Analytic distribution of the optimal cross-correlation statistic for stochastic gravitational-wave-background searches using pulsar timing arrays

Jeffrey S. Hazboun,[1,*] Patrick M. Meyers,[2] Joseph D. Romano,[3,4,†] Xavier Siemens,[1] and Anne M. Archibald

[1]*Department of Physics, Oregon State University, Corvallis, Oregon 97331, USA*
[2]*Theoretical Astrophysics Group, California Institute of Technology, Pasadena, California 91125, USA*
[3]*Department of Physics and Astronomy, Texas Tech University, Lubbock, Texas 79409-1051, USA*
[4]*Department of Physics and Astronomy, University of Texas Rio Grande Valley, One West University Boulevard, Brownsville, Texas 78520, USA*

(Received 21 June 2023; accepted 11 October 2023; published 22 November 2023)

We show via both analytical calculation and numerical simulation that the optimal cross-correlation statistic (OS) for stochastic gravitational-wave-background (GWB) searches using data from pulsar timing arrays follows a generalized chi-squared (GX2) distribution—i.e., a linear combination of chi-squared distributions with coefficients given by the eigenvalues of the quadratic form defining the statistic. This observation is particularly important for calculating the frequentist statistical significance of a possible GWB detection, which depends on the exact form of the distribution of the OS signal-to-noise ratio $\hat{\rho} \equiv \hat{A}^2_{\rm gw}/\sigma_0$ in the absence of GW-induced cross correlations (i.e., the null distribution). Previous discussions of the OS have incorrectly assumed that the analytic null distribution of $\hat{\rho}$ is well approximated by a zero-mean unit-variance Gaussian distribution. Empirical calculations show that the null distribution of $\hat{\rho}$ has "tails" which differ significantly from those for a Gaussian distribution but which follow (exactly) a GX2 distribution. Thus, a correct analytical assessment of the statistical significance of a potential detection requires the use of a GX2 distribution.

DOI: 10.1103/PhysRevD.108.104050

I. INTRODUCTION

Pulsar timing array (PTA) collaborations have recently reported evidence for a stochastic background of nanohertz gravitational waves [1–3]. These findings were corroborated by a meta-analysis of these separate data sets by the International Pulsar Timing Array (IPTA) [4]. All of these studies confirmed the presence of a common-spectrum red-noise process across pulsars, along with evidence for the quadrupolar spatial correlations [5] necessary to attribute the signal to a GWB.

A computationally efficient technique used by the pulsar timing community to calculate the significance of the cross correlations involves the so-called optimal statistic (OS) [6–8]. This statistic, denoted $\hat{A}^2_{\rm gw}$, is an unbiased estimator for the squared GWB amplitude $A^2_{\rm gw}$ derived by maximizing the logarithm of the likelihood ratio. The corresponding signal-to-noise ratio $\hat{\rho} \equiv \hat{A}^2_{\rm gw}/\sigma_0$, where σ_0^2 is the variance of the estimator $\hat{A}^2_{\rm gw}$ in the absence of GW-induced spatial correlations, can be related to the Bayesian odds ratio between a model with correlations and a model without correlations via the Laplace approximation [9]. The work in this paper concerns the calculation of the probability distribution for $\hat{\rho}$.

The distribution for $\hat{\rho}$ in the absence of such spatial correlations is called the null distribution and is denoted by $p(\hat{\rho}|H_0)$. Here, H_0 is the null hypothesis—i.e., the hypothesis that there are no GW-induced spatial correlations in the data. Although H_0 assumes no spatial correlations, it does allow for the presence of a nonzero common-spectrum red-noise process, whose amplitude $A_{\rm cp}$ is determined from a joint noise analysis for the pulsars in the array. As such, the null distribution H_0 depends on the particular value of $A_{\rm cp}$—i.e., $H_0 = H_0(A_{\rm cp})$. However, to simplify the notation in what follows, we will not explicitly display the $A_{\rm cp}$ dependence of H_0, although we will investigate the dependence of the null distribution on the amplitude and spectral shape of the common-spectrum red-noise process.

Given the null distribution, we can calculate the probability that our measured signal-to-noise ratio, denoted $\hat{\rho}_{\rm obs}$, could have resulted from noise alone. This is called the p-value and is defined as

$$p \equiv {\rm Prob}(\hat{\rho} > \hat{\rho}_{\rm obs}|H_0) \equiv \int_{\hat{\rho}_{\rm obs}}^{\infty} p(\hat{\rho}|H_0){\rm d}\hat{\rho}. \tag{1}$$

*jeffrey.hazboun@oregonstate.edu
†joseph.romano@utrgv.edu

The false alarm probability α is similarly defined:

$$\alpha \equiv \mathrm{Prob}(\hat{\rho} > \hat{\rho}_{\mathrm{th}}|H_0) \equiv \int_{\hat{\rho}_{\mathrm{th}}}^{\infty} p(\hat{\rho}|H_0)\mathrm{d}\hat{\rho}, \qquad (2)$$

where $\hat{\rho}_{\mathrm{th}}$ is the detection threshold, above which one would reject the null hypothesis and claim detection of a GWB. The detection threshold is typically chosen so that the false alarm probability has a sufficiently small value, e.g., $\alpha < 1\times 10^{-3}$ above that threshold. Note that the right-hand sides of both (1) and (2) can also be written as $1 - \mathrm{CDF}(x|H_0)$, where $\mathrm{CDF}(x)$ is the standard cumulative distribution function and $x = \hat{\rho}_{\mathrm{obs}}$ and $\hat{\rho}_{\mathrm{th}}$, respectively.

By construction, both the p-value and the false alarm probability (or, equivalently, the choice of the detection threshold) depend on the detailed form of the null distribution. In previous work [6–8], appealing to the central limit theorem (and also for simplicity), this distribution was assumed to be Gaussian. Here, we show that the null distribution for the OS follows a generalized chi-squared (GX2) distribution [10], and we explore the consequences of this observation. (A paper by Cordes and Shannon [11] also noted that the distribution of the cross-correlation statistic was highly non-Gaussian and skewed, but did not identify it as a GX2 distribution.) We show that differences between the GX2 and Gaussian distributions can be significant for current pulsar timing array configurations (defined by the numbers of pulsars, observation spans, noise parameters, etc.), as well as the amplitude and spectral shape of the common-spectrum process, especially in the tails of the distribution. Thus, the GX2 distribution should be used to calculate more accurate p-values in the case of a GW detection. In particular, we show that the Gaussian distribution assumption for the null distribution of the OS leads to overestimates of the significance of a potential detection, i.e., smaller p-values, than for the GX2 distribution.

We also compare the GX2 and Gaussian distributions of the OS to the null distributions obtained from phase shifts [12] and sky scrambles [13] of the PTA data. The phase shifts are analogous to applying time shifts in ground-based GW detectors, while sky scrambles replace the actual pulsar locations with random sky locations, thus removing the dependence of the spatial correlations on the angular separation between pairs of pulsars. Both of these techniques work well at removing GW-induced correlations and are now standard methods to determine null distributions for our statistics, be it the OS or the Bayes factor, directly from our data [14]. We show that the GX2 distribution is an excellent fit to the phase-shifted or sky-scrambled data compared with the standard Gaussian approximation, especially in the tails of the distribution (where it matters most) at larger values of the OS signal-to-noise ratio $\hat{\rho}$.

The rest of the paper is organized as follows: In Sec. II, we summarize the results of [10], explaining how GX2 distributions arise whenever one has a (symmetric) quadratic combination of random variables satisfying a multivariate Gaussian distribution. We then show, in Sec. III, that the OS used for PTA searches is an example of such a quadratic combination, explicitly describing the various pieces that enter the calculation of the eigenvalues needed for the GX2 distribution. In Sec. IV, we first calculate GX2 distributions for the OS for several different sets of simulated data, in the absence of simulated GW-induced cross correlations. We show how these distributions depend on the relative contribution of red and white noise, the number of pulsars, fitting to a timing model, etc., and compare these distributions with unit (i.e., standard normal) Gaussians, which often deviate significantly from the GX2 distributions in the tails of the distributions. We then compare the GX2 distribution for the OS with the null distribution of $\hat{\rho}$ obtained by phase shifting the NANOGrav 12.5-year data [14]. Finally, we conclude in Sec. V by discussing possible extensions or modifications to the calculations presented here—e.g., which might simplify the calculation of GX2 distributions for realistic PTA data sets. The Appendix describes a simple "tail-fitting" approach that allows us to extrapolate the tail of the OS null distribution beyond empirically determined phase-shift or sky-scramble values.

II. MATHEMATICAL FORMALISM

By definition, the GX2 distribution is the probability distribution of a quadratic form of multivariate-Gaussian random variables $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$:

$$q(\mathbf{x}) \equiv \frac{1}{2}\mathbf{x}^T\mathbf{Q}_2\mathbf{x} + \mathbf{q}_1^T\mathbf{x} + q_0, \qquad (3)$$

where

$$p(\mathbf{x}|\boldsymbol{\mu},\boldsymbol{\Sigma}) = \frac{1}{\sqrt{\det(2\pi\boldsymbol{\Sigma})}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}. \qquad (4)$$

Here, q_0 is a real constant, $\mathbf{q}_1$ is a real vector of the same dimension as $\mathbf{x}$, and $\mathbf{Q}_2$ is a real symmetric matrix $\mathbf{Q}_2^T = \mathbf{Q}_2$. As we shall describe in Sec. III, the cross-correlation statistic that we are interested in has the simpler form

$$q(\mathbf{x}) \equiv \frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x}, \qquad \mathbf{x} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}). \qquad (5)$$

In this last equation, we have dropped the subscript 2 from $\mathbf{Q}_2$ to simplify the notation since there is no chance of confusing it with a linear or constant term. We will work with this simpler form for the rest of the paper.

To use the formalism of [10] to explicitly compute the GX2 distribution of $q(\mathbf{x})$, we need to write (5) as a linear superposition of standard normal distributions $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$. This is done via a series of eigenvalue-eigenvector decompositions which we summarize below.

Following [10], we start by converting $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$ to a vector of uncorrelated standard normal distributions $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ by finding the eigenvalues and eigenvectors of $\mathbf{\Sigma}$:

$$\mathbf{D} = \mathbf{E}^T \mathbf{\Sigma} \mathbf{E} \quad \text{or} \quad \mathbf{\Sigma} = \mathbf{E} \mathbf{D} \mathbf{E}^T, \tag{6}$$

where $\mathbf{D}$ is a diagonal matrix of eigenvalues $\mathbf{D} = \mathrm{diag}(\sigma_1^2, \sigma_2^2, \cdots)$ and $\mathbf{E}$ is an orthogonal matrix (i.e., $\mathbf{E}^T = \mathbf{E}^{-1}$) whose columns are the corresponding (orthonormal) eigenvectors of $\mathbf{\Sigma}$. Since $\mathbf{\Sigma}$ is a covariance matrix, we are guaranteed that its eigenvalues are all positive, hence the form σ_1^2, σ_2^2, …. We then take the square root of $\mathbf{D}$:

$$\mathbf{\Lambda} \equiv \sqrt{\mathbf{D}}, \tag{7}$$

for which

$$\mathbf{x}^T \mathbf{\Sigma}^{-1} \mathbf{x} = \mathbf{x}^T \mathbf{E} \mathbf{D}^{-1} \mathbf{E}^T \mathbf{x} = \mathbf{x}^T \mathbf{E} \mathbf{\Lambda}^{-1} \mathbf{\Lambda}^{-1} \mathbf{E}^T \mathbf{x} = \mathbf{z}^T \mathbf{1} \mathbf{z}, \tag{8}$$

where

$$\mathbf{z} \equiv \mathbf{\Lambda}^{-1} \mathbf{E}^T \mathbf{x} \Leftrightarrow \mathbf{x} = \mathbf{E} \mathbf{\Lambda} \mathbf{z}. \tag{9}$$

In terms of $\mathbf{z}$ the quadratic combination (5) has the form

$$q(\mathbf{x}) = \frac{1}{2} \mathbf{z}^T \tilde{\mathbf{Q}} \mathbf{z}, \qquad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{1}), \tag{10}$$

where

$$\tilde{\mathbf{Q}} \equiv \mathbf{\Lambda}^T \mathbf{E}^T \mathbf{Q} \mathbf{E} \mathbf{\Lambda}. \tag{11}$$

The final step is to diagonalize $\tilde{\mathbf{Q}}$, by finding its eigenvalues $(\tilde{e}_1, \tilde{e}_2, \cdots)$ and the orthogonal matrix of eigenvectors $\mathbf{U}$. This gives

$$\tilde{\mathbf{Q}} = \mathbf{U} \mathrm{diag}(\tilde{e}_1, \tilde{e}_2, \ldots) \mathbf{U}^T, \tag{12}$$

for which

$$q(\mathbf{x}) = \frac{1}{2} \sum_i \tilde{e}_i v_i^2, \quad \text{where } \mathbf{v} \equiv \mathbf{U}^T \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{1}). \tag{13}$$

(Note that $\mathbf{v}$ is standard normal since $\mathbf{U}$ is an orthogonal matrix and $\mathbf{z}$ is standard normal.) Since the probability distribution of the square of a standard normal distribution is chi-squared distributed with 1 degree of freedom (DOF), it follows that $q(\mathbf{x})$ is a general linear combination of χ_1^2 distributions, which is the form of the generalized chi-squared distribution discussed in Ref. [10]. The analytic form of this distribution for the quadratic form (5) is completely specified by the eigenvalues of $\tilde{\mathbf{Q}}$ defined by (6), (7), and (11).

Note that the mean and variance of $q(\mathbf{x})$ can be simply written in terms of sums (and sums of squares) of the eigenvalues used to construct the optimal statistic:

$$\mu_q \equiv \langle q \rangle = \frac{1}{2} \sum_i \tilde{e}_i, \quad \sigma_q^2 \equiv \langle q^2 \rangle - \langle q \rangle^2 = \frac{1}{2} \sum_i \tilde{e}_i^2, \tag{14}$$

where $\langle\rangle$ denotes the expectation value. The above results follow from the expansion (13) with $v_i^2 \sim \chi_1^2$ being statistically independent of one another, and each having mean = 1 and variance = 2.

III. APPLICATION TO THE OPTIMAL STATISTIC FOR PTA SEARCHES FOR GWBs

Now we show that different forms of the OS used for PTA searches for GWBs are examples of GX2 distributions. For more details regarding the OS, we refer the reader to [7], with which this work shares notation, and also [6,8].

A. Optimal statistic signal-to-noise ratio and GWB amplitude estimator

The optimal statistic signal-to-noise ratio for PTA searches for GWBs is typically written as

$$\hat{\rho} \equiv \hat{A}_{\mathrm{gw}}^2 / \sigma_0, \tag{15}$$

where

$$\hat{A}_{\mathrm{gw}}^2 \equiv \mathcal{N} \sum_{a<b} \mathbf{r}_a^T \mathbf{P}_a^{-1} \tilde{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \mathbf{r}_b, \tag{16}$$

$$\sigma_0 \equiv \mathcal{N}^{1/2}, \tag{17}$$

$$\mathcal{N} \equiv \left(\sum\nolimits_{a<b} \mathrm{tr}[\mathbf{P}_a^{-1} \tilde{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \tilde{\mathbf{S}}_{ba}] \right)^{-1}. \tag{18}$$

In the above expressions, $\hat{A}_{\mathrm{gw}}^2$ is an estimator of the squared amplitude of the GW signal, σ_0^2 is its variance in the absence of GW-induced spatial correlations, and $\mathcal{N}$ is a normalization factor constructed from terms involving the total autocorrelated power and cross-correlated power in pulsars labeled by a, b (more about these expressions below). More compactly,

$$\hat{\rho} \equiv \sum_{a<b} \mathbf{r}_a^T \mathbf{Q}_{ab} \mathbf{r}_b = \frac{1}{2} \sum_{a,b} \mathbf{r}_a^T \mathbf{Q}_{ab} \mathbf{r}_b, \tag{19}$$

where

$$\mathbf{Q}_{aa} \equiv \mathbf{0}, \qquad \mathbf{Q}_{ab} \equiv \mathcal{N}^{1/2}\mathbf{P}_a^{-1}\tilde{\mathbf{S}}_{ab}\mathbf{P}_b^{-1} \tag{20}$$

define the (symmetric) quadratic form for $\hat{\rho}$. Note that summations denoted by $\sum_{a<b}$ run over *distinct* pulsar pairs, while $\sum_{a,b} \equiv \sum_a \sum_b$ double counts the pulsar pairs [hence the factor of $1/2$ in (19)] and also includes the autocorrelations (which do not contribute to $\hat{\rho}$ since $\mathbf{Q}_{aa} = \mathbf{0}$).

The data $\mathbf{r}$ are zero-mean multivariate Gaussian random variables defined by

$$p(\mathbf{r}|\vec{\theta}) = \frac{1}{\sqrt{\det(2\pi\boldsymbol{\Sigma})}}\exp\left(-\frac{1}{2}\mathbf{r}^T\boldsymbol{\Sigma}^{-1}\mathbf{r}\right), \tag{21}$$

where

$$\boldsymbol{\Sigma} \equiv \langle \mathbf{r}\mathbf{r}^T \rangle = \begin{pmatrix} \mathbf{P}_1 & \mathbf{S}_{12} & \cdots & \mathbf{S}_{1M} \\ \mathbf{S}_{21} & \mathbf{P}_2 & \cdots & \mathbf{S}_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{S}_{M1} & \mathbf{S}_{M2} & \cdots & \mathbf{P}_M \end{pmatrix} \tag{22}$$

and

$$\mathbf{P}_a \equiv \mathbf{G}_a^T\mathbf{N}_a\mathbf{G}_a, \qquad \mathbf{S}_{ab} \equiv \mathbf{G}_a^T\mathbf{X}_{ab}\mathbf{G}_b,$$
$$a, b = 1, 2, \ldots, M. \tag{23}$$

Here, M denotes the number of pulsars, and $\mathbf{G}_a$ is the G-matrix [15] for pulsar a, which encodes information about the timing-model fit

$$\mathbf{r}_a = \mathbf{G}_a^T\boldsymbol{\delta}\mathbf{t}_a, \quad a = 1, 2, \ldots, M, \tag{24}$$

where $\boldsymbol{\delta}\mathbf{t}_a$ are the timing residuals for pulsar a. If we denote the number of TOAs for pulsar a by $N_{\mathrm{TOA},a}$ and the number of timing model parameters by $N_{\mathrm{par},a}$, then $\mathbf{G}_a$ has dimensions $N_{\mathrm{TOA},a} \times (N_{\mathrm{TOA},a} - N_{\mathrm{par},a})$, and $\mathbf{r}_a$ is a vector with components

$$[\mathbf{r}_a]_{\alpha_a}, \quad \alpha_a = 1, 2, \ldots, N_a \equiv (N_{\mathrm{TOA},a} - N_{\mathrm{par},a}). \tag{25}$$

The covariance matrix $\boldsymbol{\Sigma}$ is thus a symmetric block matrix and has overall dimension $(N_1 + N_2 + \cdots + N_M) \times (N_1 + N_2 + \cdots + N_M)$.

The diagonal terms of the covariance matrix involve the autocorrelations

$$\mathbf{N}_a \equiv \langle \boldsymbol{\delta}\mathbf{t}_a\boldsymbol{\delta}\mathbf{t}_a^T \rangle = \int_0^{f_{\mathrm{Nyq}}} \mathrm{d}f \cos[2\pi f\boldsymbol{\tau}_{aa}]\mathcal{P}_a(f) + \mathcal{F}_a\mathbf{W}_a + \mathcal{Q}_a^2\mathbf{1}, \tag{26}$$

where the last two terms specify the white-noise contributions, and

$$\mathcal{P}_a(f) \equiv \mathcal{P}_a^{\mathrm{red}}(f) + \mathcal{P}_{\mathrm{gw}}(f) \tag{27}$$

consists of both intrinsic pulsar red-noise and a potential common-spectrum red-noise process contribution most likely from the GWB. We assume that both of these red-noise contributions can be described by power-law spectra

$$\mathcal{P}_a^{\mathrm{red}}(f) \equiv \frac{A_a^2}{12\pi^2 f^3}\left(\frac{f}{f_{\mathrm{ref}}}\right)^{2\alpha_a},$$
$$\mathcal{P}_{\mathrm{gw}}(f) \equiv \frac{A_{\mathrm{gw}}^2}{12\pi^2 f^3}\left(\frac{f}{f_{\mathrm{ref}}}\right)^{2\alpha_{\mathrm{gw}}}. \tag{28}$$

For the GWB formed from the superposition of signals from inspiraling supermassive binary black holes in the centers of merging galaxies, $\alpha_{\mathrm{gw}} = -2/3$. Finally, $\boldsymbol{\tau}_{aa}$ is the time-lag matrix, whose components are given by $[\boldsymbol{\tau}_{aa}]_{ij} \equiv t_{i_a} - t_{j_a}$, which are the differences of the TOAs of the pulses from pulsar a.

The off-diagonal terms in the covariance matrix are assumed to have only a GWB contribution,

$$\mathbf{X}_{ab} \equiv \langle \boldsymbol{\delta}\mathbf{t}_a\boldsymbol{\delta}\mathbf{t}_b^T \rangle = \Gamma_{ab}\int_0^{f_{\mathrm{Nyq}}} \mathrm{d}f \cos[2\pi f\boldsymbol{\tau}_{ab}]\mathcal{P}_{\mathrm{gw}}(f), \tag{29}$$

where

$$\Gamma_{ab} \equiv \frac{1}{2} + \frac{3}{2}\left(\frac{1-\cos\xi_{ab}}{2}\right)\left[\ln\left(\frac{1-\cos\xi_{ab}}{2}\right) - \frac{1}{6}\right] + \frac{1}{2}\delta_{ab} \tag{30}$$

are the values of the Hellings-and-Downs function, $\Gamma_{ab} \equiv \Gamma(\xi_{ab})$, evaluated for two pulsars a and b separated by the angle ξ_{ab} (see [5]). The quantity

$$\tilde{\mathbf{S}}_{ab} \equiv \mathbf{G}_a^T\tilde{\mathbf{X}}_{ab}\mathbf{G}_b, \tag{31}$$

which enters the expression for the quadratic form $\mathbf{Q}_{ab}$, (20), is a normalized version of $\mathbf{S}_{ab}$ defined in terms of

$$\tilde{\mathbf{X}}_{ab} \equiv \mathbf{X}_{ab}/A_{\mathrm{gw}}^2. \tag{32}$$

Note that these cross-correlations are proportional to the spectral shape of the GWB—i.e., they do not depend on its amplitude.

Finally, it is a simple matter to show that the GWB amplitude estimator $\hat{A}_{\mathrm{gw}}^2$ can also be written as a (symmetric) quadratic combination of the multivariate Gaussian random variables $\mathbf{r}$ with quadratic form

$$\mathbf{K}_{ab} \equiv \mathcal{N}^{1/2}\mathbf{Q}_{ab}. \tag{33}$$

Thus, according to the discussion in Sec. II, both $\hat{\rho}$ and $\hat{A}_{\mathrm{gw}}^2$ will be described by GX2 distributions. For calculating the distributions of these statistics in the absence of

GW-induced spatial correlations (i.e., null distributions), we should set $\mathbf{X}_{ab} = \mathbf{0}$ in the definition of $\mathbf{\Sigma}$ and replace the GW contribution $\mathcal{P}_{\rm gw}(f)$ to $\mathbf{N}_a$ by a potential common-spectrum red-noise process $\mathcal{P}_{\rm cp}(f)$ (with amplitude $A_{\rm cp}$ and spectral index $\alpha_{\rm cp}$), which is common to all pulsars. (The normalized cross-correlation terms $\tilde{\mathbf{S}}_{ab}$ in $\mathbf{Q}_{ab}$ do not change since they involve the normalized cross-correlations $\tilde{\mathbf{X}}_{ab}$.) These null distributions are obviously also described by GX2 distributions since they are special cases of the non-null quadratic combinations.

Using (14), it follows that the eigenvalues $\tilde{e}_i$ that specify the GX2 distribution for $\hat{\rho} \equiv \hat{A}^2_{\rm gw}/\sigma_0$ in the absence of GW-induced spatial correlations satisfy

$$\sum_i \tilde{e}_i = 0 \quad \text{and} \quad \sum_i \tilde{e}_i^2 = 2. \tag{34}$$

These are consequences of $\hat{\rho}$ having zero mean and unit variance for the null distribution case.

B. Optimal statistic pulsar-pair cross-correlation estimators

One can also construct cross-correlation estimators for individual pulsar pairs:

$$\hat{\rho}_{ab} \equiv \mathcal{N}_{ab}\mathbf{r}_a^T\mathbf{P}_a^{-1}\bar{\mathbf{S}}_{ab}\mathbf{P}_b^{-1}\mathbf{r}_b \equiv \mathbf{r}_a^T\bar{\mathbf{Q}}_{ab}\mathbf{r}_b, \quad a < b = 1, 2, \ldots, M, \tag{35}$$

where

$$\bar{\mathbf{Q}}_{ab} \equiv \mathcal{N}_{ab}\mathbf{P}_a^{-1}\bar{\mathbf{S}}_{ab}\mathbf{P}_b^{-1}, \tag{36}$$

$$\mathcal{N}_{ab} \equiv \left(\mathrm{tr}\left[\mathbf{P}_a^{-1}\bar{\mathbf{S}}_{ab}\mathbf{P}_b^{-1}\bar{\mathbf{S}}_{ba}\right]\right)^{-1}, \tag{37}$$

$$\bar{\mathbf{S}}_{ab} \equiv \mathbf{G}_a^T\bar{\mathbf{X}}_{ab}\mathbf{G}_b, \tag{38}$$

$$\bar{\mathbf{X}}_{ab} \equiv \mathbf{X}_{ab}/(\Gamma_{ab}A^2_{\rm gw}) = \int_0^{f_{\rm Nyq}} \mathrm{d}f\, \cos[2\pi f\boldsymbol{\tau}_{ab}]\bar{\mathcal{P}}_{\rm gw}(f), \tag{39}$$

$$\bar{\mathcal{P}}_{\rm gw}(f) \equiv \frac{1}{12\pi^2 f^3}\left(\frac{f}{f_{\rm ref}}\right)^{2\alpha_{\rm gw}}. \tag{40}$$

Using

$$\left\langle \mathbf{r}_a\mathbf{r}_b^T\right\rangle = \mathbf{S}_{ab}, \tag{41}$$

it follows that $\hat{\rho}_{ab}$ is an unbiased estimator of the cross-correlated power in the GWB—i.e.,

$$\langle\hat{\rho}_{ab}\rangle = \Gamma_{ab}A^2_{\rm gw}, \tag{42}$$

with variance

$$\sigma^2_{ab} = \mathcal{N}_{ab} + \left(A^2_{\rm gw}\right)^2\mathcal{N}^2_{ab}\mathrm{tr}\left[\mathbf{P}_a^{-1}\tilde{\mathbf{S}}_{ab}\mathbf{P}_b^{-1}\tilde{\mathbf{S}}_{ba}\mathbf{P}_a^{-1}\bar{\mathbf{S}}_{ab}\mathbf{P}_b^{-1}\bar{\mathbf{S}}_{ba}\right]. \tag{43}$$

In the absence of GW-induced spatial correlations, the variance simplifies to

$$\sigma^2_{ab} = \mathcal{N}_{ab} \equiv \sigma^2_{0,ab}. \tag{44}$$

Recall that Γ_{ab} are the values of the Hellings and Downs function evaluated for different angular separations between the two pulsars labeled by a and b. The matrix $\bar{\mathbf{X}}_{ab}$ is the time-domain representation of the spectral shape of the GW power spectrum, which (by its definition) is independent of the GWB amplitude $A_{\rm gw}$ and the spatial correlation coefficients Γ_{ab}. It depends on the pulsar pair ab only via the discrete times t_{i_a}, t_{j_b} of the timing residuals $\boldsymbol{\delta}\mathbf{t}_a$, $\boldsymbol{\delta}\mathbf{t}_b$, which enter the time-lag matrix $\boldsymbol{\tau}_{ab}$. This means that $\bar{\mathbf{X}}_{ab}$ is a rectangular matrix with dimensions $N_{{\rm TOA},a} \times N_{{\rm TOA},b}$.

Since $\hat{\rho}_{ab} = \hat{\rho}_{ba}$, we can write

$$\hat{\rho}_{ab} = \frac{1}{2}(\hat{\rho}_{ab} + \hat{\rho}_{ba}) = \frac{1}{2}\left(\mathbf{r}_a^T\bar{\mathbf{Q}}_{ab}\mathbf{r}_b + \mathbf{r}_b^T\bar{\mathbf{Q}}_{ba}\mathbf{r}_a\right) = \frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x}, \tag{45}$$

where

$$\mathbf{x}^T \equiv \begin{bmatrix}\mathbf{r}_a^T & \mathbf{r}_b^T\end{bmatrix}, \quad \mathbf{Q} \equiv \begin{bmatrix}\mathbf{0} & \bar{\mathbf{Q}}_{ab}\\ \bar{\mathbf{Q}}_{ba} & \mathbf{0}\end{bmatrix}, \quad \mathbf{x} \equiv \begin{bmatrix}\mathbf{r}_a\\ \mathbf{r}_b\end{bmatrix}. \tag{46}$$

Note that $\mathbf{Q}$ is a symmetric matrix since $\bar{\mathbf{Q}}_{ab}^T = \bar{\mathbf{Q}}_{ba}$. Note also that $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$, where

$$\mathbf{\Sigma} \equiv \left\langle \mathbf{x}\mathbf{x}^T\right\rangle = \begin{bmatrix}\mathbf{P}_a & \mathbf{S}_{ab}\\ \mathbf{S}_{ba} & \mathbf{P}_b\end{bmatrix}. \tag{47}$$

Thus, given that $\hat{\rho}_{ab}$ is a symmetric quadratic combination of multivariate Gaussian random variables, it obeys a GX2 distribution.

If desired, one can express both the OS signal-to-noise ratio $\hat{\rho}$ and the GWB amplitude estimator $\hat{A}^2_{\rm gw}$ in terms of the pulsar-pair cross-correlation estimators $\hat{\rho}_{ab}$:

$$\hat{\rho} = \frac{\sum_{a<b}\Gamma_{ab}\hat{\rho}_{ab}/\sigma^2_{0,ab}}{\sqrt{\sum_{c<d}\Gamma^2_{cd}/\sigma^2_{0,cd}}}, \quad \hat{A}^2_{\rm gw} = \frac{\sum_{a<b}\Gamma_{ab}\hat{\rho}_{ab}/\sigma^2_{0,ab}}{\sum_{c<d}\Gamma^2_{cd}/\sigma^2_{0,cd}}. \tag{48}$$

These results are a consequence of $\tilde{\mathbf{S}}_{ab} = \Gamma_{ab}\bar{\mathbf{S}}_{ab}$ and the definitions (15), (16), and (35) of $\hat{\rho}$, $\hat{A}^2_{\rm gw}$, and $\hat{\rho}_{ab}$. Thus, $\hat{\rho}$ and $\hat{A}^2_{\rm gw}$ are simple noise-weighted and Γ_{ab}-matched linear combinations of $\hat{\rho}_{ab}$.

IV. DEMONSTRATION OF THE GX2 DISTRIBUTION ON PTA DATA

The exact form of the GX2 distribution varies with the amplitude of the GW signal and the noise properties of the data. The characteristics of PTA noise have been studied extensively in the literature [16–18]. The considerations for PTA detector characterization vary widely from radiometer noise at the telescope receiver to possible intrinsic spin instabilities of the neutron star. However, they can be split into a few categories including white noise, red noise, the fit of the deterministic timing model for the radio pulse times of arrival, and the number and sky position of the pulsars. In this section we demonstrate the influence that these characteristics of PTAs have on the GX2 distribution.

Fig. 1 shows various GX2 distributions constructed from the parameters defining some simple (fictitious) PTAs, with the spectral properties of the pulsars calculated using the pulsar spectral characterization software HASASIA [19]. The panels demonstrate how differences in the GW signal (spatial correlations, spectral index), noise, and timing parameters change the shape of the distribution. Note that we have not included a comparison of different white-noise levels because the noise power spectral densities cancel out of the numerator and denominator of $\tilde{\mathbf{Q}}$, which defines the GX2 distribution.[1]

The interplay of the parameters makes it difficult to predict exactly what the distribution will look like. However, the trends for any individual parameter follow the simple rule that as the sensitivity of the PTA is increased, the tail of the GX2 distribution becomes smaller at larger values of $\hat{\rho}$. A couple other general observations are as follows: (i) adding sufficiently large red noise to the pulsars quickly obscures the subtle differences in the GX2 distributions, as seen in panel (b) of Fig. 1, and (ii) the sky positions of the pulsars matter via the $\tilde{\mathbf{S}}_{ab}$ terms in the expression for $\hat{\rho}$.

Next, we present a realistic GX2 distribution calculated using the noise and sensitivity parameters of the NANOGrav 12.5-year dataset [20,21] to demonstrate the usefulness of an accurate analytic GX2 distribution in calculating p-values or false-alarm probabilities. Diagonalizing the various matrices described in Sec. III over full PTA data sets is challenging because of the length of the data sets. Nonetheless, by using the salient noise and sensitivity parameters of the NANOGrav 12.5 data set, we obtain fairly reasonable agreement with the phase-shift method for determining the null distribution of the optimal statistic. This is illustrated in Fig. 2, which compares the analytic GX2 distribution with a histogram of $\hat{\rho}$ values for 1000 different phase shifts of the NANOGrav 12.5-year data. For reference, we also show the unit (standard normal) Gaussian distribution.

[1]See (11), (20), and (22), noting that $\mathbf{\Sigma}$ is block diagonal for the null distribution case.

By looking at the right-hand panel of Fig. 2, one immediately sees the inaccuracy that would arise if the p-value was calculated assuming that the null distribution was Gaussian. Table I gives the p-values for $\hat{\rho} = 5$ and $\hat{\rho} = 1.3$ calculated using the analytic form of the GX2 distribution and the Gaussian distribution, and from phase shifts and sky scrambles of the NANOGrav 12.5-year data, the latter as described in [12–14]. (The value $\hat{\rho} = 1.3$ is what was measured in the NANOGrav 12.5-year data set.) While the agreement between the various methods[2] is reasonable for the lower value of $\hat{\rho}$, assuming that the null distribution is Gaussian for $\hat{\rho} = 5$ leads to a p-value that is more than 1000 times smaller than it should be.

V. DISCUSSION

We have demonstrated that a GX2 distribution is the correct analytical distribution for the optimal cross-correlation statistic (OS) used for analyzing PTA data sets. Although we focused on the null distribution of the OS signal-to-noise ratio $\hat{\rho}$ for this paper, our analyses in Secs. III A and III B were sufficiently general to show that GX2 distributions also apply to the optimal estimator of the squared-amplitude $\hat{A}^2_{\rm gw}$ and pulsar-pair cross-correlations $\hat{\rho}_{ab}$ in the presence of a signal.

We applied the general formalism to calculate the GX2 distribution for $\hat{\rho}$ using parameters appropriate for NANOGrav's 12.5-year data set and showed that it agreed quite well with the empirical null distribution that was obtained by phase shifting the NANOGrav data. We also calculated the GX2 distribution for several different sets of simulated data in the absence of a GWB cross-correlation, varying, in turn, the number of pulsars, the relative contribution of red and white noise, etc., to see how these affected the shape of the resulting distributions.

Generically, the GX2 distributions we obtained differed from the best-fit (standard normal) Gaussian distribution by having a mode less than their mean and having "fatter" tails at high values of the statistic. (Both distributions have zero mean and unit variance for the null distribution case.) The fatter tails are especially important when calculating p-values for the null distribution, which is needed to assess the statistical significance of a possible detection.

As mentioned in Sec. IV, constructing the quadratic form for the GX2 distribution—which requires solving for the eigenvectors of large matrices—is challenging for realistic data. Using the full NANOGrav 12.5-year data set would require solving the eigenvalue problem for an $N = 410064$ dimensional matrix twice. An additional complication is that the detailed dispersion-measure-variation (DMX) model that NANOGrav uses [23] has a large effect on the transmission function of the pulsars [18,24]. This shows up

[2]Note that sky scrambles are currently being revisited as a robust method for retrieving the null distribution empirically [22]. We have included the values here for historical completeness.

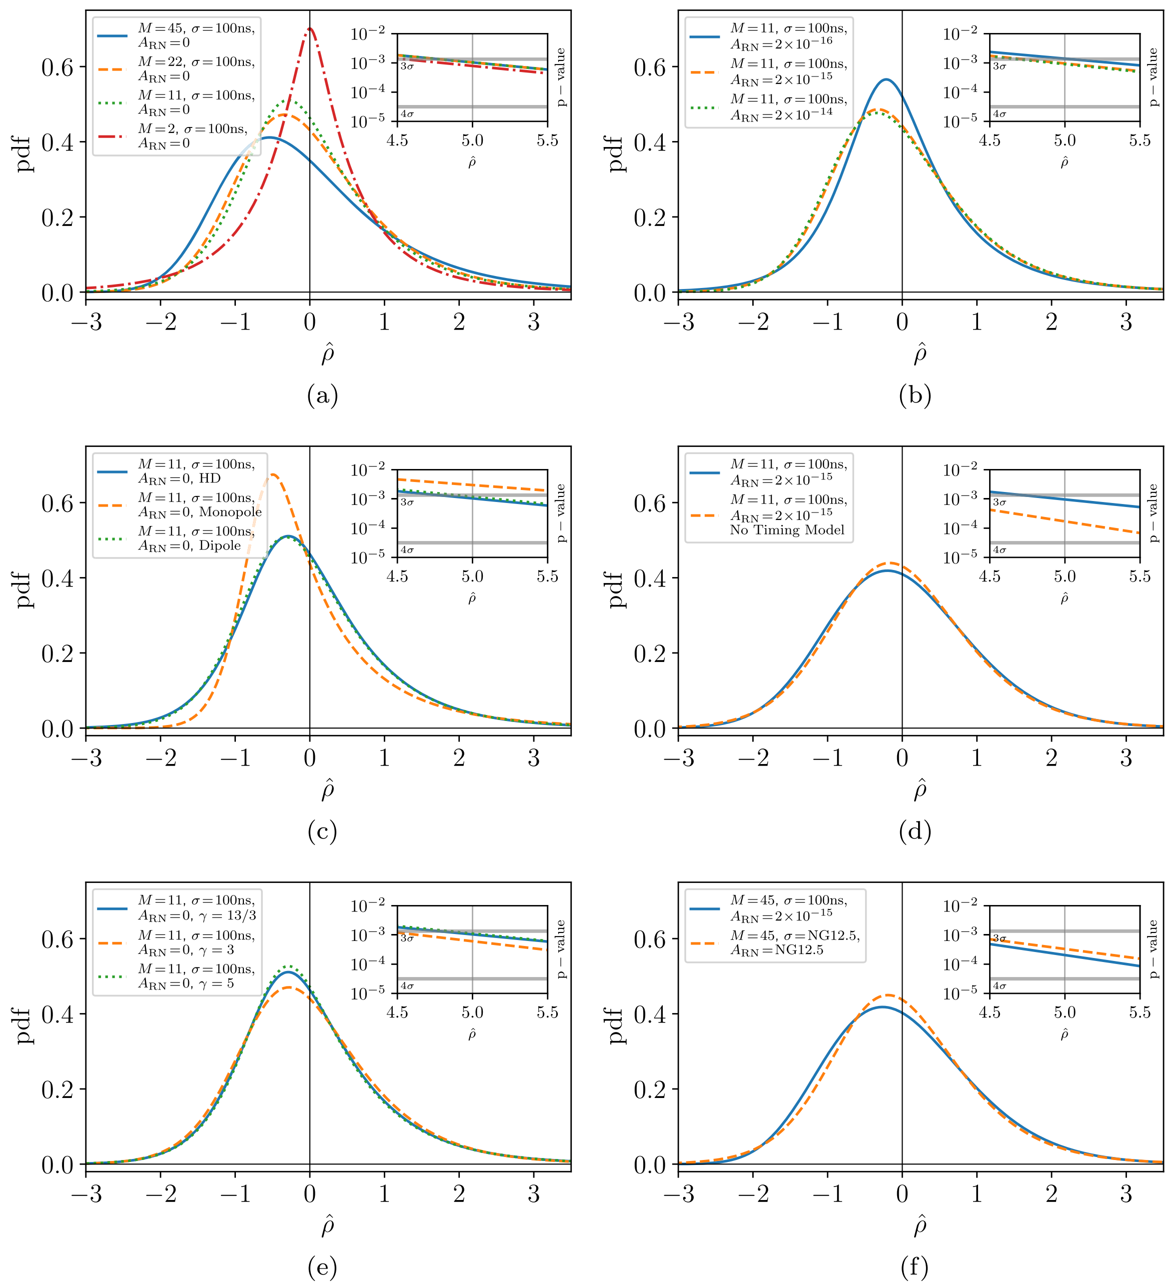

FIG. 1. Comparison of GX2 distributions when varying different PTA parameters: (a) varying the number of pulsars for fixed pulsar white noise, (b) varying the amplitude of the common-spectrum red-noise process for fixed pulsar white noise and number of pulsars, (c) comparing distributions with different spatial correlations for fixed pulsar white noise and number of pulsars, (d) comparing distributions having a nontrivial and trivial (i.e., identity) timing models, (e) varying the spectral index of the GWB search (here written in terms of the spectral index in timing residuals, $\gamma \equiv 3 - 2\alpha$) for fixed white noise and number of pulsars, and (f) comparing a set of realistic NANOGrav parameters, with a simple white-noise plus common-spectrum red-noise process. The insets show plots of the p-values of the various GX2 distributions as functions of $\hat{\rho}$. The gray bands show the p-values for traditional 3σ and 4σ detection significances based on a unit (standard normal) Gaussian distribution. Note that the 11, 22, and 45 pulsar cases all fall along the same p-values in (a).

in the timing model fit, which enters the quadratic form via the G-matrix, e.g., $\mathbf{r}_a = \mathbf{G}_a^T \boldsymbol{\delta}\mathbf{t}_a$.

As such, the key to constructing a valid GX2 distribution for a realistic PTA data set is to find a reduced set of parameters that faithfully describes the spectral properties of the PTA pulsars and the corresponding timing model of the array. We used HASASIA to calculate the spectra for all of the NG12.5 pulsars, using their full data sets [18]. The noise power spectral density (from the spectrum) and red-noise parameters (from the Bayesian noise analyses) were

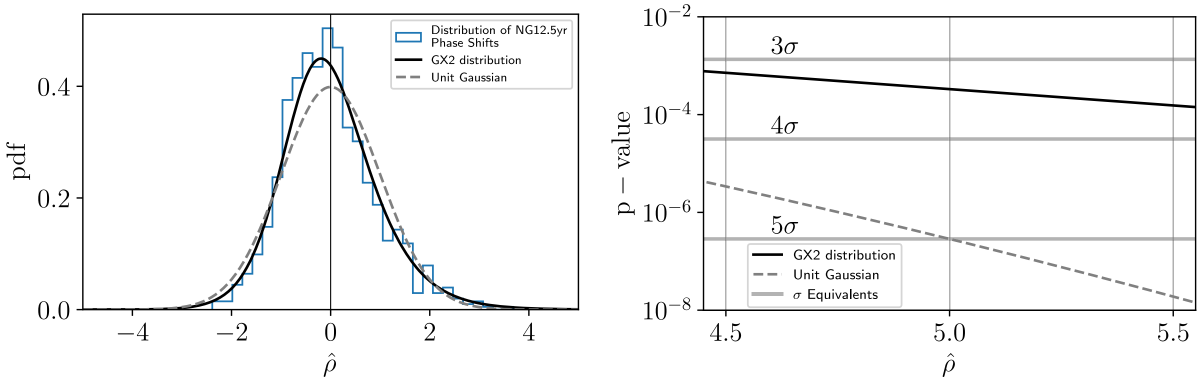

FIG. 2. Comparison of various distributions of the OS signal-to-noise ratio $\hat{\rho}$ in the absence of GW-induced spatial correlations. Shown are (i) an empirical null distribution for the NANOGrav 12.5-year data, obtained from performing 1000 phase shifts (blue histogram), (ii) the analytic GX2 distribution (black solid line), and (iii) the best-fit (standard normal) Gaussian distribution (gray dashed line). The right panel is a plot of the p-value as a function of $\hat{\rho}$. The phase shifts do not show up in the right-panel plot out to these large values of $\hat{\rho}$.

then used to construct shorter data sets with similar properties, including the broadband effect of the DMX model. Spectra of these simulated data sets were taken to iteratively find the correct effective level of white-noise power spectral density to inject into the pulsars to match the spectra from the full data sets.

This process proved expeditious to obtain a fairly accurate realization of a GX2 distribution on a modest laptop. However, it is computationally inefficient for generating GX2 distributions for several different choices of noise parameters, for example. Fortunately, PTA calculations are usually carried out using a rank-reduced formalism [25] that drastically reduces the size of matrices dealt with in the analysis. This reduced representation of the data alleviates the main problem discussed previously. For example, we are currently developing techniques to use a frequency-domain implementation of the optimal statistic to speed up the calculation, taking advantage of the rank-reduced matrix, which is only $2N_{\mathrm{freq}} \times 2N_{\mathrm{freq}}$, as opposed to $N_{\mathrm{TOA}} \times N_{\mathrm{TOA}}$. However, we leave that discussion for future work.

An alternative to using noise estimates from a previous Bayesian inference run to calculate the optimal statistic is to marginalize over the noise parameters. As described in [8], marginalizing over the red-noise parameters tends to remove biases that would otherwise exist due to correlations between the noise estimates and the timing residual data used to construct the optimal statistic. For the analyses described in this paper, we used maximum-likelihood estimates of noise parameters from a Bayesian analysis to construct the quadratic forms needed for calculating the GX2 distributions. We did not investigate any source of bias that might have been introduced by using noise estimates as opposed to noise marginalization. However, we are currently investigating the possibility of marginalizing over the noise for future uses of the GX2 distributions.

TABLE I. The p-values calculated using various methods in the context of the NANOGrav 12.5-year data set.

$\hat{\rho}$	Analytic GX2	Gaussian	Phase shifts	Sky scrambles
5	3.3×10^{-4}	2.87×10^{-7}	⋯	⋯
1.3	0.0983	0.0951	0.091 [14]	0.082 [14]

Finally, optimal cross-correlation statistics are also used when analyzing data from ground-based GW detectors like Advanced LIGO, Virgo, and KAGRA. However, for this case, the optimal statistics are well described by Gaussian distributions, so GX2 distributions are not needed. This is because the data from pairs of detectors are analyzed in roughly 100-sec segments (to account for potential nonstationarities in the detector noise power) and then averaged together over $\gtrsim 10^5$ such segments, corresponding to a typical year-long observation. The cross-correlation estimates of the amplitude of the GWB for each 100-sec segment are GX2 distributed. However, the final averaged optimal-statistic value (inverse-noise weighted by the variance of the individual estimates) is well described by a Gaussian distribution due to the central-limit theorem.

ACKNOWLEDGMENTS

J. S. H., P. M. M., J. D. R., and X. S. acknowledge support from the NSF NANOGrav Physics Frontier Center (NSF Grants No. PHY-1430284 and No. PHY-2020265). J. D. R. acknowledges support from start-up funds from Texas Tech University. J. S. H. acknowledges support from start-up funds from Oregon State University. The authors thank `Math.StackExchange` user River Li, whose answer [26] introduced us to the technique of diagonalization to derive probability distributions for inner products. Lastly, we also thank the authors of [10] for helpful discussions regarding the MatLab code [27] they developed to calculate GX2 distributions.

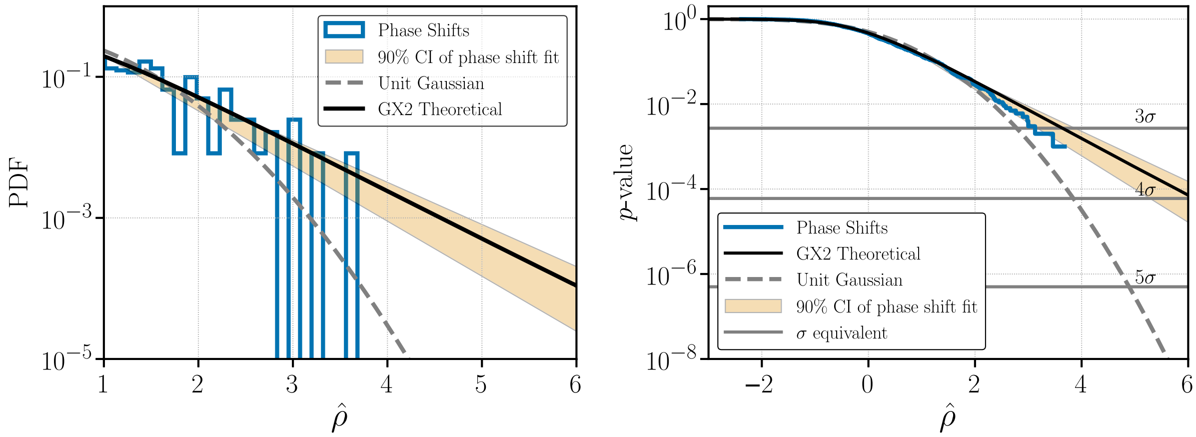

FIG. 3. Comparison of GX2 fit to empirical phase shifts and tail-fitting procedure. Left: 90% credible interval of fit (orange) to the empirical PDF (blue). The theoretical distribution from the GX2 distribution is shown in black. Right: same as the left but for the p-value ($1 - \mathrm{CDF}$). In both cases we can see that fitting the empirical distribution from phase shifts with an exponential gives a reasonable approximation of the GX2 in the region in which we are interested.

APPENDIX: EMPIRICAL TAIL FITTING AND EXTRAPOLATION

The GX2 is the analytic null distribution for the optimal statistic, but it depends upon specific characteristics of the data set, e.g., white-noise and red-noise parameters. In the case of red noise, a Bayesian analysis is used to obtain a posterior distribution on the amplitude and spectral index; using different draws from that posterior to calculate the GX2 changes the resulting distribution. Choosing the best parameters to use requires care and potentially comparison with empirical distributions. In this appendix we introduce a technique that can be used in the interim. In situations where we need to evaluate the null distribution in a place where we have few empirical simulations, but we do not yet have the full GX2 method, we can empirically fit the tail of the empirical simulations with an exponential function. This method has been used to extrapolate the tail of the null distribution in, e.g., searches for GWs from rapidly rotating neutron stars [28]. The uncertainty on the fit parameters can then be used to bound our estimate of the p-value or false alarm probability.

We begin by choosing a point beyond which $p(\hat{\rho})$ looks like an exponential. This choice is arbitrary, but the procedure can be done using several choices, picking the one that looks most reasonable. We call this value $\hat{\rho}_{\mathrm{tail}}$. Therefore, for $\hat{\rho} > \hat{\rho}_{\mathrm{tail}}$, we have

$$p(\hat{\rho}|\hat{\rho} > \hat{\rho}_{\mathrm{tail}}) = \lambda \exp\left[-\lambda(\hat{\rho} - \hat{\rho}_{\mathrm{tail}})\right]. \tag{A1}$$

The shape parameter λ is a free parameter that we fit using the values in the tail of the empirical distribution. Since there is only a single free parameter, we can fit this using a brute-force Bayesian approach, with a uniform prior on λ

$$p\left[\lambda|\{\hat{\rho}_i\}_{i=1}^{N_{\mathrm{tail}}}\right] = \frac{1}{\lambda_{\max}}\lambda^{N_{\mathrm{tail}}} \exp\left\{\left[-\lambda\sum_{i=1}^{N_{\mathrm{tail}}}\left(\hat{\rho}_i - \hat{\rho}_{\mathrm{tail}}\right)\right]\right\}, \tag{A2}$$

where N_{tail} is the number of empirical distribution values satisfying $\hat{\rho}_i > \hat{\rho}_{\mathrm{tail}}$, and $\lambda_{\max}$ is the upper bound on the uniform prior on λ. In this notation, $\{\hat{\rho}_i\}_{i=1}^{N_{\mathrm{tail}}}$ denotes the set of empirical distribution results that populate the tail.

The posterior on λ can then be used to quantify the uncertainty in the tail of the distribution on $\hat{\rho}$. We show an example of fits to the tail of the empirical distribution in Fig. 3 using 300 different draws from the posterior on λ (the 90% credible interval is shown in orange), along with the GX2 estimate (black line) and the empirical distribution (blue histogram). The blue histogram uses the phase shifts from the NANOGrav 12.5-year analysis that were used to generate Fig. 2. For the sake of comparison, we also include the best-fit Gaussian PDF (gray dashed), which clearly underestimates $p(\hat{\rho})$ in the tail. The exponential fit agrees with the GX2 estimate out to $\hat{\rho} = 6$, indicating it can be effectively used to extrapolate the tail and still be consistent with the analytic distribution. We can also quantify the uncertainty in the tails of the empirical distribution due to having a small number of points in the tail of the empirical distribution.

Once we have a fit for $p(\hat{\rho})$, we can analytically calculate the cumulative distribution function (CDF), $p(\hat{\rho} < \rho)$, which can be used to estimate the false alarm probability. The CDF is given by

$$p(\hat{\rho} < \rho|\lambda) = \begin{cases} \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\hat{\rho}_i \le \rho} & \rho \le \hat{\rho}_{\text{tail}}, \\ \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\hat{\rho}_i \le \hat{\rho}_{\text{tail}}} + \mathcal{N}\left(1 - e^{-\lambda(\rho - \hat{\rho}_{\text{tail}})}\right) & \rho > \hat{\rho}_{\text{tail}}, \end{cases} \tag{A3}$$

where $\mathbb{1}$ is the indicator function that equals 1 if the subscript condition is met and equals zero otherwise. In this expression, N is the total number of empirical trials, and $\hat{\rho}_i$ is now any one of those N trials (as opposed to being taken only from the tail). The normalization factor $\mathcal{N}$ is chosen such that the CDF evaluates to 1 at infinity.

[1] Gabriella Agazie *et al.* (NANOGrav Collaboration), The NANOGrav 15 yr Data Set: Evidence for a gravitational-wave background, Astrophys. J. Lett. **951**, L8 (2023).

[2] J. Antoniadis *et al.*, The second data release from the European Pulsar Timing Array III. Search for gravitational wave signals, Astron. Astrophys. **678**, A50 (2023).

[3] Daniel J. Reardon *et al.*, Search for an isotropic gravitational-wave background with the Parkes Pulsar Timing Array, Astrophys. J. Lett. **951**, L6 (2023).

[4] G. Agazie *et al.* (The International Pulsar Timing Array Collaboration), Comparing recent PTA results on the nanohertz stochastic gravitational wave background, arXiv:2309.00693.

[5] R. W. Hellings and G. S. Downs, Upper limits on the isotropic gravitational radiation background from pulsar timing analysis, Astrophys. J. Lett. **265**, L39 (1983).

[6] M. Anholm, S. Ballmer, J. D. E. Creighton, L. R. Price, and X. Siemens, Optimal strategies for gravitational wave stochastic background searches in pulsar timing data, Phys. Rev. D **79**, 084030 (2009).

[7] S. J. Chamberlin, J. D. E. Creighton, X. Siemens, P. Demorest, J. Ellis, L. R. Price, and J. D. Romano, Time-domain implementation of the optimal cross-correlation statistic for stochastic gravitational-wave background searches in pulsar timing data, Phys. Rev. D **91**, 044048 (2015).

[8] Sarah J. Vigeland, Kristina Islo, Stephen R. Taylor, and Justin A. Ellis, Noise-marginalized optimal statistic: A robust hybrid frequentist-Bayesian statistic for the stochastic gravitational-wave background in pulsar timing arrays, Phys. Rev. D **98**, 044003 (2018).

[9] Joseph D. Romano and Neil. J. Cornish, Detection methods for stochastic gravitational-wave backgrounds: A unified treatment, Living Rev. Relativity **20**, 2 (2017).

[10] Abhranil Das and Wilson S. Geisler, A method to integrate and classify normal distributions, J. Vis. **21**, 1 (2021).

[11] J. M. Cordes and R. M. Shannon, Minimum requirements for detecting a stochastic gravitational wave background using pulsars, Astrophys. J. **750**, 89 (2012).

[12] S. R. Taylor, L. Lentati, S. Babak, P. Brem, J. R. Gair, A. Sesana, and A. Vecchio, All correlations must die: Assessing the significance of a stochastic gravitational-wave background in pulsar-timing arrays, Phys. Rev. D **95**, 042002 (2017).

[13] N. J. Cornish and L. Sampson, Towards robust gravitational wave detection with pulsar timing arrays, Phys. Rev. D **93**, 104047 (2016).

[14] Zaven Arzoumanian *et al.*, The NANOGrav 12.5-year data set: Search for an isotropic stochastic gravitational-wave background, arXiv:2009.04496.

[15] Rutger van Haasteren and Yuri Levin, Understanding and analysing time-correlated stochastic signals in pulsar timing, Mon. Not. R. Astron. Soc. **428**, 1147 (2013).

[16] M. T. Lam, M. A. McLaughlin, J. M. Cordes, S. Chatterjee, and T. J. W. Lazio, Optimal frequency ranges for submicrosecond precision pulsar timing, Astrophys. J. **861**, 12 (2018).

[17] M. T. Lam, Optimizing pulsar timing array observational cadences for sensitivity to low-frequency gravitational-wave sources, Astrophys. J. **868**, 33 (2018).

[18] Jeffrey S. Hazboun, Joseph D. Romano, and Tristan L. Smith, Realistic sensitivity curves for pulsar timing arrays, Phys. Rev. D **100**, 104028 (2019).

[19] Jeffrey Hazboun, Joseph Romano, and Tristan Smith, HASASIA: A PYTHON package for pulsar timing array sensitivity curves, J. Open Source Software **4**, 1775 (2019).

[20] Md F. Alam *et al.*, The NANOGrav 12.5-year data set: Observations and narrowband timing of 47 millisecond pulsars, Astrophys. J. Suppl. Ser. **252**, 1 (2020).

[21] Zaven Arzoumanian *et al.* (NANOGrav Collaboration), The NANOGrav 12.5 yr data set: Search for an isotropic stochastic gravitational-wave background, Astrophys. J. Lett. **905**, L34 (2020).

[22] Valentina Di Marco, Andrew Zic, Matthew T. Miles, Daniel J. Reardon, Eric Thrane, and Ryan M. Shannon, Toward robust detections of nanohertz gravitational waves, Astrophys. J. **956**, 14 (2023).

[23] Z. Arzoumanian *et al.* (The NANOGrav Collaboration), The NANOGrav nine-year data set: Observations, arrival time measurements, and analysis of 37 millisecond pulsars, Astrophys. J. **813**, 65 (2015).

[24] Paul B. Demorest, Measuring the gravitational wave background using precision pulsar timing, Ph.D. thesis, University of California, Berkeley, 2007.

[25] R. van Haasteren and M. Vallisneri, Low-rank approximations for large stationary covariance matrices, as used in the Bayesian and generalized-least-squares analysis of

pulsar-timing data, Mon. Not. R. Astron. Soc. **446**, 1170 (2015).
[26] https://math.stackexchange.com/a/3716276.
[27] https://www.mathworks.com/matlabcentral/fileexchange/85028-generalized-chi-square-distribution.
[28] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal *et al.*, Search for continuous gravitational waves from 20 accreting millisecond x-ray pulsars in O3 LIGO data, Phys. Rev. D **105**, 022002 (2022).