$, $i = 1, \dots, n$, corresponding to disjoint subsets (or epochs) of the total dataset, each with its own Fisher matrix $\Gamma_{\ell m, \ell' m'}^{(i)}$. The optimal C_ℓ estimate is then formed by summing over all nonequal pairs of maps,

$$\hat{C}_\ell^{(\text{opt})} = \frac{1}{2\ell + 1} \frac{1}{n(n-1)} \sum_m \sum_{i \neq j} \hat{\Omega}_{\ell m}^{(i)} \hat{\Omega}_{\ell m}^{(j)*}. \quad (16)$$

One can then explicitly show, by a very similar calculation to those above, that this is unbiased,

$$\langle \hat{C}_\ell^{(\text{opt})} \rangle_{N, S, \Omega} = C_\ell + \delta_{\ell 0} 4\pi \bar{\Omega}^2. \quad (17)$$

We can calculate the variance of this estimator, using the fact that the spherical harmonic estimates $\hat{\Omega}_{\ell m}^{(i)}$ are Gaussian by the central limit theorem (since they are each formed by summing over many independent data segments). The general expression is very lengthy, but in the *weak-signal regime* where $\Gamma^{-1} \gg C_\ell, \bar{\Omega}^2, W_\tau$ it simplifies to give

$$\begin{aligned} \text{Var}[\hat{C}_\ell^{(\text{opt})}] &\simeq \frac{2}{(2\ell + 1)^2} \frac{1}{n^2(n-1)^2} \\ &\times \sum_{m, m'} \sum_{i \neq j} (\Gamma^{-1})_{\ell m, \ell m'}^{(i)} (\Gamma^{-1})_{\ell m', \ell m}^{(j)}. \end{aligned} \quad (18)$$

Let us consider the simplest case, where the Fisher matrices for each of the subsets of the data are equal to one another, and are thus given by $1/n$ times the total Fisher matrix,

$$\Gamma_{\ell m, \ell' m'}^{(i)} \simeq \frac{1}{n} \Gamma_{\ell m, \ell' m'}. \quad (19)$$

We can then simplify Eq. (18) to give

$$\text{Var}[\hat{C}_\ell^{(\text{opt})}] \simeq \frac{2}{(2\ell + 1)^2} \frac{n}{n-1} \sum_{m, m'} |(\Gamma^{-1})_{\ell m, \ell m'}|^2. \quad (20)$$

This is equal to the variance found in Ref. [27] for the current estimator, times a factor of $n/(n-1)$. We therefore see that our unbiased estimator is equally efficient in the limit of large n , just as in Ref. [22].

III. VALIDATION ON MOCK DATA

A. O3 sensitivity

1. Pipeline and settings

To run our anisotropic analysis, we make use of the publicly available code `PyStoch` [30,31]. This pipeline has been designed and tested on O3 folded data [32] (CSD and PSD), for which it computes the dirty map and Fisher matrix either in pixel space or spherical harmonic space. The user specifies an assumed spectral index, α , a frequency range (the whole GWB frequency range in the standard LVK analysis corresponds to [20, 1726] Hz), a detector baseline, and the number of pixels or maximum multipole they want to probe, ℓ_{max} .

In our case, we adopt $\alpha = 2/3$ to consider an astrophysical background [33] and run `PyStoch` in spherical harmonic space assuming the LIGO Hanford-Livingston (HL) baseline. We produce mock CSDs¹ made up of the following components:

¹In this exercise, only the CSDs are drawn randomly to represent individual data realizations. The PSD for each realization is kept fixed, essentially reducing the Fisher matrix to a known weight in order to model the ideal scenario of equal-length datasets. This approach simplified computations.

- (i) *Detector noise.* This is generated as random Gaussian realizations (different at every element in our time-frequency CSD array) of the folded data [34] at O3 sensitivity. For increased LIGO sensitivity (see Sec. III B), we use the expected amplitude spectral densities (ASDs) from Ref. [35].
- (ii) *Shot noise.* We generate shot-noise sky maps (different at every time segment in our data) with Healpy [36,37] as random realizations of the same flat C_ℓ spectrum (corresponding to white noise). We then multiply each map with the ORF [given by Eq. (4)] to produce a shot-noise CSD.
- (iii) *Mock AGWB signal.* We simulate a sky map that represents the AGWB signal. Again, this is done with Healpy, using a model of the cosmological C_ℓ spectrum from Ref. [19] (see Sec. III A 3). We subsequently multiply this map with the ORF [Eq. (4)] to produce a CSD for the GWB.

We produce dirty maps, Fisher matrices and clean maps from our simulated CSDs by means of Eqs. (5) and (6). In the results that follow, we have not applied any regularization schemes to the Fisher matrix inversion. See Appendix B for a discussion on the impact of regularization in our study. Given that we are working in the weak-signal regime, we expect the associated statistical uncertainty in the C_ℓ calculation to be large compared to the injected values for the short-duration (~ 1 yr) datasets simulated here.

Our PyStoch-based analysis is sufficiently fast that we can simulate many independent realizations of shot noise and detector noise. We achieve this by choosing appropriate values for the maximum spherical multipole ($\ell_{\max} = 8$) and the analyzed frequency range ($[20, 520]$ Hz), which allow us to perform our simulations relatively cheaply while still capturing the key features of the problem. We stress however that our methods are general, and can be applied to a broader range of multipoles and frequencies in a full analysis.

2. Validation with existing estimator

First we confirm that the existing estimator ($\hat{C}_\ell^{(\text{curr})}$) performs as expected in the weak-signal regime by running PyStoch on simulated detector noise, without injecting shot noise or an underlying cosmological signal. In Fig. 1, we plot the clean-map angular power spectra of 1,000 random detector noise realizations (blue curves), along with their mean (black curve). As expected, this mean curve exhibits excellent agreement with the weak-signal limit of the bias, $\mathcal{N}_\ell^{(\text{lim})}$, as given by Eq. (14) (dashed yellow curve). In Fig. 2 we plot the estimator $\hat{C}_\ell^{(\text{curr})}$, which corresponds to the same spectra but with the weak-signal limit of the bias subtracted. At O3 sensitivity, the shot-noise and GWB contributions to the bias in the current estimator are much smaller than the uncertainty of the estimate, so that the resulting curve is




FIG. 1. The angular power spectra of 1,000 clean maps (blue curves), as recovered from simulations with detector noise only at O3 sensitivity (mock 1-month data sets). The mean $\hat{C}_\ell^{(\text{raw})}$ spectrum (black curve) matches the expected bias in the weak-signal limit, $\mathcal{N}_\ell^{(\text{lim})}$ (dashed yellow curve).

consistent with zero for pure detector noise. In the same figure, we plot the expected variance of this current estimator from our calculations [Eq. (20) in the $n \rightarrow \infty$ limit], as well as the sample variance observed in our ensemble of 1,000 noise realizations. The excellent agreement between these is a useful consistency check for our simulation pipeline.

3. Modeling shot noise and large-scale structure

We now turn to the full problem, in which the simulated data contain a GWB signal with intrinsic angular correlations due to cosmic large-scale structure, as well as spurious correlations due to temporal shot noise. We simulate the maps that we inject in random detector noise realizations using realistic values for the shot noise and LSS (monopole and anisotropies) cited in Ref. [19]. The reported values in that work for the angular power spectrum are in units of Ω_{GW}^2 at a frequency of $f_{\text{ref}} = 65$ Hz. We consider the shot noise and LSS values estimated for a subtracted foreground of $r_* = 250$ Mpc (some degree of foreground subtraction must be assumed in order to avoid a divergence in the expected shot noise power at short distances). See Appendix A for a detailed conversion of these results into the right units and f_{ref} for our analysis.

In Fig. 3, we plot the predicted bias for the existing estimator $\hat{C}_\ell^{(\text{curr})}$ in the weak-signal limit, as given by Eq. (14), along with the expected angular power spectra for the shot noise and underlying AGWB. Everything has been normalized with respect to the monopole, $C_\ell \rightarrow C_\ell / (4\pi\Omega^2)$. We note that for the detector noise and the shot noise, the assumed amount of time over which each data set has been collected is one year. For O3 sensitivity, the detector noise is orders of magnitude above

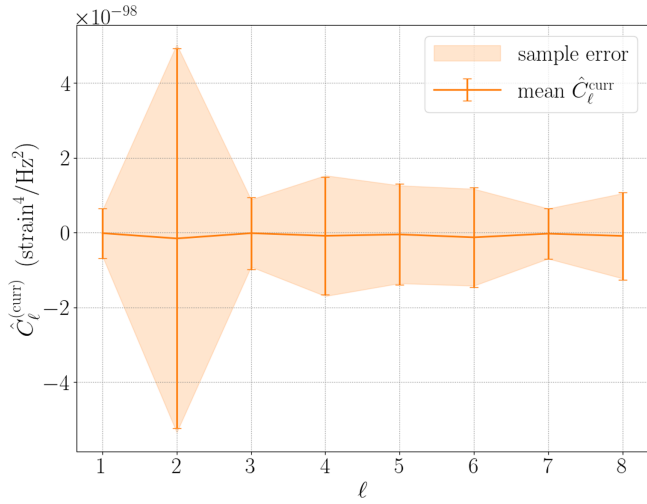


FIG. 2. The mean $\hat{C}_\ell^{(\text{curr})}$ estimator of 1,000 simulations with detector noise only, at O3 sensitivity. We plot the predicted variance of the estimator with error bars, and the sample variance in our ensemble of simulations as a shaded region.

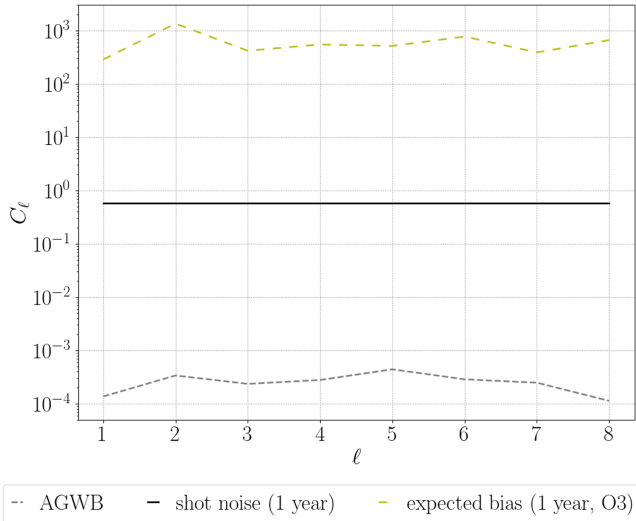


FIG. 3. Expected angular power spectrum of the temporal shot noise (black curve) and the AGWB (dashed gray curve). The dashed yellow curve denotes the expected bias in the weak-signal limit, given by the O3 (HL) detector noise. Both the detector noise and the shot noise correspond to 1 year of data. All curves have been normalized with respect to the AGWB monopole.

the shot noise and AGWB signal, so that we are unambiguously in the weak-signal regime. It is also clear that for O3 sensitivity, we cannot expect to probe the angular power spectrum of the AGWB. This is consistent with the non-detection of such a signal in O3 data [26].

4. Injected map simulations and angular power spectra

We generate detector noise and shot noise assuming one-month datasets, such that both vary randomly at every time




FIG. 4. The angular power spectra of 1,200 clean maps (blue curves), normalized to the monopole, as recovered from simulations with O3 detector noise, shot noise and the AGWB (mock 1-month datasets). The mean $\hat{C}_\ell^{(\text{raw})}$ spectrum (black curve) matches the expected (O3) bias in the weak-signal regime (dashed yellow curve).

segment in our datasets. We then inject the same, fixed underlying LSS signal to each dataset.

In total, we produce 1,200 clean maps corresponding to one month of data each. In Fig. 4, we plot the angular power spectra for these clean maps, along with the ensemble mean. This once again matches the bias in Eq. (14), as expected in the weak-signal regime.

We subsequently distribute the clean maps in 100 sets of 12 maps each to construct one-year datasets.² For each set, we calculate the $\hat{C}_\ell^{(\text{curr})}$ and $\hat{C}_\ell^{(\text{opt})}$ estimators. The results are plotted in Fig. 5, along with the predicted variance [given by Eq. (20) for $\hat{C}_\ell^{(\text{opt})}$, and the same expression in the $n \rightarrow \infty$ limit for $\hat{C}_\ell^{(\text{curr})}$] and sample variance for both estimators, which are in good agreement in both cases. Due to the weak signal and short time span (relative to O3 sensitivity), both estimators have large uncertainties, and are both equally consistent with the injected signal. (The variance of the unbiased estimator is larger by a factor of 12/11, as expected; this factor can be reduced by dividing the data into a larger number of subsets.) This confirms that the effects of shot noise and the bias in the existing analysis method are both negligibly small at O3 sensitivity. In the section below, we show that this will no longer be true at design sensitivity.

²The corresponding dirty map and Fisher matrix for each one-year dataset that are used in the $\hat{C}_\ell^{(\text{curr})}$ estimator calculation are obtained by adding up the 12 different dirty maps and Fisher matrices, respectively.




FIG. 5. Comparison of the performance of the $\hat{C}_\ell^{(\text{curr})}$ (orange) and $\hat{C}_\ell^{(\text{opt})}$ (blue) estimators in a realistic O3-like scenario. We plot the expected variances as error bars and the sample variances as error bands. Everything has been normalized with respect to the AGWB monopole. For 1 year of O3 data, the AGWB (dashed gray curve) is very small compared to the uncertainties in the angular power spectrum estimate.

B. O5 sensitivity

Advanced LIGO is expected to reach design sensitivity, A+, during the O5 observing run. We run simulations for this detector sensitivity to see how the two C_ℓ estimators perform in this setting. We follow the procedure that is outlined in section III A, with the sole difference that the detector noise is generated for A+ instead of O3 sensitivity.

Again, we produce 1,200 clean maps corresponding to one month of data each. We plot the angular power spectra of all clean maps in Fig. 6. This time, their mean clearly lies above the weak-signal expression for the bias, $\mathcal{N}_\ell^{(\text{lim})}$. This is unsurprising, as the monopole and shot noise terms in Eq. (13) are significantly larger compared to the detector noise at A+ sensitivity than at O3 sensitivity. As a result, the limiting-case expression for the expected bias is no longer sufficient to describe the actual bias in the analysis. Therefore, the estimator $\hat{C}_\ell^{(\text{curr})}$ we obtain after subtracting the detector bias will still be biased by the subleading contributions, which are no longer negligible. In contrast, the optimal estimator introduced in this work is always unbiased; hence, we expect the corresponding curves to lie below those of the current estimator for these simulations.

As before, we distribute the clean maps in 100 sets of 12 maps each and plot the current and optimal estimators in Fig. 7, along with their expected and sample variances. This time, there is a lack of agreement between the sample variance for each estimator and the theoretical prediction given by Eq. (20), with the latter being consistently smaller. This can be understood as being due to the breakdown of the weak-signal approximation; one now needs to account




FIG. 6. The angular power spectrum of 1,200 clean maps (blue curves), normalized to the monopole, as recovered from simulations with O5 detector noise, shot noise and the AGWB (mock 1-month datasets). The mean $\hat{C}_\ell^{(\text{raw})}$ spectrum (black curve) and the expected (O5) bias in the weak-signal regime (dashed yellow curve) are also plotted. The discrepancy between the black and yellow curves indicates the presence of bias due to beyond-weak-signal-limit effects.




FIG. 7. Comparison of the performance of the $\hat{C}_\ell^{(\text{curr})}$ (orange curve) and $\hat{C}_\ell^{(\text{opt})}$ (blue) estimators. We plot the expected variances as error bars and the sample variances as error bands. Everything has been normalized with respect to the AGWB monopole. For 1 year of O5 data, the AGWB spectrum (dashed gray curve) is still small compared to the uncertainties in the angular power spectrum calculation.

for additional contributions to the variance coming from the shot noise and the AGWB. We also see that the optimal estimator $\hat{C}_\ell^{(\text{opt})}$ returns a spectrum that is, on average, significantly below that of the current estimator $\hat{C}_\ell^{(\text{curr})}$, and is thus closer to the injected C_ℓ signal. This reflects the bias inherent to the current estimator beyond the weak-signal

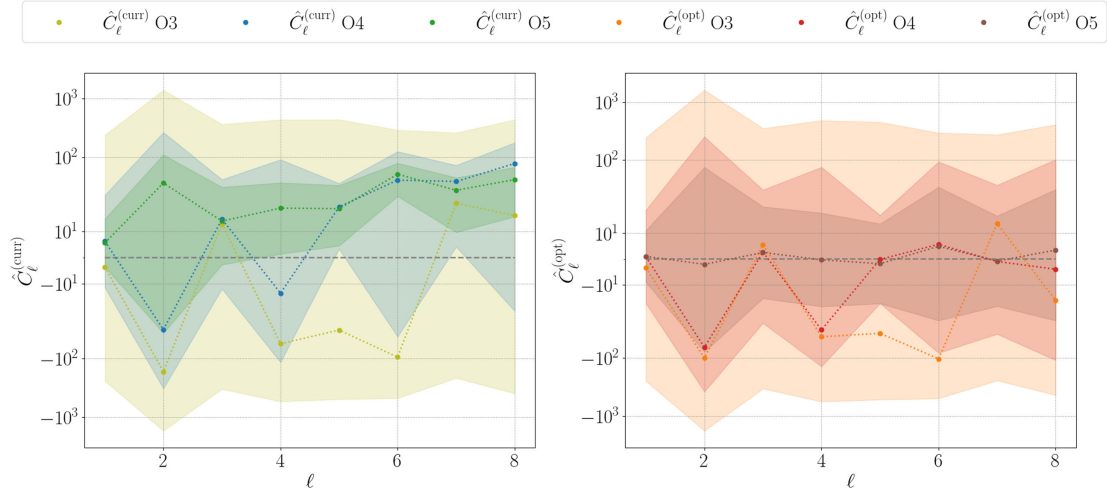


FIG. 8. Summary of the $\hat{C}_\ell^{(\text{curr})}$ (left panel) and $\hat{C}_\ell^{(\text{opt})}$ (right panel) estimator performance for the O3, O4 and O5 sensitivities. The sample variances are plotted as error bands. The dashed gray curve denotes the AGWB. It is clear that increasing the detector sensitivity results in the optimal estimator recovering the true AGWB spectrum more precisely, and the bias in the current estimator becoming more evident. Note: because of the symlog scale for the y-axis, error bands with the same physical size but at different vertical positions in the plot span different ranges of values.

regime, and illustrates that it will be crucial to use the unbiased estimator $\hat{C}_\ell^{(\text{opt})}$ once A + sensitivity is reached.

C. Discussion

We close this section with a comparison of our mock simulations for O3, O4, and O5 sensitivity. We repeat our analysis for the projected O4 sensitivity as well to include a forecast of the effect on present-day data. Much like for the case of A+, we notice that the mean C_ℓ curve for the 1,200 produced clean maps lies above the detector bias. However, the discrepancy is not so clear for O4, starting as negligible at lower multipoles, and only becoming measurable at higher multipoles. Looking at the performance of the two estimators, we therefore find that the optimal estimator performs visibly better than the current one only for $\ell \gtrsim 5$.

In Fig. 8, we summarise our results for the current (left panel) and optimal (right panel) estimators, for the three different detector sensitivities that we assumed during our

analysis. The different C_ℓ estimators are plotted as dotted curves of different colors (summarized in the figure legend). Again, the error bands denote the sample variances, while the AGWB (dashed gray curve) is also plotted.

For both estimators, increasing the detector sensitivity results in narrower error bands, as expected. In the case of the optimal estimator, these bands are always centered on the injected AGWB signal; however, for the current estimator, they deviate away from the AGWB, towards positive values (particularly for higher multipoles). Hence, while the optimal estimator yields more precise and accurate C_ℓ measurements for increased detector sensitivity, the bias in the current estimator becomes increasingly evident.

We quantify the deviation of the mean of the $\hat{C}_\ell^{(\text{curr})}$ estimator from the AGWB signal in terms of the sample standard deviations difference between them, which is summarized in Table I.

TABLE I. The difference between the mean of the $\hat{C}_\ell^{(\text{curr})}$ estimator and the AGWB signal at each ℓ mode in units of the sample standard deviation (STD).

Sample STD difference								
ℓ mode								
Detector sensitivity	1	2	3	4	5	6	7	8
O3	-0.015	-0.109	0.037	-0.116	-0.071	-0.25	0.088	0.039
O4	0.352	-0.11	0.552	-0.131	1.204	0.484	1.12	0.791
O5	0.612	0.492	0.838	1.078	1.318	1.85	1.556	1.612

IV. SUMMARY AND OUTLOOK

In this study, we have investigated methods for measuring GW background anisotropies in the presence of both detector noise and temporal shot noise. We started by describing the C_ℓ estimator currently used in LIGO-Virgo-KAGRA analyses, and demonstrated that this estimator is biased due to contributions from the shot noise and from the target GW signal itself. We then presented a new estimator which avoids this bias by only cross-correlating between data segments at unequal times, validating our predictions on simulations using mock data. The two estimators produce similar results for O3 sensitivity, where we confidently lie in the weak-signal regime. However, we find that for LIGO A+ sensitivity it is crucial to use our new estimator to obtain accurate inferences of the GWB angular power spectrum, with some improvement over the current estimator already visible at O4 sensitivity.

In the future, we plan to extend this work by testing both estimators on real O4 data. Without the computational cost of running the analysis on thousands of simulated datasets, it will be possible to probe higher spherical harmonic ℓ modes, in the whole GWB frequency range ([20, 1726] Hz). It will also be straightforward to include more detectors in the analysis—such as Virgo and KAGRA—and perform the analysis on multiple detector baselines. This will improve the conditioning of the Fisher matrices involved in the angular power spectrum estimation, allowing access to higher angular resolution.

In addition, it will be interesting to test the capabilities of planned ‘third-generation’ detectors in probing GW anisotropies. It is already clear that the large bias in the angular power spectrum for future detectors will render the current search method unsuitable. On the other hand, the optimal estimator will be able to measure the LSS anisotropies with precision orders of magnitude better than for current detectors.

ACKNOWLEDGMENTS

We would like to thank Shivaraj Kandhasamy for carefully reviewing this manuscript (LIGO-Document No. P2300434) as a part of the LVK Collaboration’s internal review process. This material is based upon work supported by NSF’s LIGO Laboratory, which is a major facility fully funded by the National Science Foundation. The authors acknowledge computational resources provided by the LIGO Laboratory and supported by NSF Grants No. PHY-0757058 and No. PHY-0823459. We thank Deepali Agarwal and Stavros Venikoudis for useful comments, and in particular Jishnu Suresh for many useful discussions on how to run PyStoch. N. K. is supported by King’s College London through an NMES Funded Studentship. A. C. J. is supported by the Science and Technology Facilities Council (STFC) through the UKRI Quantum Technologies for Fundamental Physics

Programme [Grant No. ST/T005904/1]. A. I. R. is supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 101064542, and acknowledges support from the NSF Award No. PHY-1912594. J. D. R. acknowledges support from National Science Foundation (NSF) Grant No. PHY-2207270. M. S. acknowledges support from the Science and Technology Facilities Council (STFC), UK, under the research Grant No. ST/X000753/1. This work was partly enabled by the UCL Cosmoparticle Initiative.

APPENDIX A: UNIT CONVERSION

The GW energy density is related to the GW strain power by [38]

$$\Omega_{\text{GW}}(f, \theta) = \frac{2\pi^2}{3H_0^2} f^3 P(f, \theta), \quad (\text{A1})$$

where $H_0 = 67.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [39]. The GW energy density and power can then be factored into spectral and angular terms,

$$\Omega_{\text{GW}}(f, \theta) = \Omega_{\text{GW}}(\theta) \left(\frac{f}{f_{\text{ref}}} \right)^\alpha, \quad (\text{A2})$$

$$P(f, \theta) = P(\theta) H(f), \quad (\text{A3})$$

where $H(f) = (f/f_{\text{ref}})^{\alpha-3}$. This leads to

$$P(\theta) = A \Omega_{\text{GW}}(\theta), \quad (\text{A4})$$

where $A = 3H_0^2 / (2\pi^2 f_{\text{ref}}^3)$ (this should be further reduced by a factor of 2 to account for taking the one-sided spectrum).

This expression allows us to convert a sky map from units of Ω_{GW} to units of P . Finally, we convert the results in Ref. [19] (which correspond to a frequency of 65 Hz) to the appropriate reference frequency by means of Eq. (A2).

APPENDIX B: IMPACT OF REGULARIZATION

In this paper, we have chosen not to employ a regularization scheme when inverting the Fisher matrix, as the choice for regularization is not obvious and the community has yet to reach a consensus on the appropriate way to regularise the matrix for present-day detector networks (see e.g. studies presented in [40–43]). As shown in [41], the choice of regularization scheme depends on whether the search prioritizes signal detection versus characterization; this goes beyond the scope of our work, and we do not expect this to qualitatively change the results of this case study. Nevertheless, for completeness, we assess the impact that regularizing the Fisher matrix has on our results in this appendix. We consider two different regularization schemes on the Fisher matrix inversion:

- (i) a condition number cutoff (CNC) of 10^{-5} (as done in [23,24,44,45]);
- (ii) a reassignment of 1/3 of the Fisher matrix eigenvalues (RE) to infinity (as customarily done in LVK Collaboration efforts [25,26,46]).

This CNC approach is generally less aggressive than the RE, as it allows for no eigenvalues to be discarded during matrix inversion, while the RE approach blindly discards 1/3 of the eigenvalues, no matter their actual amplitude. The latter can lead to an earlier detection but can also underestimate the measurement errors, while the former can lead to a large amount of noise leaking into the measurement. When employing a regularization scheme, formally the inverse of Γ , Γ^{-1} , should be replaced by the regularized inverse $(\Gamma')^{-1}$. Specifically, this should be done in Eqs. (5), (13)–(15), (18), (20).

We find that the ability of the current estimator to recover the AGWB signal (first shown in Fig. 2) is not affected by our regularization schemes. Similarly, the performance of the current estimator in the presence of detector noise, shot noise, and the AGWB is not affected significantly in the

case of O3 sensitivity (first tested in Fig. 5). However, the result changes significantly for O4 and O5 sensitivity. In Fig. 9, we plot the current (top) and optimal (bottom) estimator, for O4 (left) and O5 (right) sensitivity and the CNC and RE regularization schemes. The CNC scheme has an almost negligible impact on the Fisher matrix inversion and the current and optimal estimators (same for O3), giving almost the same results as reported in Sec. III. This is because the chosen CNC is very close to the Fisher matrix's condition number in this case. On the other hand, imposing the RE scheme significantly affects the Fisher matrix inversion and results in an enhancement of the bias in the current estimator. As for the optimal estimator, this performs equally well for both regularization schemes. Furthermore, the discrepancy among the theoretical uncertainties and sample variances for O4 and O5 sensitivity appears greater for the RE scheme. We conclude that when using a more aggressive regularization scheme, error estimates are reduced and hence the noise bias becomes more evident in the current estimator, and the weak-signal approximation loses validity.

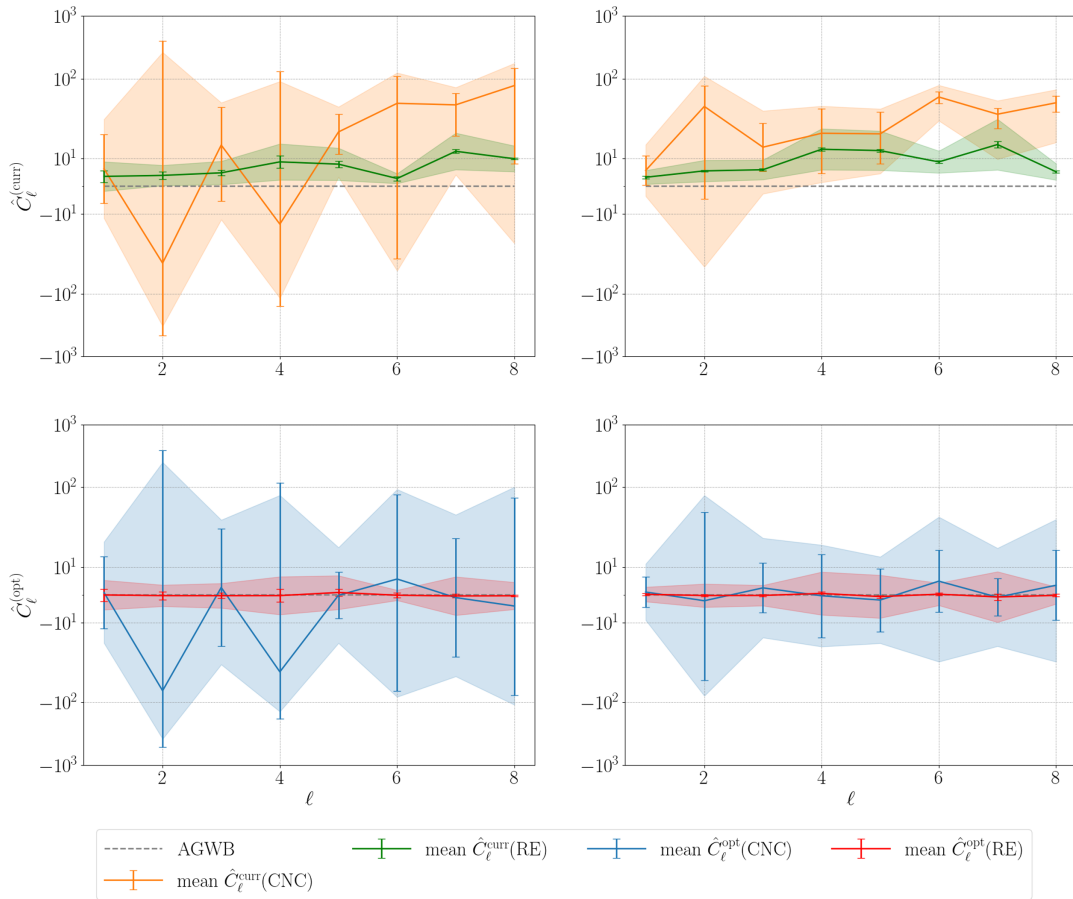


FIG. 9. Performance of the $\hat{C}_\ell^{(\text{curr})}$ (top) and $\hat{C}_\ell^{(\text{opt})}$ (bottom) estimators, for O4 (left) and O5 (right) detector sensitivity and the CNC and RE regularization schemes. We plot the expected variances as error bars and the sample variances as error bands. Everything has been normalized with respect to the AGWB monopole. Imposing a more aggressive regularization scheme results in the bias in the $\hat{C}_\ell^{(\text{curr})}$ estimator becoming more evident.

- [1] LIGO Scientific and Virgo Collaborations, GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs, *Phys. Rev. X* **9**, 031040 (2019).
- [2] LIGO Scientific and Virgo Collaborations, GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, *Phys. Rev. X* **11**, 021053 (2021).
- [3] LIGO Scientific, Virgo, and KAGRA Collaborations, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, *Phys. Rev. X* **13**, 041039 (2023).
- [4] LIGO Scientific, Virgo, and KAGRA Collaborations, Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, *Living Rev. Relativity* **21**, 3 (2018).
- [5] B. Allen, The stochastic gravity-wave background: Sources and detection, in *Proceedings, Les Houches School of Physics: Astrophysical Sources of Gravitational Radiation*, edited by J.-A. Marck and J.-P. Lasota, Cambridge Contemporary Astrophysics (1997) pp. 373–417, [arXiv:gr-qc/9604033](https://arxiv.org/abs/gr-qc/9604033).
- [6] N. Christensen, Stochastic gravitational wave backgrounds, *Rep. Prog. Phys.* **82**, 016903 (2018).
- [7] EPTA and InPTA Collaborations, The second data release from the European Pulsar Timing Array—III. Search for gravitational wave signals, *Astron. Astrophys.* **678**, A50 (2023).
- [8] NANOGrav Collaboration, The NANOGrav 15 yr data set: Evidence for a gravitational-wave background, *Astrophys. J. Lett.* **951**, L8 (2023).
- [9] PPTA Collaboration, Search for an isotropic gravitational-wave background with the parkes pulsar timing array, *Astrophys. J. Lett.* **951**, L6 (2023).
- [10] C. Caprini and D. G. Figueroa, Cosmological backgrounds of gravitational waves, *Classical Quantum Gravity* **35**, 163001 (2018).
- [11] T. Regimbau, The astrophysical gravitational wave stochastic background, *Res. Astron. Astrophys.* **11**, 369 (2011).
- [12] C. R. Contaldi, Anisotropies of gravitational wave backgrounds: A line of sight approach, *Phys. Lett. B* **771**, 9 (2017).
- [13] G. Cusin, C. Pitrou, and J.-P. Uzan, Anisotropy of the astrophysical gravitational wave background: Analytic expression of the angular power spectrum and correlation with cosmological observations, *Phys. Rev. D* **96**, 103019 (2017).
- [14] G. Cusin, I. Dvorkin, C. Pitrou, and J.-P. Uzan, First predictions of the angular power spectrum of the astrophysical gravitational wave background, *Phys. Rev. Lett.* **120**, 231101 (2018).
- [15] A. C. Jenkins, M. Sakellariadou, T. Regimbau, and E. Slezak, Anisotropies in the astrophysical gravitational-wave background: Predictions for the detection of compact binaries by LIGO and Virgo, *Phys. Rev. D* **98**, 063501 (2018).
- [16] A. C. Jenkins, R. O’Shaughnessy, M. Sakellariadou, and D. Wysocki, Anisotropies in the astrophysical gravitational-wave background: The impact of black hole distributions, *Phys. Rev. Lett.* **122**, 111101 (2019).
- [17] G. Cusin, I. Dvorkin, C. Pitrou, and J.-P. Uzan, Properties of the stochastic astrophysical gravitational wave background: Astrophysical sources dependencies, *Phys. Rev. D* **100**, 063004 (2019).
- [18] G. Capurri, A. Lapi, C. Baccigalupi, L. Boco, G. Scelfo, and T. Ronconi, Intensity and anisotropies of the stochastic gravitational wave background from merging compact binaries in galaxies, *J. Cosmol. Astropart. Phys.* **11** (2021) 032.
- [19] A. C. Jenkins and M. Sakellariadou, Shot noise in the astrophysical gravitational-wave background, *Phys. Rev. D* **100**, 063508 (2019).
- [20] S. Mukherjee and J. Silk, Time-dependence of the astrophysical stochastic gravitational wave background, *Mon. Not. R. Astron. Soc.* **491**, 4690 (2020).
- [21] M. R. Sah and S. Mukherjee, Non-stationary astrophysical stochastic gravitational-wave background: a new probe to the high-redshift population of binary black holes, *Mon. Not. R. Astron. Soc.* **527**, 4100 (2023).
- [22] A. C. Jenkins, J. D. Romano, and M. Sakellariadou, Estimating the angular power spectrum of the gravitational-wave background in the presence of shot noise, *Phys. Rev. D* **100**, 083501 (2019).
- [23] A. I. Renzini and C. R. Contaldi, Gravitational-wave background sky maps from Advanced LIGO O1 data, *Phys. Rev. Lett.* **122**, 081102 (2019).
- [24] A. Renzini and C. Contaldi, Improved limits on a stochastic gravitational-wave background and its anisotropies from Advanced LIGO O1 and O2 runs, *Phys. Rev. D* **100**, 063527 (2019).
- [25] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Directional limits on persistent gravitational waves using data from Advanced LIGO’s first two observing runs, *Phys. Rev. D* **100**, 062001 (2019).
- [26] LIGO Scientific, Virgo, and KAGRA Collaborations, Search for anisotropic gravitational-wave backgrounds using data from Advanced LIGO and Advanced Virgo’s first three observing runs, *Phys. Rev. D* **104**, 022005 (2021).
- [27] E. Thrane, S. Ballmer, J. D. Romano, S. Mitra, D. Talukder, S. Bose, and V. Mandic, Probing the anisotropies of a stochastic gravitational-wave background using a network of ground-based laser interferometers, *Phys. Rev. D* **80**, 122002 (2009).
- [28] LIGO Scientific, Virgo, and KAGRA Collaborations, Upper limits on the isotropic gravitational-wave background from Advanced LIGO and Advanced Virgo’s third observing run, *Phys. Rev. D* **104**, 022004 (2021).
- [29] L. Valbusa Dall’Armi, A. Ricciardone, and D. Bertacca, The dipole of the astrophysical gravitational-wave background, *J. Cosmol. Astropart. Phys.* **11** (2022) 040.
- [30] A. Ain, J. Suresh, and S. Mitra, Very fast stochastic gravitational wave background map making using folded data, *Phys. Rev. D* **98**, 024001 (2018).
- [31] J. Suresh, A. Ain, and S. Mitra, Unified mapmaking for an anisotropic stochastic gravitational wave background, *Phys. Rev. D* **103**, 083024 (2021).
- [32] A. Ain, P. Dalvi, and S. Mitra, Fast gravitational wave radiometry using data folding, *Phys. Rev. D* **92**, 022003 (2015).

- [33] E. S. Phinney, A practical theorem on gravitational wave backgrounds, [arXiv:astro-ph/0108028](https://arxiv.org/abs/astro-ph/0108028).
- [34] LIGO Scientific, Virgo, and KAGRA Collaborations, Folded data for first three observing runs of Advanced LIGO and Advanced Virgo, [10.5281/zenodo.6326656](https://zenodo.org/record/6326656) (2022).
- [35] LIGO Scientific, Virgo, and KAGRA Collaborations, Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, *Living Rev. Relativity* **21**, 3 (2018).
- [36] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski, `healpy`: Equal area pixelization and spherical harmonics transforms for data on the sphere in PYTHON, *J. Open Source Software* **4**, 1298 (2019).
- [37] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere, *Astrophys. J.* **622**, 759 (2005).
- [38] B. Allen and J.D. Romano, Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities, *Phys. Rev. D* **59**, 102001 (1999).
- [39] Planck Collaboration, Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [40] L. Xiao, A. I. Renzini, and A. J. Weinstein, Model-independent search for anisotropies in stochastic gravitational-wave backgrounds and application to LIGO-Virgo's first three observing runs, *Phys. Rev. D* **107**, 122002 (2023).
- [41] E. Floden, V. Mandic, A. Matas, and L. Tsukada, Angular resolution of the search for anisotropic stochastic gravitational-wave background with terrestrial gravitational-wave detectors, *Phys. Rev. D* **106**, 023010 (2022).
- [42] S. Panda, S. Bhagwat, J. Suresh, and S. Mitra, Stochastic gravitational wave background mapmaking using regularized deconvolution, *Phys. Rev. D* **100**, 043541 (2019).
- [43] D. Agarwal, J. Suresh, S. Mitra, and A. Ain, Upper limits on persistent gravitational waves using folded data and the full covariance matrix from Advanced LIGO's first two observing runs, *Phys. Rev. D* **104**, 123018 (2021).
- [44] A. I. Renzini and C. R. Contaldi, Mapping incoherent gravitational wave backgrounds, *Mon. Not. R. Astron. Soc.* **481**, 4650 (2018).
- [45] C. R. Contaldi, M. Pieroni, A. I. Renzini, G. Cusin, N. Karnesis, M. Peloso, A. Ricciardone, and G. Tasinato, Maximum likelihood map-making with the Laser Interferometer Space Antenna, *Phys. Rev. D* **102**, 043502 (2020).
- [46] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Directional limits on persistent gravitational waves from Advanced LIGO's first observing run, *Phys. Rev. Lett.* **118**, 121102 (2017).