

Generalized optimal statistic for characterizing multiple correlated signals in pulsar timing arrays

Shashwat C. Sardesai and Sarah J. Vigeland 

*Center for Gravitation, Cosmology and Astrophysics, University of Wisconsin—Milwaukee,
P.O. Box 413, Milwaukee, Wisconsin 53201, USA*

Kyle A. Gersbach and Stephen R. Taylor

*Department of Physics and Astronomy, Vanderbilt University,
2301 Vanderbilt Place, Nashville, Tennessee 37235, USA*



(Received 21 March 2023; accepted 4 December 2023; published 29 December 2023)

Pulsar timing arrays are sensitive to low-frequency gravitational waves (GWs), including the low-frequency stochastic gravitational wave background (GWB), which induces correlated changes in millisecond pulsars' timing residuals described by the Hellings-Downs curve. Some sources of noise can also induce correlated changes in pulsar timing residuals, albeit with different correlation signatures. A spatial correlation that differs from Hellings-Downs could also be indicative of non-Einsteinian GW polarizations. It is therefore crucial that we be able to characterize the spatial correlation in order to distinguish between the GWB and sources of noise. The optimal statistic (OS) is a frequentist estimator for the amplitude and significance of a spatially correlated signal in PTA data, and it is widely used to search for the GWB. However, the OS cannot perfectly distinguish between different spatial correlations. In this paper, we introduce the multiple component optimal statistic (MCOS): a generalization of the OS that allows for multiple correlations to be simultaneously fit to the data. We use simulated data to show that this method more accurately recovers injected spatially correlated signals, and in particular eliminates the problem of overestimating the amplitude of correlations that are not present in the data. We also demonstrate that this method can be used to recover multiple correlated signals.

DOI: [10.1103/PhysRevD.108.124081](https://doi.org/10.1103/PhysRevD.108.124081)

I. INTRODUCTION

When massive bodies such as black holes orbit each other, they radiate away energy in the form of gravitational waves (GWs), causing their orbit to shrink until they coalesce. Supermassive black hole binaries (SMBHBs), which form in galaxy mergers, emit low-frequency GWs over many orders of magnitude in frequency as they go from subparsec separations to coalescence. The incoherent superposition of GWs from a cosmological population of SMBHBs produces a stochastic GW background (GWB) in the nanohertz frequency range ($\sim 10^{-9}$ – 10^{-7} Hz) [1–3].

Pulsars are rapidly spinning neutron stars that emit radio waves that can be observed by radio telescopes. Pulsars have very stable spin rates, which makes them ideal for detecting long-period gravitational waves [4–6]. In pulsar timing arrays, pulsars essentially act as astronomical clocks, allowing us to measure small fluctuations in spacetime over years to decades. There are currently four regional PTA experiments: the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) [7], the European Pulsar Timing Array (EPTA) [8], the Parkes Pulsar Timing Array (PPTA) [9], and the Indian Pulsar Timing Array (InPTA) [10]. All of these groups collaborate

and share data under the umbrella of the International Pulsar Timing Array (IPTA) [11].

The presence of GWs induces correlated changes in the pulse times of arrival, with the correlation between two different pulsars depending on their angular separation on the sky according to the Hellings-Downs (HD) curve [12]. This characteristic correlation allows us to distinguish between the GWB and other astrophysical and terrestrial effects that could affect the pulse times of arrival of many pulsars, such as the solar wind, instrumentation errors, a clock error, ephemeris errors, etc. [13–15]. A correlated signal with a correlation pattern that differs from the HD curve could also indicate the presence of non-Einsteinian GW polarizations [16–19].

It is therefore crucial to be able to distinguish between different cross-correlations. One method for doing this is using the OS, a frequentist estimator of the amplitude and significance of a correlated stochastic process, where the correlation function of interest is determined by specifying the overlap reduction function (ORF) [20–22]. One limitation of the OS is that for any real PTA, it is not possible to perfectly distinguish between different ORFs. In general, the ORFs for different cross-correlations are not orthogonal,

with the overlap between them depending on the number of pulsars and their sky locations [23,24]. In this paper, we address this issue by introducing a modified version of the OS that allows us to simultaneously search for multiple correlated OS signals (MCOS). As we show in this paper, this allows us to more accurately characterize the presence of correlated signals in PTA data.

This paper is organized as follows. In Sec. II, we discuss the OS and derive a modified version that can simultaneously fit multiple correlation functions. In Sec. III, we describe the methods, signal models, and software used to generate and analyze simulated PTA data. We generated three types of simulated datasets: one containing a GWB (i.e., a common-spectrum process with HD correlations), another containing a common-spectrum process with GW-like monopole correlations, and finally a model with both a GWB, and GW-like monopole signal injected. In Sec. IV we use the standard OS and our MCOS to analyze the simulated datasets. We find that simultaneously fitting multiple correlations using the MCOS prevents overestimating the presence of correlations that are not present, while not significantly affecting parameter estimation of correlated signals. We also demonstrate that the MCOS allows us to recover multiple correlated signals from the data. We summarize the paper and discuss future work in Sec. V.

II. OPTIMAL STATISTIC

The OS is a frequentist estimator of the amplitude and significance of the stochastic background. It only considers the cross-correlations between two different pulsars and not auto-correlation terms of individual pulsars. The OS is very fast to compute compared to performing an equivalent Bayesian analysis; however, the OS can be biased due to covariance between individual pulsar red noise and a common stochastic process. The noise-marginalized OS is a hybrid Bayesian-frequentist method that marginalizes over the pulsars' intrinsic red noise [23]. It has been shown that this method more accurately estimates the amplitude of the GWB in the presence of intrinsic pulsar red noise compared to computing the OS with fixed pulsar red noise parameters.

One method to derive the OS is to consider fitting the cross-correlations between pulsars to an ORF Γ_{ab} with amplitude A^2 [21,22]. Let \mathbf{r}_a be the residuals for pulsar a . The cross-correlations between pulsars a and b , ρ_{ab} , and their associated uncertainties, σ_{ab} , can be written as

$$\rho_{ab} = \mathcal{N}_{ab} \mathbf{r}_a^T \mathbf{P}_a^{-1} \bar{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \mathbf{r}_b, \quad (1)$$

$$\sigma_{ab}^2 = \mathcal{N}_{ab} = (\text{tr}[\mathbf{P}_a^{-1} \bar{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \bar{\mathbf{S}}_{ba}])^{-1}, \quad (2)$$

where \mathbf{P}_a is the autocovariance matrix and contain contributions from white noise, intrinsic red noise, and the common stochastic process red noise from pulsar a ; and

$\mathbf{S}_{ab} = A^2 \Gamma_{ab} \bar{\mathbf{S}}_{ab}$ is the cross-covariance matrix and only considers contributions from the common stochastic process between pulsars a and b .

We can fit the measurement cross-correlations to an ORF. The chi-squared is given by

$$\chi^2 = \sum_{ab, a < b} \left(\frac{\rho_{ab} - A^2 \Gamma_{ab}}{\sigma_{ab}} \right)^2, \quad (3)$$

where σ_{ab} are the uncertainties in the cross-correlations and contain contributions from both the common process and intrinsic noise, as shown in [22]. The OS is the value of A^2 that minimizes the chi-squared:

$$\hat{A}^2 = \frac{\sum_{ab} (\rho_{ab} \Gamma_{ab} / \sigma_{ab}^2)}{\sum_{ab} (\Gamma_{ab}^2 / \sigma_{ab}^2)}. \quad (4)$$

The uncertainty in \hat{A}^2 can be found through simple propagation of errors:

$$\sigma_{\hat{A}^2}^2 = \left[\sum_{ab} \left(\frac{\Gamma_{ab}^2}{\sigma_{ab}^2} \right) \right]^{-1}. \quad (5)$$

The corresponding signal-to-noise ratio (S/N) is

$$\text{S/N} = \frac{\hat{A}^2}{\sigma_{\hat{A}^2}}. \quad (6)$$

The OS as derived above assumes the presence of only one spatially correlated process in our data. It is possible that PTA data could contain multiple spatially correlated processes, e.g., a GWB and a common correlated source of noise. As discussed in Tiburzi *et al.* [14], some noise sources can produce spatially correlated signals—a clock error appears in PTA data as a common process with monopolar correlations between different pulsars, while an ephemeris error produces dipolar correlations.

If more than one spatially correlated process is present, the chi-squared becomes

$$\chi^2 = \sum_{ab} \left(\frac{\rho_{ab} - \sum_{\alpha} A_{\alpha}^2 \Gamma_{ab}^{\alpha}}{\sigma_{ab}} \right)^2, \quad (7)$$

where α indexes the different spatially correlated signals. Minimizing the chi-squared with respect to the amplitude squared value A_{α}^2 gives

$$C^{\beta} - \sum_{\alpha} \hat{A}_{\alpha}^2 B^{\alpha\beta} = 0, \quad (8)$$

where

$$C^{\beta} \equiv \sum_{ab} \frac{\rho_{ab} \Gamma_{ab}^{\beta}}{\sigma_{ab}^2}, \quad (9)$$

$$B^{ab} \equiv \sum_{ab} \frac{\Gamma_{ab}^\alpha \Gamma_{ab}^\beta}{\sigma_{ab}^2}. \quad (10)$$

Therefore, \hat{A}_α^2 is given by

$$\hat{A}_\alpha^2 = \sum_{\beta} B_{\alpha\beta} C^\beta, \quad (11)$$

where $B_{\alpha\beta}$ is the inverse of $B^{\alpha\beta}$. This is a more general form of Eq. (4): when we generalize to multiple ORFs, the numerator of Eq. (4) becomes C^β , while the denominator becomes $B_{\alpha\beta}$. The uncertainty in \hat{A}_α^2 is

$$\sigma_{\hat{A}_\alpha^2}^2 = B_{\alpha\alpha}. \quad (12)$$

III. METHODS AND SOFTWARE

In this section we describe how to construct the pulsar timing model, including the intrinsic and common red noise, and discuss how to retrieve the injected signals. We based our simulated datasets on the NANOGrav 12.5-year dataset [25], with the pulsar positions, observing spans, cadences, and white and red noise all taken from that dataset. In addition, we inject two different kinds of stochastic backgrounds: one with HD correlations (i.e., a GWB), and one with GW-like monopolar correlations. The methods used are similar to those used in the NANOGrav 12.5-year GWB paper [26]: below we briefly describe them, highlighting how our methods differ from those in [26].

A. PTA model

We model the timing residuals for a pulsar δt as

$$\delta t = M\epsilon + F\mathbf{a} + \mathbf{n}, \quad (13)$$

where $M\epsilon$ describes linear perturbations to the timing model, $F\mathbf{a}$ describes red noise, including both individual pulsar red noise and a common stochastic process affecting all of the pulsars, and \mathbf{n} is the uncorrelated white noise. We model the red noise using a Fourier series with frequencies $f = 1/T, 2/T, \dots$, where T is the span of the observations, and use a power-law model for the power spectral density

$$P(f) = \frac{A^2}{12\pi^2} \left(\frac{f}{f_{\text{yr}}} \right)^{-\gamma} \text{yr}^3, \quad (14)$$

where A is the amplitude, defined at a reference frequency of $f_{\text{yr}} = 1/(1 \text{ year})$, and γ is the spectral index.

B. Common stochastic process

In addition to pulsar red noise and white noise, we also include the presence of a common stochastic process. For all types of stochastic backgrounds, we model the power

spectrum as a power law as described in Eq. (14) with $A_{\text{gw}} = 2 \times 10^{-15}$ and $\gamma = 13/3$. The value of γ corresponds to that for a GWB made from GWs from circular SMBBHs evolving only due to GW emission [27].

We model two different types of stochastic process: one with HD correlations, as is expected for a GWB; one with spatial correlations described by a GW-like monopole (GWMO). The ORF for HD correlations is [12,28]

$$\Gamma_{ab}^{\text{HD}}(\xi) = \frac{\delta_{ab}}{2} + \frac{1}{2} - \frac{(1 - \cos \xi)}{4} \left[\frac{1}{2} - 3 \ln \left(\frac{1 - \cos \xi}{2} \right) \right], \quad (15)$$

where Γ_{ab}^{HD} is the HD ORF for pulsars with indices a and b , and ξ is the angular separation between them. Note that the maximum correlation between two different pulsars (i.e., $a \neq b$) is $1/2$: this is because the GWB induces an Earth term and a pulsar term in each pulsar's residuals, and only the Earth terms are correlated. The GWMO was introduced in [24] and the ORF takes the form

$$\Gamma_{ab}^{\text{GWMO}}(\xi) = \frac{1}{2} + \frac{\delta_{ab}}{2}. \quad (16)$$

The cross-correlation for two different pulsars is $1/2$ rather than 1 because we are assuming the signal has two components: an ‘‘Earth term’’ that is correlated and a ‘‘pulsar term’’ that is uncorrelated (hence why it is described as ‘‘GW-like’’). The GWMO does not correspond to any physical process, but it is similar to the ORF produced by a scalar-tensor mode.

C. Simulated datasets

We use LIBSTEMPO [29] to generate our simulated datasets. The simulated datasets are based on the NANOGrav 12.5-year dataset [25]. We create two types of simulated datasets. One contains a common stochastic process with HD correlations. The second contains two common stochastic processes: one with HD correlations and another with GWMO correlations. For both, we use the linearized timing models for the pulsars in the 12.5-year dataset, as well as the dates and times of each observation. We first idealize our pulsar timing residuals and then add uncorrelated white noise equal to the measured TOA uncertainties. We add pulsar intrinsic red noise, with the maximum likelihood values taken from a Bayesian run of the NANOGrav 12.5-year dataset with the values given (see Appendix B).

D. Optimal statistic calculation

In order to compute the OS, we must specify values for the pulsars' intrinsic white noise and red noise. The white noise parameters are the maximum-likelihood values from Bayesian noise analyses where each pulsar is run individually. The red noise parameters come from a Bayesian analysis of all the pulsars that simultaneously fits for pulsar

intrinsic red noise and a common uncorrelated stochastic process. The white noise parameters are not searched over in the Bayesian analysis and are instead fixed to the maximum-likelihood values. We use uniform priors for the red noise parameters, $\log_{10}A \in [-20, -11]$ and $\gamma \in [0, 7]$, and the common process amplitude, $\log_{10}A_{\text{gw}} \in [-18, -11]$, while the common process γ_{gw} is fixed to 13/3. We use ENTERPRISE [30] and ENTERPRISE_EXTENSIONS [31] to implement the models and compute the likelihood, and we use Markov Chain Monte Carlo (MCMC) methods to obtain samples from the posterior as implemented in PTMCMCSampler [32].

We then compute the OS, using the results of the Bayesian analysis to determine the pulsar intrinsic red noise and common uncorrelated stochastic process. There are two ways of doing this. The *fixed-noise* version computes the OS at a single noise realization using the red noise values that maximize the likelihood of the Bayesian analysis. The *noise-marginalized* version pulls red noise values from the posteriors obtained by the Bayesian analysis in order to compute the OS at many different noise realizations. This results in distributions for \hat{A}^2 and σ .

We can compute the OS for different correlations by using different ORFs. In addition to the HD and GWMO ORFs given in Eqs. (15) and (16), we also consider a standard dipole ORF and monopole ORF:

$$\Gamma_{ab}^{\text{dipole}}(\xi) = \cos \xi, \quad (17)$$

$$\Gamma_{ab}^{\text{monopole}}(\xi) = 1. \quad (18)$$

In Fig. 1, we plot all four ORFs used in this paper. Note that the monopole and GWMO are both constant as a function of angular separation, but the monopole as a value of 1 while the GWMO as a value of 1/2. This means that any signal that fits one will also fit the other, but the inferred




FIG. 1. Comparison between HD (blue), monopole (red), dipole (green), and GWMO (purple) ORFs as a function of angular separation between pulsars. The solid black line indicates a zero amplitude correlation. Note that a GWMO has the same functional form as the standard monopole, but a different normalization.

TABLE I. Match statistic \bar{M} for HD, monopole, dipole, and GW-like monopole correlations, as defined by Eq. (19), for our simulated PTA. The monopole and GW-like monopole have $\bar{M} = 1$ since they only differ by a normalization factor.

| Correlation | HD | Monopole | Dipole | GW-like monopole |
|------------------|-----|----------|--------|------------------|
| HD | 1 | 0.255 | 0.435 | 0.255 |
| Monopole | ... | 1 | 0.395 | 1 |
| Dipole | ... | ... | 1 | 0.395 |
| GW-like monopole | ... | ... | ... | 1 |

amplitudes will differ because the two ORFs have different normalizations.

We can measure the overlap between different ORFs by computing the “unweighted match statistic” [33],

$$\bar{M} = \frac{\sum_{a,b \neq a} \Gamma_{ab} \Gamma'_{ab}}{\sqrt{(\sum_{a,b \neq a} \Gamma_{ab} \Gamma_{ab})(\sum_{a,b \neq a} \Gamma'_{ab} \Gamma'_{ab})}}, \quad (19)$$

where Γ and Γ' are two different ORFs. This overlap depends on the number of pulsars in the PTA and their sky locations. Table I lists the match statistics for the four ORFs considered in this paper for our simulated PTA, which is based on the NANOGrav 12.5-year dataset.

IV. RESULTS

Here we present the results of analyzing our simulated datasets using the standard OS calculation and the modified MCOS. We show two types of simulations: one with an injected GWB (i.e., a common stochastic process with HD correlations), and one with both HD and GWMO correlations. We present the recovered OS values and corresponding signal-to-noise ratios (S/N) for different spatial correlations, and compare them to the injected values. We note the fraction of realizations with $S/N \geq 3$ in Appendix A.

We also look at how accurately the OS and MCOS recover the amplitude of the correlated process. If a non-negligible correlated signal is present in the data, then the expressions for \hat{A}^2 and $\sigma_{\hat{A}^2}$ given in Eqs. (4), (5), (11), and (12) must be revised to include interpulsar pair correlation [34,35]. For more details, see Appendix C. We use the recovered \hat{A}^2 and $\sigma_{\hat{A}^2}$ to compute the percentile of the injected value, assuming a Gaussian distribution,¹ and compare the fraction of simulations for.

¹As shown in [36], the optimal statistic actually follows a generalized chi-squared distribution and not a Gaussian distribution. In this work, we approximate the distribution as a Gaussian, owing to the computational expense of constructing the generalized chi-squared distribution for each simulation. This is a reasonably good approximation, and has been used in previous work [23,35].

Finally, we look at how much the data preferred models with multiple correlated signals using the Akaike Information Criterion (AIC) [37], which is a frequentist analog for the Bayes factor. It is given by

$$\text{AIC} = 2k + \chi^2, \quad (20)$$

where k is the number of model parameters (i.e., the number of correlated signal amplitudes, or the number of ORFs), and the chi-squared given by Eq. (3) for a single correlated signal and Eq. (7) for multiple correlated signals. The factor of $2k$ acts as an Occam's penalty for adding more parameters to a model. We compute χ^2 using the maximum-likelihood red noise parameters. The relative probability of two models is then given by

$$p(\text{AIC}) = e^{\frac{(\text{AIC} - \text{AIC}_{\min})}{2}}, \quad (21)$$

where AIC_{\min} represents the minimum AIC value of all the different models. The model with the minimum AIC will thus be the most preferred model, with a $p(\text{AIC}) = 1$. Some of the models are equally preferred, with the difference in their $\text{AIC} \sim 0.01$, so we set a threshold value of $p(\text{AIC}) \geq 0.99$ to account for these cases. These results are also presented in Appendix A.

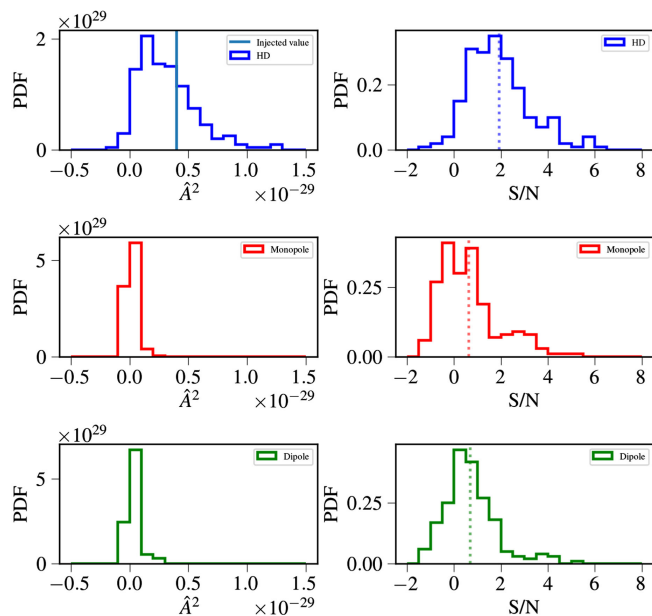


FIG. 2. Noise-marginalized OS analysis of 200 simulations with an injected HD. We show results using HD (blue), monopole (red), and dipole (green) ORFs. The left hand columns show the distributions of \hat{A}^2 , whereas the right hand columns show the S/N histograms. The vertical blue line in the top left plot indicates the injected value of $A_{\text{HD}}^2 = 4 \times 10^{-30}$. The colored dashed lines in the S/N plots are the means of the respective ORFs. We find a mean $\hat{A}^2 > 0$ and mean $\text{S/N} > 0$ for monopole and dipole correlations even though neither were injected in the data.

A. HD simulation

We created 200 simulations with an injected HD which we analyzed using the traditional OS and the MCOS. A complete summary of the results can be found in Table II in Appendix A. When we search for one spatial correlation at a time, we find a $\text{S/N} > 3$ for HD correlations in 17% of the simulations, but we also find $\text{S/N} > 3$ for monopole and dipole correlations in 6% and 5% of the simulations, respectively, even though there are no monopole-correlated or dipole-correlated signals present in the datasets. In Fig. 2, we show histograms for the recovered distributions of \hat{A}^2 and the S/N for HD, monopole, and dipole correlations. On average, we recover higher S/N of HD correlations than for monopole or dipole correlations; however, we recover $\text{S/N} > 0$ for monopole and dipole correlations in more than half of the simulations, as shown by the dotted lines in Fig. 2.

Figure 3 shows $p-p$ plots for the recovered single component OS \hat{A}^2 using HD, monopole, and dipole correlations. For each simulation, we use the marginalized chain values to calculate \hat{A}^2 and $\sigma_{\hat{A}^2}$, taking pulsar pair

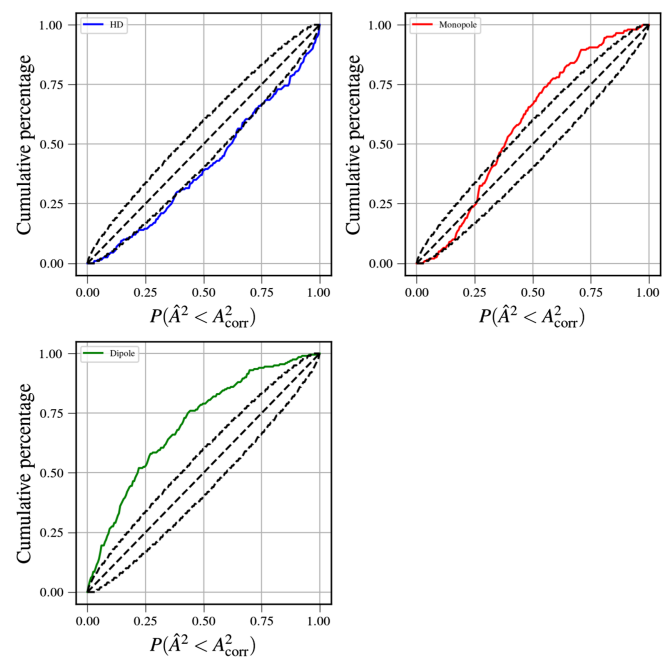


FIG. 3. $p-p$ plot for the HD injected simulations with single component OS. The horizontal axis is the percentile of each simulation to the injected value. The vertical axis represents the cumulative fraction of simulations that are recovered at that particular percentile. The models shown here are HD (top left), monopole (top right), and dipole (bottom left). If the parameters are recovered accurately, the two quantities would be equal (black dashed line). The curved black lines indicate the 95% confidence intervals of the $p-p$ plots. We find that the amplitude of HD correlations is underestimated, while the amplitude of monopole and dipole correlations (which are not present) are significantly overestimated.

cross covariance into consideration, and then compute the percentile of the injected value for that type of correlation. We then plot the cumulative fraction of realizations for which the injected value is at that percentile. If the parameters are accurately recovered, we would expect the cumulative fraction of realizations to be equal to the cumulative percentile. If it is not, that indicates the parameter is being overestimated (over the central dashed black line) or underestimated (under the central dashed black line). We find that the amplitude of HD correlations is underestimated, while the amplitudes of monopole and dipole correlations (which were not injected) are overestimated. The underestimation of the amplitude of HD correlations is likely due to two factors: first, there is covariance between the pulsars' intrinsic red noise and common red noise, and second, the simulations inject the intrinsic red noise and common red noise over more frequency components than we use to recover them.

In contrast, the MCOS more accurately recovers the amplitude of monopole and dipole correlations. We show the recovered distributions of \hat{A}^2 and the S/N for models consisting of multiple correlations in Fig. 4. The distributions for the amplitude and S/N of HD correlations are similar to those obtained when fitting for only HD correlations (Fig. 2, top row), but now the distributions for the amplitude and S/N of monopole and dipole




FIG. 4. MCOS analysis of 200 simulations for HD simulations for the HD + monopole + dipole ORFs. The top figure contains histograms of \hat{A}^2 and the bottom figure contains histograms for the S/N. The solid light blue line is the injected value of $A_{\text{HD}}^2 = 4 \times 10^{-30}$. The colored dashed lines in the S/N plots are the means of the respective ORFs. The recovered amplitudes and S/N for monopole and dipole correlations are centered around zero, unlike in Fig. 2.

correlations are centered around zero. The $p - p$ plots, shown in Fig. 5, show the recovery of the MCOS \hat{A}^2 . We find that MCOS analysis recovers the injected parameters for the monopole and dipole ORFs more accurately than the OS. The amplitude of HD correlations is slightly underestimated, as it was when using the OS.

Table II lists the percentage of ORFs where an S/N greater than 3 was recovered, as well as the percentage of cases where the particular model had the lowest AIC value (most preferential model). In 51.5% of realizations, the standard OS, HD ORF model had the lowest AIC value, as expected, considering our injection has only one signal present. The standard monopole and monopole + HD model have nearly the same preference. We note that these simulations include relatively low-significance HD correlations, as can be seen by the fact that the HD S/N > 3 in only 17% of simulations. But the percentage of models with an HD S/N > 3 does not change drastically between the individual and MCOS models. This means that it is difficult to distinguish between different correlations. Nevertheless, the MCOS analysis allows us to search for the presence of spatial correlations and compute the relative probability of different models.

B. HD + GWMO simulations

Finally, we looked at how well the MCOS could recover multiple correlated signals. We generated 200 simulations that contained two stochastic processes: one with HD correlations and one with GWMO correlations. We used




FIG. 5. $p - p$ plot for the MCOS HD injected simulations. The amplitudes shown here are recovered using a model that simultaneously fits for HD, monopole, and dipole correlations. The horizontal axis is the recovered percentile of each simulation to the injected value. The vertical axis represents the cumulative fraction of realizations that recover the particular percentile. The central black dashed line indicates accurate recovery, and the upper and lower dashed black lines indicate the 95% confidence interval. We find that the amplitudes of monopole and dipole correlations are accurately recovered, while the amplitude of HD correlations is underestimated.

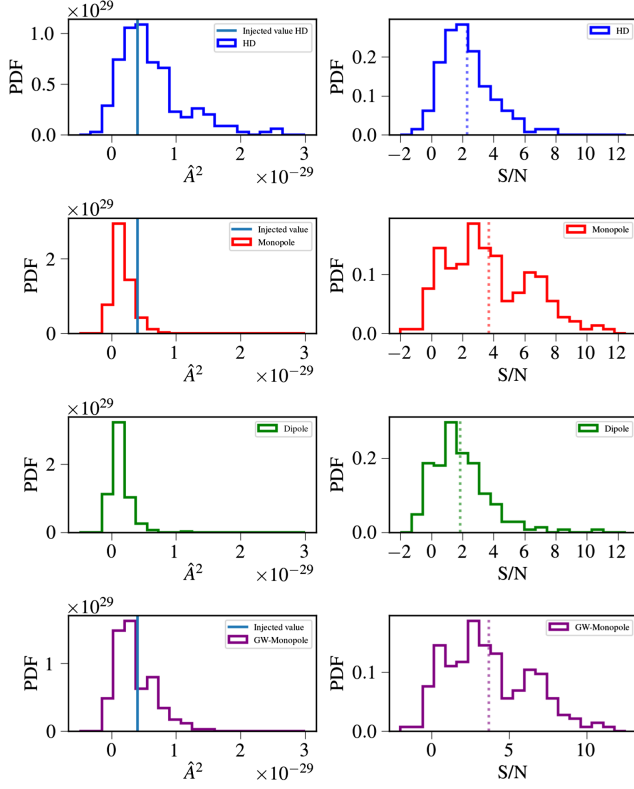


FIG. 6. Individual correlated signal analysis for 200 HD + GWMO simulations. Similar to Fig. 2, histograms of \hat{A}^2 and S/N on the left and right hand columns respectively. From top to bottom, the ORFs are HD (blue), monopole (red), dipole (green), and GWMO (purple). The solid light blue line indicates the injected OS of $A_{\text{HD}}^2 = 4 \times 10^{-30}$ and $A_{\text{GWMO}}^2 = 4 \times 10^{-30}$. The colored dashed lines in the S/N plots are the means of the respective ORFs.

the same amplitude and spectral index $\gamma = 13/3$ for both processes. We summarize the results of analyzing these simulations in Table III in Appendix A.

When we analyze the simulations for a single correlation, we find $S/N > 3$ for HD correlations in 27% of simulations and for GWMO in 56% of simulations. As shown in Fig. 6, we recover the GWMO signal with higher mean S/N than the HD signal. We note that the correlation coefficient for HD correlations is less than the correlation coefficient for GWMO correlations for all angular separations (except for zero), which is why the OS recovers the GWMO signal with a higher S/N than the HD signal—even though the HD and GWMO signals have the same amplitude, the GWMO signal has significantly more power in the cross-correlations than the HD signal. We also recover a dipole signal with $S/N > 3$ in 21% of simulations even though no such signal has been injected.

Figure 7 shows p-p plots for HD, monopole, dipole, and GWMO correlations. Since the data contain both HD and GWMO-correlated processes, but we are only fitting one process at a time, both signals are being incorrectly associated with a single correlation, resulting in inaccurate




FIG. 7. The $p - p$ plots of the HD + GWMO simulations for the individual analysis (clockwise from top-left) HD, standard monopole, GWMO, and dipole. The horizontal axis is the recovered percentile of each simulation to the injected value, and the vertical axis is the cumulative fraction of realizations that recover the particular percentile. If the percentiles are recovered accurately, the two quantities would be equal (central dashed line). The upper and lower dashed black lines indicate the 95% confidence interval. Since we are only fitting for one type of correlation, but multiple correlated signals are present in the data, we are not able to accurately recover the amplitudes of the HD or monopole signals. We also overestimate the amplitude of dipole correlations, which are not present.




FIG. 8. HD and GWMO multiple correlated signal analysis for 200 HD + GWMO simulations. Similar to Fig. 2, histograms of the \hat{A}^2 and S/N on the left and right hand columns respectively. The solid light blue line indicates the injected OS of $A_{\text{HD}}^2 = A_{\text{GWMO}}^2 = 4 \times 10^{-30}$. The colored dashed lines in the S/N plots are the means of the respective ORFs.




FIG. 9. The $p-p$ plots for the MCOS analysis of the HD + GWMO injections. The left figure depicts the HD + GWMO + DI model, and the right figure shows the HD + GWMO model. The horizontal axis is the recovered percentile of each simulation to the injected value, and the vertical axis is the cumulative fraction of realizations that recover the particular percentile. The central black dashed line indicates perfectly accurate recovery of percentiles, while the upper and lower black dashed lines indicate the 95% confidence intervals. We are able to more accurately recover the amplitude of HD and GWMO correlations using the MCOS, although the amplitudes of both are underestimated. We also note that including dipole correlations in the model does not bias parameter estimation of the other signals, and we accurately recover that no dipolar signal is present.

parameter estimation. When we fit to HD correlations only, this results in a significant overestimate of the HD-correlated amplitude. When we fit to GWMO correlations only, the amplitude estimate is not as affected, but the uncertainty is underestimated. We also see that the amplitude of a dipolar process is overestimated, even though one is not present.

When we use the MCOS, we find that we are able to more accurately recover the amplitudes of all the correlated processes. Figure 8 shows the recovered amplitudes and S/N, while Fig. 9 shows $p-p$ plots for the recovered amplitudes. When fitting for both HD and GWMO signals, the amplitudes are more accurately recovered than when fitting for only one signal at a time, although the amplitudes are slightly underestimated. This is possibly because of covariance between the pulsar noise and the common signals. We also find that when we include dipole correlations in our model, we accurately recover that no dipole signal is present, and the addition of the dipole signal does not affect the recovery of either the HD and GWMO signals.

V. CONCLUSIONS

In this paper, we present a generalization of the OS that can simultaneously search for multiple correlated signals. As we have shown, this allows us to better distinguish between different spatial correlations. For a real PTA, the ORFs corresponding to different correlations are not orthogonal as shown in Table I. The overlap between them can lead to nonzero S/N recovery of noninjected signals. It is also useful in cases where multiple correlated signals,

e.g., a GWB and correlated noise, are present in the data. We have also performed model selection by computing the AIC, which acts as a pseudo-Bayes factor.

We have tested this method using two simulated datasets. The first contained a stochastic signal with HD correlations. When fitting only a single correlation, we found that we overestimated the presence of correlations that were not present, but using the MCOS eliminated this problem. We also generated simulated datasets that contained both HD and GWMO-correlated signals. We demonstrated that the MCOS allows both signals to be recovered while ruling out any noninjected signals.

In this paper, we have considered four types of correlations: HD, monopole, dipole, and GWMO. This method can be used with any number and type of correlations. In the NANOGrav 15-year gravitational wave background search [38], the MCOS was used to search for evidence of monopolar or dipolar processes, which could be associated with correlated noise sources such as clock errors or ephemeris errors [15,39,40], and to do a model-independent reconstruction of the correlations using Legendre polynomials [41]. It was also used in the search of the NANOGrav 15-year dataset for evidence of non-Einsteinian polarizations [42]. The modular nature of the MCOS makes it a valuable tool for characterizing the nature of correlations in PTA datasets.

ACKNOWLEDGMENTS

We thank Justin Ellis, Jeff Hazboun, David Kaplan, Dusty Madison, Pat Meyers, Gabe Freedman, Joe Romano, Xavier Siemens, Abhimanyu Susobhanan, and Michele Vallisneri for useful discussions. We thank Nima Laal for help creating simulated datasets with monopole injections. This work was supported by National Science Foundation (NSF) grant No. PHY-2011772. S. R. T. acknowledges support from NSF Grants No. AST-2007993 and No. AST-2307719. S. R. T. and K. A. G. also acknowledge support from an NSF CAREER No. 2146016. The authors are members of the NANOGrav collaboration, which receives support from NSF Physics Frontiers Center award No. 1430284 and 2020265.

APPENDIX A: SIMULATION RESULTS

This appendix contains the results of the MCOS analysis on the HD (Table II), and HD + GWMO (Table III) simulations. The middle three columns refer to the percentage of realizations where the ORF listed in the top row had a S/N greater than or equal to three. The last column shows the percentage of realizations where the particular model listed was the most favored, with a $p(AIC) = 1$. Since certain realizations had AIC values very similar to one another, implying that more than one model was equally favored, we have set our $p(AIC)$ threshold to 0.99.

TABLE II. Summary of results of GWB simulations.

| Model | % of Simulations with $S/N > 3$ | | | % of Simulations with $p(\text{AIC}) \geq 0.99$ |
|------------------------|---------------------------------|----------|--------|---|
| | HD | Monopole | Dipole | |
| HD only | 17.5% | ... | ... | 40% |
| Monopole only | ... | 7% | ... | 18% |
| Dipole only | ... | ... | 5% | 14% |
| HD + monopole | 16.5% | 2% | ... | 15.5% |
| HD + dipole | 16% | ... | 1% | 8.5% |
| Monopole + dipole | ... | 3% | 2% | 1% |
| HD + monopole + dipole | 14.5% | 2% | 0.5% | 5% |

TABLE III. Summary of results of simulations with both GWB and GWMO stochastic processes.

| Model | % of Simulations with $S/N > 3$ | | | | % of Simulations with $p(\text{AIC}) \geq 0.99$ |
|------------------------|---------------------------------|-------|-----|-------|---|
| | HD | GWMO | DI | MO | |
| HD | 27.5% | ... | ... | ... | 13.5% |
| GWMO | ... | 53.5% | ... | ... | 39% |
| Dipole | ... | ... | 21% | ... | 7.5% |
| Monopole | ... | ... | ... | 53.5% | 39% |
| HD + GWMO + Dipole | 10.5% | 44% | 2% | ... | 18.5% |
| HD + Monopole + Dipole | 10.5% | ... | 2% | 44% | 18.5% |
| HD + GWMO | 13.5% | 48% | ... | ... | 19.5% |
| HD + Monopole | 13.5% | ... | ... | 48% | 19.5% |
| HD + Dipole | 16.5% | ... | 10% | ... | 2.5% |
| GWMO + Dipole | ... | 46.5% | 5% | ... | 0% |

APPENDIX B: PULSAR RED NOISE PARAMETERS

Our simulated datasets include intrinsic red noise in each pulsar. We model the red noise with a power-law power spectrum $P(f) = \frac{A^2}{12\pi^2} \left(\frac{f}{f_{\text{yr}}}\right)^{-\gamma} \text{yr}^3$ (as shown in Equation (14)). The values of A and γ used for each pulsar are listed in Table IV, and are based on the noise properties of the NANOGrav 12.5-year dataset.

TABLE IV. Pulsar red noise parameters A and γ used in our simulated datasets.

| Pulsar | $\log_{10} A$ injected | γ injected |
|--------------|------------------------|-------------------|
| B1855 + 09 | -14.3053 | 5.5879 |
| B1937 + 21 | -13.4731 | 3.6189 |
| B1953 + 29 | -12.7813 | 1.1485 |
| J0023 + 0923 | -13.2146 | 0.2572 |
| J0030 + 0451 | -14.8837 | 5.1835 |
| J0340 + 4130 | -17.2031 | 0.3121 |
| J0613 - 0200 | -17.8287 | 6.9765 |
| J0636 + 5128 | -16.6117 | 0.4274 |
| J0645 + 5158 | -19.3925 | 6.2542 |
| J0740 + 6620 | -18.3882 | 0.8377 |

(Table continued)

TABLE IV. (Continued)

| Pulsar | $\log_{10} A$ injected | γ injected |
|--------------|------------------------|-------------------|
| J0931 - 1902 | -19.3076 | 3.1884 |
| J1012 + 5307 | -12.7970 | 0.9968 |
| J1024 - 0719 | -17.2779 | 2.5507 |
| J1125 + 7819 | -12.4459 | 1.5990 |
| J1453 + 1902 | -18.9550 | 4.9088 |
| J1455 - 3330 | -19.2175 | 6.2343 |
| J1600 - 3053 | -13.3480 | 0.0627 |
| J1614 - 2230 | -19.8669 | 2.5180 |
| J1640 + 2224 | -16.3220 | 5.6796 |
| J1643 - 1224 | -12.2480 | 1.1554 |
| J1713 + 0747 | -14.0354 | 0.3368 |
| J1738 + 0333 | -19.5594 | 5.2048 |
| J1741 + 1351 | -14.3474 | 1.9857 |
| J1744 - 1134 | -13.3903 | 1.5453 |
| J1747 - 4036 | -12.7570 | 3.5500 |
| J1832 - 0836 | -13.1021 | 1.6074 |
| J1853 + 1303 | -13.3632 | 1.3637 |
| J1903 + 0327 | -12.2570 | 1.8621 |
| J1909 - 3744 | -13.9843 | 1.7386 |
| J1910 + 1256 | -16.0932 | 2.6105 |
| J1911 + 1347 | -13.4836 | 1.7034 |
| J1918 - 0642 | -19.7173 | 4.0390 |
| J1923 + 2515 | -17.7267 | 1.7844 |

(Table continued)

TABLE IV. (Continued)

| Pulsar | $\log_{10} A$ injected | γ injected |
|--------------|------------------------|-------------------|
| J1944 + 0907 | -13.2678 | 1.2939 |
| J2010 - 1323 | -18.9887 | 0.8175 |
| J2017 + 0603 | -17.8317 | 1.9894 |
| J2033 + 1734 | -17.3183 | 3.8762 |
| J2043 + 1711 | -17.5979 | 2.8240 |
| J2145 - 0750 | -12.9290 | 1.2728 |
| J2214 + 3000 | -16.9390 | 4.6045 |
| J2229 + 2643 | -17.4349 | 6.1385 |
| J2234 + 0611 | -13.5827 | 3.6588 |
| J2234 + 0944 | -18.1712 | 0.7777 |
| J2302 + 4442 | -14.7765 | 2.8715 |
| J2317 + 1439 | -18.1829 | 6.5325 |

APPENDIX C: INCORPORATING PAIR COVARIANCE INTO THE OS AND MCOS

In the weak-signal regime, our pulsar pair cross-correlations ρ_{ab} and their associated uncertainties σ_{ab} as described in Eqs. (1) and (2), are given by [21–23,43,44]

$$\rho_{ab} = \mathcal{N}_{ab} \mathbf{r}_a^T \mathbf{P}_a^{-1} \bar{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \mathbf{r}_b \equiv \mathbf{r}_a^T \mathbf{Q}_{ab} \mathbf{r}_b, \quad (\text{C1})$$

$$\sigma_{ab}^2 = \mathcal{N}_{ab} = (\text{tr}[\mathbf{P}_a^{-1} \bar{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \bar{\mathbf{S}}_{ba}])^{-1}, \quad (\text{C2})$$

where \mathbf{P}_a is the autocovariance matrix of pulsar a , $\mathbf{S}_{ab} = A_{\text{gw}}^2 \Gamma_{ab} \bar{\mathbf{S}}_{ab}$ is the cross-covariance matrix between pulsars a and b , and $\mathbf{Q}_{ab} \equiv \mathcal{N}_{ab} \mathbf{P}_a^{-1} \bar{\mathbf{S}}_{ab} \mathbf{P}_b^{-1}$. The autocorrelation matrices \mathbf{P}_a contain contributions from white noise, intrinsic red noise, and the common process red noise, while the cross-correlation matrices \mathbf{S}_{ab} only considers contributions from the common process. The definition of the cross-correlation uncertainties σ_{ab} in Eq. (C2) is only valid in the weak-signal regime and does not include correlations between different pairs of pulsars arising from the fact that some of the pulsar pairs will have pulsars in common (e.g., the correlations ρ_{ab} and ρ_{ac} both have the pulsar a in common).

If a significant correlated signal is present, then σ_{ab} is an underestimate of the cross-correlation uncertainties. When we consider the covariances between different pairs of pulsars, our uncertainties σ_{ab} must include an extra term for inter-pulsar pair covariances. For two pulsar pairs ab and cd [35],

$$\mathbf{\Sigma}_{ab,cd} = \langle \rho_{ab} \rho_{cd} \rangle - \langle \rho_{ab} \rangle \langle \rho_{cd} \rangle \quad (\text{C3})$$

$$= \text{tr}[\mathbf{Q}_{ba} \mathbf{S}_{ac} \mathbf{Q}_{cd} \mathbf{S}_{db}] + \text{tr}[\mathbf{Q}_{ba} \mathbf{S}_{ad} \mathbf{Q}_{dc} \mathbf{S}_{cb}]. \quad (\text{C4})$$

Note that $\mathbf{\Sigma}_{ab,cd}$ is dependent on the ORF and amplitude of the common process, due to the contributions from \mathbf{S}_{ab} .

For the standard single component OS,

$$\mathbf{\Sigma}_{ab,cd}^{\text{OS}} = A_{\text{CURN}}^4 \Gamma_{ac} \Gamma_{db} \text{tr}[\mathbf{Q}_{ba} \bar{\mathbf{S}}_{ac} \mathbf{Q}_{cd} \bar{\mathbf{S}}_{db}] + A_{\text{CURN}}^4 \Gamma_{ad} \Gamma_{cb} \text{tr}[\mathbf{Q}_{ba} \bar{\mathbf{S}}_{ad} \mathbf{Q}_{dc} \bar{\mathbf{S}}_{cb}], \quad (\text{C5})$$

where Γ is the ORF for the correlation, and A_{CURN}^2 is the amplitude of the common uncorrelated red noise found from a Bayesian analysis. We chose to use the amplitude of the common uncorrelated red noise here instead of the amplitude of the correlated process because it. When we include pair covariance, the expressions for the OS and its variance become

$$\hat{A}^2 = \frac{\sum_{ab,cd} \Gamma_{ab} (\mathbf{\Sigma}_{ab,cd}^{\text{OS}})^{-1} \rho_{cd}}{\sum_{ab,cd} \Gamma_{ab} (\mathbf{\Sigma}_{ab,cd}^{\text{OS}})^{-1} \Gamma_{cd}}, \quad (\text{C6})$$

$$\sigma_{\hat{A}^2} = \left[\sum_{ab,cd} \Gamma_{ab} (\mathbf{\Sigma}_{ab,cd}^{\text{OS}})^{-1} \Gamma_{cd} \right]^{-\frac{1}{2}}. \quad (\text{C7})$$

For the MCOS, we generalize Eqs. (C6) and (C7) to allow for arrays of A^2 and Γ :

$$\mathbf{S}_{ab} = \bar{\mathbf{S}}_{ab} \Gamma_{ab} A^2 = \bar{\mathbf{S}}_{ab} \sum_i \Gamma_{ab}^i A_i^2, \quad (\text{C8})$$

where i is the index of the ORF. To obtain the amplitudes of the correlated signals A_i^2 , we first calculate the MCOS values neglecting pair covariance using Eq. (11). If any of the values for \hat{A}^2 are negative, we set it equal to zero. We then set A_i^2 to be equal to a weighted fraction of the common uncorrelated red noise amplitude,

$$A_i^2 = A_{\text{CURN}}^2 \times \frac{A_{\text{MCOS},i}^2}{\sum_j A_{\text{MCOS},j}^2}. \quad (\text{C9})$$

Then the MCOS covariance matrix becomes

$$\mathbf{\Sigma}_{ab,cd}^{\text{MCOS}} = \left(\sum_i A_i^2 \Gamma_{ac}^i \right) \left(\sum_j A_j^2 \Gamma_{db}^j \right) \text{tr}[\mathbf{Q}_{ba} \bar{\mathbf{S}}_{ac} \mathbf{Q}_{cd} \bar{\mathbf{S}}_{db}] + \left(\sum_i A_i^2 \Gamma_{ad}^i \right) \left(\sum_j A_j^2 \Gamma_{cb}^j \right) \times \text{tr}[\mathbf{Q}_{ba} \bar{\mathbf{S}}_{ad} \mathbf{Q}_{dc} \bar{\mathbf{S}}_{cb}], \quad (\text{C10})$$

then, as we did in Sec. II, we can replace the σ_{ab}^2 with our new covariance $\mathbf{\Sigma}_{ab,cd}^{\text{MCOS}}$ and construct our new MCOS B and C matrices

$$C^\beta \equiv \sum_{ab,cd} \rho_{ab} [\mathbf{\Sigma}_{ab,cd}^{\text{MCOS}}]^{-1} \Gamma_{cd}^\beta, \quad (\text{C11})$$

$$B^{\alpha\beta} \equiv \sum_{ab,cd} \Gamma_{ab}^\alpha [\mathbf{\Sigma}_{ab,cd}^{\text{MCOS}}]^{-1} \Gamma_{cd}^\beta, \quad (\text{C12})$$

$$\hat{A}_\alpha^2 = \sum_\beta B_{\alpha\beta} C^\beta, \quad (\text{C13})$$

$$\sigma_{\hat{A}_\alpha^2}^2 = B_{\alpha\alpha}. \quad (\text{C14})$$

- [1] A. Sesana, *Classical Quantum Gravity* **30**, 244009 (2013).
- [2] A. Sesana, F. Haardt, P. Madau, and M. Volonteri, *Astrophys. J.* **611**, 623 (2004).
- [3] S. Burke-Spolaor, S. R. Taylor, M. Charisi, T. Dolch, J. S. Hazboun, A. M. Holgado, L. Z. Kelley, T. J. W. Lazio, D. R. Madison, N. McManm *et al.*, *Astron. Astrophys. Rev.* **27**, 5 (2019).
- [4] M. V. Sazhin, *Sov. Astron.* **22**, 36 (1978), <https://ui.adsabs.harvard.edu/abs/1978SvA....22...36S>.
- [5] S. Detweiler, *Astrophys. J.* **234**, 1100 (1979).
- [6] R. S. Foster and D. C. Backer, *Astrophys. J.* **361**, 300 (1990).
- [7] S. Ransom, A. Brazier, S. Chatterjee, T. Cohen, J. M. Cordes, M. E. DeCesar, P. B. Demorest, J. S. Hazboun, M. T. Lam, R. S. Lynch *et al.* (The NANOGrav Collaboration), *Bull. Am. Astron. Soc.* **51**, 195 (2019), <https://ui.adsabs.harvard.edu/abs/2019BAAS...51g.195R>.
- [8] G. Desvignes, R. N. Caballero, L. Lentati, J. P. W. Verbiest, D. J. Champion, B. W. Stappers, G. H. Janssen, P. Lazarus, S. Osłowski, S. Babak *et al.*, *Mon. Not. R. Astron. Soc.* **458**, 3341 (2016).
- [9] D. J. Reardon, G. Hobbs, W. Coles, Y. Levin, M. J. Keith, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, S. Dai, M. Kerr *et al.*, *Mon. Not. R. Astron. Soc.* **455**, 1751 (2016).
- [10] P. Tarafdar, K. Nobleson, P. Rana, J. Singha, M. A. Krishnakumar, B. C. Joshi, A. K. Paladi, N. Kolhe, N. D. Batra, N. Agarwal *et al.*, *Pub. Astron. Soc. Aust.* **39**, e053 (2022).
- [11] B. B. P. Perera, M. E. DeCesar, P. B. Demorest, M. Kerr, L. Lentati, D. J. Nice, S. Osłowski, S. M. Ransom, M. J. Keith, Z. Arzoumanian *et al.*, *Mon. Not. R. Astron. Soc.* **490**, 4666 (2019).
- [12] R. W. Hellings and G. S. Downs, *Astrophys. J.* **265**, L39 (1983).
- [13] G. Hobbs, W. Coles, R. N. Manchester, M. J. Keith, R. M. Shannon, D. Chen, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, D. Champion *et al.*, *Mon. Not. R. Astron. Soc.* **427**, 2780 (2012).
- [14] C. Tiburzi, G. Hobbs, M. Kerr, W. A. Coles, S. Dai, R. N. Manchester, A. Possenti, R. M. Shannon, and X. P. You, *Mon. Not. R. Astron. Soc.* **455**, 4339 (2016).
- [15] G. Hobbs, L. Guo, R. N. Caballero, W. Coles, K. J. Lee, R. N. Manchester, D. J. Reardon, D. Matsakis, M. L. Tong, Z. Arzoumanian *et al.*, *Mon. Not. R. Astron. Soc.* **491**, 5951 (2020).
- [16] K. J. Lee, F. A. Jenet, and R. H. Price, *Astrophys. J.* **685**, 1304 (2008).
- [17] S. J. Chamberlin and X. Siemens, *Phys. Rev. D* **85**, 082001 (2012).
- [18] J. R. Gair, J. D. Romano, and S. R. Taylor, *Phys. Rev. D* **92**, 102003 (2015).
- [19] N. J. Cornish, L. O’Beirne, S. R. Taylor, and N. Yunes, *Phys. Rev. Lett.* **120**, 181101 (2018).
- [20] M. Anholm, S. Ballmer, J. D. E. Creighton, L. R. Price, and X. Siemens, *Phys. Rev. D* **79**, 084030 (2009).
- [21] P. B. Demorest, R. D. Ferdman, M. E. Gonzalez, D. Nice, S. Ransom, I. H. Stairs, Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin *et al.*, *Astrophys. J.* **762**, 94 (2013).
- [22] S. J. Chamberlin, J. D. E. Creighton, X. Siemens, P. Demorest, J. Ellis, L. R. Price, and J. D. Romano, *Phys. Rev. D* **91**, 044048 (2015).
- [23] S. J. Vigeland, K. Iso, S. R. Taylor, and J. A. Ellis, *Phys. Rev. D* **98**, 044003 (2018).
- [24] Z. Arzoumanian, P. T. Baker, H. Blumer, B. Bécsy, A. Brazier, P. R. Brook, S. Burke-Spolaor, M. Charisi, S. Chatterjee, S. Chen *et al.*, *Astrophys. J. Lett.* **923**, L22 (2021).
- [25] M. F. Alam, Z. Arzoumanian, P. T. Baker, H. Blumer, K. E. Bohler, A. Brazier, P. R. Brook, S. Burke-Spolaor, K. Caballero, R. S. Camuccio *et al.*, *Astrophys. J. Suppl. Ser.* **252**, 4 (2021).
- [26] Z. Arzoumanian, P. T. Baker, H. Blumer, B. Bécsy, A. Brazier, P. R. Brook, S. Burke-Spolaor, S. Chatterjee, S. Chen, J. M. Cordes *et al.*, *Astrophys. J. Lett.* **905**, L34 (2020).
- [27] E. S. Phinney, [arXiv:astro-ph/0108028](https://arxiv.org/abs/astro-ph/0108028).
- [28] F. A. Jenet and J. D. Romano, *Am. J. Phys.* **83**, 635 (2015).
- [29] M. Vallisneri, *Astrophysics Source Code Library p.ascl:2002.017* (2020), <https://ui.adsabs.harvard.edu/abs/2020ascl.soft02017V>.
- [30] J. A. Ellis, M. Vallisneri, S. R. Taylor, and P. T. Baker, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, Zenodo (2020), [10.5281/zenodo.4059815](https://zenodo.org/record/4059815).
- [31] S. R. Taylor, P. T. Baker, J. S. Hazboun, J. Simon, and S. J. Vigeland, *enterprise_extensions* (2021), v2.3.3, https://github.com/nanograv/enterprise_extensions.
- [32] J. Ellis and R. van Haasteren, PTMCMCSampler: Parallel tempering MCMC sampler package written in Python, *Astrophysics Source Code Library*, record ascl:1912.017 (2019).
- [33] N. J. Cornish and L. Sampson, *Phys. Rev. D* **93**, 104047 (2016).
- [34] B. Allen and J. D. Romano, *Phys. Rev. D* **108**, 043026 (2023).
- [35] A. D. Johnson, P. M. Meyers, P. T. Baker, N. J. Cornish, J. S. Hazboun, T. B. Littenberg, J. D. Romano, S. R. Taylor, M. Vallisneri, S. J. Vigeland *et al.*, [arXiv:2306.16223](https://arxiv.org/abs/2306.16223).
- [36] J. S. Hazboun, P. M. Meyers, J. D. Romano, X. Siemens, and A. M. Archibald, *Phys. Rev. D* **108**, 104050 (2023).
- [37] H. Akaike, *Information Theory and an Extension of the Maximum Likelihood Principle* (Springer, New York, NY, 1998), pp. 199–213, ISBN: 978-1-4612-1694-0.
- [38] G. Agazie, A. Anumarlapudi, A. M. Archibald, Z. Arzoumanian, P. T. Baker, B. Bécsy, L. Blecha, A. Brazier, P. R. Brook, S. Burke-Spolaor *et al.*, *Astrophys. J. Lett.* **956**, L3 (2023).
- [39] M. Vallisneri *et al.*, *Astrophys. J.* **893**, 112 (2020).
- [40] E. Roebber *et al.*, *Astrophys. J.* **876**, 55 (2019).
- [41] J. Gair, J. D. Romano, S. Taylor, and C. M. F. Mingarelli, *Phys. Rev. D* **90**, 082001 (2014).
- [42] G. Agazie *et al.* (NANOGrav Collaboration), [arXiv:2310.12138](https://arxiv.org/abs/2310.12138).
- [43] N. Pol, S. R. Taylor, and J. D. Romano, *Astrophys. J.* **940**, 173 (2022).
- [44] X. Siemens, J. Ellis, F. Jenet, and J. D. Romano, *Classical Quantum Gravity* **30**, 224015 (2013).