



Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array

Daniel J. Reardon^{1,2}, Andrew Zic^{3,4}, Ryan M. Shannon^{1,2}, George B. Hobbs³, Matthew Bailes^{1,2},
Valentina Di Marco^{5,6}, Agastya Kapur^{3,4}, Axl F. Rogers⁷, Eric Thrane^{5,6}, Jacob Askew^{1,2}, N. D. Ramesh Bhat⁸,
Andrew Cameron^{1,2}, Małgorzata Curyło⁹, William A. Coles¹⁰, Shi Dai¹¹, Boris Goncharov^{12,13}, Matthew Kerr¹⁴,
Atharva Kulkarni^{1,2}, Yuri Levin^{5,15,16}, Marcus E. Lower³, Richard N. Manchester³, Rami Mandow^{3,4},
Matthew T. Miles^{1,2}, Rowina S. Nathan^{5,6}, Stefan Osłowski¹⁷, Christopher J. Russell¹⁸, Renée Spiewak¹⁹,
Songbo Zhang^{3,20}, and Xing-Jiang Zhu²¹

¹ Centre for Astrophysics and Supercomputing, Swinburne University of Technology, P.O. Box 218, Hawthorn, Victoria 3122, Australia; dreardon@swin.edu.au

² OzGrav: The Australian Research Council Centre of Excellence for Gravitational Wave Discovery, Hawthorn VIC 3122, Australia

³ Australia Telescope National Facility, CSIRO, Space and Astronomy, P.O. Box 76, Epping, NSW 1710, Australia; andrew.zic@csiro.au

⁴ Department of Physics and Astronomy and MQ Research Centre in Astronomy, Astrophysics and Astrophotonics, Macquarie University, NSW 2109, Australia

⁵ School of Physics and Astronomy, Monash University, VIC 3800, Australia

⁶ OzGrav: The Australian Research Council Centre of Excellence for Gravitational Wave Discovery, Clayton VIC 3800, Australia

⁷ Institute for Radio Astronomy & Space Research, Auckland University of Technology, Private Bag 92006, Auckland 1142, New Zealand

⁸ International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia

⁹ Astronomical Observatory, University of Warsaw, Aleje Ujazdowskie 4, 00-478 Warsaw, Poland

¹⁰ Electrical and Computer Engineering, University of California at San Diego, La Jolla, California, USA

¹¹ School of Science, Western Sydney University, Locked Bag 1797, Penrith South DC, NSW 2751, Australia

¹² Gran Sasso Science Institute (GSSI), I-67100 L'Aquila, Italy

¹³ INFN, Laboratori Nazionali del Gran Sasso, I-67100 Assergi, Italy

¹⁴ Space Science Division, US Naval Research Laboratory, 4555 Overlook Ave SW, Washington DC 20375, USA

¹⁵ Physics Department and Columbia Astrophysics Laboratory, Columbia University, 538 West 120th Street, New York, NY 10027, USA

¹⁶ Center for Computational Astrophysics, Flatiron Institute, 162 5th Ave, NY10011, USA

¹⁷ Manly Astrophysics, 15/41-42 East Esplanade, Manly, NSW 2095, Australia

¹⁸ CSIRO Scientific Computing, Australian Technology Park, Locked Bag 9013, Alexandria, NSW 1435, Australia

¹⁹ Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Manchester M13 9PL, UK

²⁰ Purple Mountain Observatory, Chinese Academy of Sciences, Nanjing 210008, People's Republic of China

²¹ Advanced Institute of Natural Sciences, Beijing Normal University, Zhuhai 519087, People's Republic of China

Received 2023 May 8; revised 2023 June 6; accepted 2023 June 8; published 2023 June 29

Abstract

Pulsar timing arrays aim to detect nanohertz-frequency gravitational waves (GWs). A background of GWs modulates pulsar arrival times and manifests as a stochastic process, common to all pulsars, with a signature spatial correlation. Here we describe a search for an isotropic stochastic gravitational-wave background (GWB) using observations of 30 millisecond pulsars from the third data release of the Parkes Pulsar Timing Array (PPTA), which spans 18 yr. Using current Bayesian inference techniques we recover and characterize a common-spectrum noise process. Represented as a strain spectrum $h_c = A(f/1\text{yr}^{-1})^\alpha$, we measure $A = 3.1^{+1.3}_{-0.9} \times 10^{-15}$ and $\alpha = -0.45 \pm 0.20$, respectively (median and 68% credible interval). For a spectral index of $\alpha = -2/3$, corresponding to an isotropic background of GWs radiated by inspiraling supermassive black hole binaries, we recover an amplitude of $A = 2.04^{+0.25}_{-0.22} \times 10^{-15}$. However, we demonstrate that the apparent signal strength is time-dependent, as the first half of our data set can be used to place an upper limit on A that is in tension with the inferred common-spectrum amplitude using the complete data set. We search for spatial correlations in the observations by hierarchically analyzing individual pulsar pairs, which also allows for significance validation through randomizing pulsar positions on the sky. For a process with $\alpha = -2/3$, we measure spatial correlations consistent with a GWB, with an estimated false-alarm probability of $p \lesssim 0.02$ (approx. 2σ). The long timing baselines of the PPTA and the access to southern pulsars will continue to play an important role in the International Pulsar Timing Array.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Gravitational wave astronomy (675); Millisecond pulsars (1062); Pulsar timing method (1305); Bayesian statistics (1900)

1. Introduction

The era of observational gravitational-wave (GW) astronomy commenced with the observation of a stellar-mass binary black hole merger (Abbott et al. 2016). The growing catalog of GW events from ground-based observatories covers relatively high frequencies in a band spanning $\sim 10^1$ Hz to $\mathcal{O}(10^3)$ Hz (The

LIGO-Virgo-KAGRA Collaboration et al. 2021). Pulsar timing array (PTA) experiments offer a complementary window into the GW landscape in the nanohertz-frequency band at frequencies $\mathcal{O}(10^{-9})$ Hz. Sources of GWs at such low frequencies may include supermassive black hole binary (SMBHB) systems (e.g., Rajagopal & Romani 1995; Wyithe & Loeb 2003; Sesana 2013; Ravi et al. 2015; Burke-Spolaor et al. 2019), primordial quantum fluctuations amplified by inflation (e.g., Grishchuk 2005; Lasky et al. 2016), cosmic strings (e.g., Siemens et al. 2007; Ölmaz et al. 2010), or cosmological phase transitions (e.g., Caprini et al. 2010; Xue et al. 2021). The highest-



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

amplitude (loudest) signal in this band is expected to be the stochastic background produced by the incoherent superposition of GWs radiated by inspiraling SMBHBs. An isotropic stochastic gravitational-wave background (GWB) is the primary target for GW searches by current PTA collaborations.

A PTA consists of a set of millisecond pulsars (MSPs) that are monitored regularly (cadences of many days to a few weeks), over decade timescales, to take advantage of their long-term rotational stability (Foster & Backer 1990). As GWs impart a strain on space-time, they slowly modulate the observed pulse times of arrival (ToAs) from MSPs in an achronic manner (Sazhin 1978; Detweiler 1979). The ToA modulations induced by an isotropic GWB are predicted to possess a steep power spectral density (PSD) $P \propto f^{-\gamma}$, where f is the GW frequency. This spectrum is derived from the GW strain amplitude spectrum $h_c = A(f/1\text{yr}^{-1})^\alpha$, where $\alpha = (3 - \gamma)/2$. The characteristic spectral index of an isotropic GWB from inspiraling SMBHBs is $\alpha = -2/3$ and thus $\gamma = 13/3$ (Phinney 2001). The challenge in detecting this signal is that MSPs are not perfectly stable in their spin and pulse properties and may show intrinsic low-frequency noise that could resemble a GWB for individual pulsars (Shannon & Cordes 2010). Additionally, plasma in the interstellar medium has a significant effect on radio pulses as it imparts dispersion and scattering delays that depend on the radio wavelength as λ^2 and λ^{-4} , respectively (e.g., You et al. 2007; Hemberger & Stinebring 2008). Variations in these propagation effects can also produce a red signal in the pulse arrival times (e.g., Cordes & Shannon 2010; Keith et al. 2013). Other low-frequency noise sources can include errors in the solar system ephemeris (SSE), clock errors, unmodeled variability in the solar wind, or offsets induced by the observing instrumentation (Tiburzi et al. 2016).

The distinct characteristic of an isotropic GWB, which separates it from intrinsic pulsar spin noise and other noise processes, is the spatial correlations in the arrival times. The Hellings–Downs function describes the expected correlation in the pulse arrival times between pairs of pulsars, as a function of their sky separation angle (Hellings & Downs 1983). It is the *fingerpr*int of the GWB and is presently the most sought-after signal for a PTA experiment. The sensitivity of a PTA to such spatial correlations depends primarily on the number of pulsar pairs and their sky distribution (Siemens et al. 2013; Taylor et al. 2016), which motivates the combination of global PTA data sets under the International Pulsar Timing Array (IPTA; Hobbs et al. 2010; Antoniadis et al. 2022) project. Future observations in the high signal-to-noise regime of the GWB may reveal anisotropy resulting from an unresolved local population of SMBHBs (Taylor & Gair 2013; Mingarelli et al. 2017). Individual supermassive black hole sources may also be observed, with orbital properties inferred through template matching of the GW waveform (Jenet et al. 2004; Corbin & Cornish 2010; Ellis 2013; Zhu et al. 2015).

The Hellings–Downs signature cross correlation is relatively weak, with the average amplitude of the cross correlation being well under 0.25, depending on the distribution of the pulsars on the sky. As the number of pulsars in the array increases, along with the time span and advances in instrumentation that improve the timing precision, the sensitivity of a PTA also increases. We would expect, eventually, to see the same GWB spectrum in the power spectrum for each pulsar, rising above the noise level (which is different for each pulsar). However, the cross correlations may not become evident until the sensitivity of

the PTA increases significantly (Pol et al. 2021; Romano et al. 2021). Exactly this situation occurred when the North American Nanohertz Observatory for gravitational waves (NANOGrav) detected a “common-spectrum” process in their pulsars using 12.5 yr of MSP timing data, but was unable to establish a significant cross correlation between pulsar pairs (Arzoumanian et al. 2020). A similar common-spectrum process was also detected by the other PTAs without obvious cross correlations (Goncharov et al. 2021b; Chen et al. 2021). The IPTA, using a union of earlier data releases from each of these PTAs, also identified a common-spectrum process, notably with higher significance than the individual PTA constituents of the data set (Antoniadis et al. 2022).

The origin of the common-spectrum process is unclear, in part because these recent results were in apparent tension with the 95% confidence upper limits placed by some PTA analyses using earlier data releases (e.g., Shannon et al. 2015; Arzoumanian et al. 2018). Unmodeled SSE errors (Vallisneri et al. 2020), choices of noise models and priors (Hazboun et al. 2020a), and finite numbers of pulsars (Johnson et al. 2022) have been proposed to have been factors in these earlier upper-limit results. The NANOGrav analysis, and earlier works (Arzoumanian et al. 2018; Vallisneri et al. 2020), demonstrated that SSE errors could be contributing to the measured common noise, and consequently, techniques were developed to account for such errors. The PPTA analysis found that the “common process” model does not discriminate between common and spectrally similar noise processes (Zic et al. 2022), which was addressed in Goncharov et al. (2022). Depending on the noise modeling framework, a recovered common-spectrum process could contain various non-GW sources of noise. For example, the common spectrum could include pulsar timing noise and/or errors in the correction of interstellar medium effects (dispersion and scattering), which would be uncorrelated between pulsars and could have a similar power-law red spectrum. It could also include SSE errors, which would be correlated between pulsars and could have a wide range of spectral distributions (Guo et al. 2019; Vallisneri et al. 2020). Evidently, the origin of the common spectrum will not be resolved until significant cross correlation is obtained.

For this work we have used the third data release of the PPTA (Zic et al. 2023) to search for and characterize a common-spectrum stochastic process and spatial correlations, with particular focus on the GWB. In a coordinated effort, the NANOGrav, EPTA and Indian Pulsar Timing Array (Joshi et al. 2018), and Chinese Pulsar Timing Array (Lee 2016) collaborations have also recently searched their respective data sets for a GWB (Antoniadis et al. 2023; Arzoumanian et al. 2023; Xu et al. 2023). In Section 2, we describe the data analysis methods and the models we consider. The results are presented in Section 3, and the implications are discussed further in Section 4. We conclude in Section 5.

2. Methods

The analysis presented here is performed on the third PPTA data release (PPTA-DR3) and individual pulsar noise analyses. The data release adds ~ 3 yr of timing baseline compared to PPTA-DR2 (Kerr et al. 2020) and includes observations taken with the new ultrawide-band low receiver (UWL; Hobbs et al. 2020), which offers wider bandwidths than previous observations. This is useful for measuring interstellar medium effects and obtaining arrival times with higher precision. The data

release (Zic et al. 2023) and noise analyses (Reardon et al. 2023) are presented in companion papers. While PPTA-DR3 contains observations from 32 pulsars, in this work we only consider 30 of them. The globular cluster MSP PSR J1824–2452A was excluded as it contains steep-spectrum red noise that is too strong for this pulsar to be sensitive to a common process. The noise is likely intrinsic to the pulsar (Shannon & Cordes 2010), although globular-cluster dynamics may contribute as well. We also exclude PSR J1741 + 1351 from the analysis. It was only added to the PPTA as the UWL was commissioned, and observed with low priority. In the current data release, only 16 observations were available, resulting in a data set that would not be sensitive to any GWB signal. These observations could, however, become useful as part of future IPTA data sets.

2.1. Inference and PTA Likelihood

We use Bayesian inference²² to search for and characterize noise and signals in our data set, following previous PTA analyses (Arzoumanian et al. 2020; Goncharov et al. 2021b; Chen et al. 2021; Antoniadis et al. 2022). PTAs observe offsets between the measured pulse arrival times δt and the arrival times predicted by the models. The time-domain likelihood of the observed residuals given a model is described by a Gaussian likelihood function, $\mathcal{L}(\delta t|\theta)$, where θ is the vector of model parameters (van Haasteren et al. 2009). The likelihood is multivariate with respect to the number of observations of timing residuals. The implementation of the likelihood is described in Arzoumanian et al. (2016) and Taylor et al. (2017), and it is given by

$$\mathcal{L}(\delta t|\theta) = \frac{\exp\left(-\frac{1}{2}(\delta t - \mu)^T C^{-1}(\delta t - \mu)\right)}{\sqrt{\det(2\pi C)}}, \quad (1)$$

where $\mu(\theta)$ represents the model prediction for timing residuals, and the covariance matrix is $C(\theta)$. The diagonal elements in the covariance matrix represent the temporally uncorrelated noise, which is referred to as white. Temporally correlated stochastic processes, referred to as red, could have been represented by off-diagonal elements in C . However, to avoid a computationally expensive matrix inversion, temporally correlated stochastic processes are modeled as Gaussian processes (Lentati et al. 2013; van Haasteren & Vallisneri 2014) in μ :

$$\mu = Fa + M\epsilon, \quad (2)$$

where the design matrix M and the corresponding coefficients ϵ are the timing model contributions, and the red processes are modeled using the Fourier sine and cosine basis functions in F and amplitudes a . We emphasize that $(a$ and $\epsilon)$ are a subset of the parameters in the model θ . They are nuisance parameters and can be analytically marginalized over a Gaussian prior $\pi(a, \epsilon|\theta')$. It is common to assume that there is a relationship between the amplitudes of the Fourier components. In this case, θ' are hyperparameters that govern the spectra of temporal correlations, such as the power-law amplitude and the slope of $P(f)$ of pulsar-intrinsic red noise, or the gravitational-wave background. The prior on a joins single-pulsar likelihoods into a

joint posterior distribution

$$\mathcal{P}(\theta, \theta'|\delta t) \propto \mathcal{L}(\delta t|\theta)\pi(a, \epsilon|\theta')\pi(\theta'). \quad (3)$$

The covariance matrix of $\pi(a, \epsilon|\theta')$ has elements

$$\pi_{(ai),(bj)} = P_{ai}\delta_{ab}\delta_{ij} + \Gamma_{ab}P_i\delta_{ij}, \quad (4)$$

where a and b are pulsar indices, and i and j are frequency indices. P_{ai} is the PSD of the noise intrinsic to pulsar a at a frequency bin i , P_i is the PSD of a spatially correlated signal at the frequency i , and Γ_{ab} is the overlap reduction function that determines the degree of this correlation between pulsars a and b . In case of an isotropic GWB from circular SMBHBs, Γ_{ab} , also referred to as the overlap reduction function, is given by the Hellings & Downs (1983) curve:

$$\Gamma_{ab} = \frac{1}{2}\delta_{ab} + \frac{1}{2} - \frac{x_{ab}}{4} + \frac{3}{2}x_{ab}\ln x_{ab}, \quad (5)$$

where $x_{ab} = (1 - \cos \zeta_{ab})/2$, δ_{ab} is the Kronecker delta function, and ζ_{ab} is the sky separation angle for a given pair of pulsars.

We evaluate the posterior probability $\mathcal{P}(\theta, \theta'|\delta t)$ using the ENTERPRISE package (Ellis et al. 2019) and a parallel-tempered Markov Chain Monte Carlo (MCMC) algorithm (PTMCMCSAMPLER; Ellis & van Haasteren 2019). For more details about modeling Gaussian processes in PTAs, see Lentati et al. (2013) and van Haasteren & Vallisneri (2014).

Previous GWB analyses have assessed the properties of the dominant temporally correlated and spatially uncorrelated *common-spectrum* process. This process was considered as the null hypothesis in searches for the spatially correlated GWB signal. The PSD of temporal correlations can be modeled as a power law parameterized by an amplitude (A) and spectral index:

$$P(f) = \frac{A^2}{12\pi^2} \left(\frac{f}{f_{\text{yr}}} \right)^{-\gamma}, \quad (6)$$

where $\gamma > 0$. The models of common noise in this form are added to the PTA likelihood and analyzed simultaneously with all red-noise processes identified in single pulsars. The support for the proposed common-noise processes is quantified primarily through the Bayesian odds ratio (equivalent to the Bayes factor $\mathcal{B}$ for equal prior odds), which can be estimated using the product-space sampling method (Arzoumanian et al. 2016; Carlin & Chib 2018).

The number of Fourier frequency components used to model the common power-law processes can be set by the sensitivity of our PTA, which is white-noise dominated near $f \sim 1/(240 \text{ days})$ (as determined based on broken power-law analysis, e.g., Arzoumanian et al. 2020). Accordingly we set the number of components as $N_{\text{comp}} = \lfloor T_{\text{span}}/(240 \text{ days}) \rfloor = 28$, where $T_{\text{span}} = 6605 \text{ days}$ is the total observing span of PPTA-DR3. We note that the use of a broken power-law model to set the number of significant Fourier components (Arzoumanian et al. 2020) could highlight excess unmodeled (non-power-law) noise at higher frequencies. If the noise is correctly characterized, and the signal is a power law, then the data will not favor a turnover frequency and should instead become insensitive to a turnover for frequencies that are white-noise dominated.

Under the model of a fixed spectral index $\gamma = 13/3$ (the spectral index expected for a GWB produced by circular

²² For an overview of the principles of Bayesian inference in gravitational-wave astronomy, see Thrane & Talbot (2019) and Taylor (2021).

SMBHBs, with their inspirals driven by GW radiation at the frequencies of interest; Phinney 2001) and zero spatial correlation, the PTA likelihood can be factorized into a product of likelihoods from individual pulsars (e.g., Taylor et al. 2022). In this way, a $\gamma = 13/3$ process can be introduced to the noise models of individual pulsars and the resulting posterior probability density functions (PDFs) for the process amplitude can be multiplied to produce a posterior for the whole array. We use this factorized-likelihood approach to efficiently assess the preferred amplitude of a $\gamma = 13/3$ process in our data, and to estimate the model support under an assumed prior range using the Savage–Dickey method for computing the Bayes factor $\mathcal{B}$ (Dickey 1971).

Hereafter, when referring to specific models, we label the amplitude of an uncorrelated common-red-noise (CRN) process as A^{CRN} , and when the spectral index is fixed at $\gamma = 13/3$, we write $A_{13/3}^{\text{CRN}}$. For a $\gamma = 13/3$ process in a single pulsar (i.e., not common to the PTA), we use $A_{13/3}$. A process with Hellings–Downs correlations included in the covariance matrix is denoted A^{HD} with a varied spectral index and A^{GWB} when the index is $\gamma = 13/3$.

2.1.1. Defining Prior Ranges

The choice of prior has a dramatic effect on the Bayesian evidence of common-noise processes, and therefore Bayes factors used for model comparison (Zic et al. 2022). It is important for priors to encompass the likelihood support, but prior ranges that are too broad have been shown to induce spurious support for certain models. This problem affects the inference of common uncorrelated noise processes in PTAs, and can lead to very high Bayesian support for a common process where none exists (Zic et al. 2022). To mitigate the possibility of spurious support, we trialed different priors in our analysis. For the factorized-likelihood analysis we adopt a log-uniform prior on the amplitude in the range $-20 \leq \log_{10} A_{13/3} \leq -11$ to be consistent with earlier analyses (Arzoumanian et al. 2020; Goncharov et al. 2021b). However, our data are not sensitive to signals below $\log_{10} A_{13/3} \lesssim -16.5$ even for very steep spectral indices ($\gamma \sim 7$). We therefore adopt $\log_{10} A_{13/3} \geq -18$ as a conservative lower bound for the remainder of our analyses. The prior for the common-noise spectral index is uniform in $0 \leq \gamma \leq 7$.

To set the prior ranges for the individual pulsar noise parameters, we take the central 99.7% credible interval from our single-pulsar noise analysis posteriors (Reardon et al. 2023), and add a generous buffer region such that pulsar noise properties are allowed to vary substantially under a common-noise model. These ranges are $\log_{10} A_{0.15} - 2 \leq \log_{10} A \leq \log_{10} A_{99.85} + 1$ and $\gamma_{0.15} - 0.5 \leq \gamma \leq \gamma_{99.85} + 0.5$, where the parameter subscripts denote the percentile of the posterior distribution for this parameter. This ensures that we do not force any pulsars to have intrinsic noise if it is instead attributed to a common process and the prior range encompasses all likelihood support from the data. PSR J1643–1224 is the only pulsar with a lower bound $\log_{10} A \geq -17$, with a red noise prior constructed in this way of $-15.2 < \log_{10} A < -11.5$ as it is observed to have shallow-spectrum noise ($\gamma \leq 2.3$ with 99.7% confidence). By restricting the prior ranges in this way, we reduce the risk of spurious measurements (Zic et al. 2022) of common noise and are able to accelerate the inference. We are confident in the recovered parameters using this methodology as the shapes and values of

the inferred marginal posterior PDFs using these prior ranges are consistent with those recovered using standard priors.

2.2. Noise and Signal Models

2.2.1. Common Processes

The term common noise refers to any noise process (or potentially unmodeled deterministic signal), correlated or uncorrelated, that is characterized by the same power spectral density for all pulsars. However, common noise could be recovered from similar power spectral densities in most pulsars, and as such, examples include not only GWs but also potentially unmodeled instrumental offsets, the solar wind; or errors in the clock correction, Earth orientation parameters, or SSE. If intrinsic pulsar timing noise (rotational irregularities) is generated by the same physical mechanism in multiple pulsars, the spectral properties of this random process could be recovered via inference of a common-spectrum noise process but should be spatially uncorrelated, unlike the GWB. The variety of potential common-noise sources highlights the need to measure the signature spatial correlations in order to detect a GWB. The power spectrum of a power-law common uncorrelated process is given by Equation (6).

2.2.2. Searching for Spatial Correlations

We conduct two distinct classes of searches for the Hellings–Downs correlations inherent to a GWB. First, we replicate the methods of some previous searches (e.g., Arzoumanian et al. 2020; Goncharov et al. 2021b; Chen et al. 2021; Antoniadis et al. 2022), by constructing the global PTA likelihood for all pulsars and introducing a Hellings–Downs correlated noise term to the likelihood (see Equation (4)). For this model, both autocorrelation and cross correlation are considered and inform the parameter inference. The sensitivity of the PPTA data set is predicted to be dominated by the autocorrelations (Spiwak et al. 2022), so we could expect the spectral properties of such a model to closely resemble the uncorrelated common process. While pulsar white-noise terms are usually fixed at their maximum likelihood values (as they are less covariant with the GWB), we numerically marginalized over all red-noise terms for all pulsars in the array, as these noise terms could be highly covariant with any common signal. This results in a high-dimension parameter space that needs to be explored. In the case of PPTA-DR3 the parameter space spans 260 dimensions.

We also present a new search method that increases the efficiency of the inference with a hierarchical inference approach. This is important because with our large data volume ($\sim 1.2 \times 10^5$ ToAs) and parameter space (~ 260 free parameters) it takes of order CPU years to collect a sufficient number of independent MCMC samples for standard full-PTA correlation analyses, which makes it costly to determine false-alarm probabilities with bootstrap methods. We first search for common-spectrum noise process in pulsar pairs individually. We measure the $\log_{10} A_{13/3}$ and correlation coefficient, Γ , for each of the 435 unique pulsar pairs ($N_{\text{pair}} = N_{\text{psr}}(N_{\text{psr}} - 1)/2$, for N_{psr} pulsars) in our data set. From the posterior samples, we derive a smooth PDF, $p(\log_{10} A_{13/3}, \Gamma)$, using a two-dimensional Gaussian kernel density estimate (KDE). We account for the bounded domains $-1 \leq \Gamma \leq 1$, and $-18 \leq \log_{10} A \leq -14$ by mirroring the samples at these boundaries to partially overcome bias in the density estimator introduced by the boundary discontinuity (Schuster 1985). The

standard deviation of the Gaussian kernel (for all pairs) was $\sigma_{\text{KDE}}(\log_{10}A) = 0.05$ for amplitude (the standard deviation of our measured common process) and $\sigma_{\text{KDE}}(\Gamma) = 0.27$ for correlation coefficient (the minimum observed standard deviation for an individual pair in our sample). We also tried using a different $\sigma_{\text{KDE}}(\Gamma)$ kernel width for each pair, computed using Silverman’s rule (because we observe that the distributions are unimodal; Silverman 1986), and we observe consistent results.

From these resulting PDFs, we reweight the posterior samples (using a form of importance sampling; Gelman et al. 1995; Goncharov et al. 2022; Hourihane et al. 2023; Taylor et al. 2022) to recover the probability density for Γ that corresponds to the posterior probability density of $\log_{10}A_{13/3}^{\text{CRN}}$ that we measured from our factorized-likelihood analysis. After this reweighting, the 435 PDFs for Γ , denoted $p(\Gamma)$, assume a common distribution on $\log_{10}A$. These PDFs can then be analyzed to assess the model support for correlations such as the Hellings–Downs ORF expected of the isotropic GWB. The measurements will not be completely independent, and we address the implications of this below.

We also measure $p(\Gamma)$ for the amplitude fixed at our maximum likelihood $\log_{10}A_{13/3}^{\text{CRN}}$, which is equivalent to the reweighted measurements if the correlations do not evolve across the uncertainty of the recovered common noise. These fixed-amplitude measurements are more computationally efficient, particularly for pulsars with low likelihood support at the amplitude of interest. We therefore obtain these fixed-amplitude correlation measurements from independent MCMC chains so that they can also serve as a consistency test, sensitive to errors in the KDE due to a finite number of samples.

We quantify the support for a Hellings–Downs ORF over uncorrelated common noise by computing the likelihood of each model. The Hellings–Downs likelihood is defined to be the product of the probabilities of observing $\Gamma = \Gamma_{ab}$ for each pair

$$\mathcal{L}^{\text{HD}} = \prod_{k=1}^{N_{\text{pair}}} p_k(\Gamma = \Gamma_{ab}), \quad (7)$$

where the k^{th} pair involves pulsars a and b , $N_{\text{pair}} = 435$, and Γ_{ab} is given by Equation (5). Similarly, the likelihood of zero correlations for the observed common noise is

$$\mathcal{L}^{\text{CRN}} = \prod_{k=1}^{N_{\text{pair}}} p_k(\Gamma = 0). \quad (8)$$

The support for Hellings–Downs over zero correlations can be quantified using the likelihood ratio $\Delta\mathcal{L}_{\text{CRN}}^{\text{HD}} = \mathcal{L}^{\text{HD}}/\mathcal{L}^{\text{CRN}}$. This could be interpreted as an estimated Bayes Factor $\mathcal{B}_{\text{CRN}}^{\text{HD}}$ because the ORFs are effectively zero-parameter models. However, because of the weaknesses in this method (e.g., assuming the measured correlations are independent, and KDE biases), the *ad-hoc* estimator $\Delta\mathcal{L}_{\text{CRN}}^{\text{HD}}$ is not as well-defined as a Bayes factor, and we instead interpret its statistical significance by generating noise realizations from the data itself.

The method allows for the rapid computation of false-alarm probabilities by recalculating the likelihood statistics after assuming random positions for the pulsars on the sky. This process is referred to as sky scrambling (Cornish & Sampson 2016; Taylor et al. 2017). The number of possible independent skies (leading to orthogonal sets of samples from the Hellings–Downs ORF) is limited to $N_{\text{sky}} = N_{\text{psr}}(N_{\text{psr}} - 1)$, for N_{psr} pulsars

with equal sensitivity. However, we are able to generate a larger number of skies by relaxing the strict requirement for independence. We instead consider randomized skies where the normalized inner product of the ORF samples from any two skies is ≤ 0.2 (e.g., Cornish & Sampson 2016; Taylor et al. 2017).

Our hierarchical inference scheme can also be simply extended to search for anisotropy in the GWB (Taylor & Gair 2013), or to search for a signal with a completely different spatial correlation signature. It could also be used to estimate the amplitude of a correlated signal independently from the autocorrelations (provided a reliable reconstruction of the full $(\log_{10}A, \Gamma)$ parameter space can be made for each pair).

Pulsars have different sensitivities to common-noise processes due to variations in their timing precision and intrinsic red-noise properties. As a result, the effective number of pulsar pairs in the array is less than the total. We can also use the PDFs for Γ to compute the number of effective pulsar pairs n_{eff} using

$$n_{\text{eff}} = \frac{\left(\sum_{k=1}^{N_{\text{pair}}} \sigma_k^{-2}\right)^2}{\sum_{k=1}^{N_{\text{pair}}} \sigma_k^{-4}}, \quad (9)$$

where σ_k is the uncertainty in the angular correlation for the k^{th} pulsar pair. We define σ_k as the half-width of the 68% credibility interval of the correlation coefficient posterior for pulsar pair k . This is more conservative than using the standard deviation of the posterior samples as σ_k because the distributions are often asymmetric.

3. Results

3.1. Factorized-likelihood Analysis

We first searched for a purported common red process using the factorized-likelihood technique (Taylor et al. 2022). For each pulsar, we include an achromatic process with a fixed spectral index of $\gamma = 13/3$, while the (log) amplitude for this process, with a prior in the range $-20 \leq \log_{10}A_{13/3} \leq -11$, was sampled simultaneously with the full single-pulsar noise models. The posterior probability density for each pulsar (light-gray and colored lines) is shown in Figure 1. The probability density for the log amplitude of a *common process* is found by taking the product of these individual pulsar posteriors. In this way, the log amplitude can be interpreted as a probability-weighted mean. The resulting distribution gives $\log_{10}A_{13/3}^{\text{CRN}} = -14.69 \pm 0.05$, where, here and throughout, the uncertainties represent the central 68% credibility interval unless otherwise stated.

The PPTA data as a whole support the inclusion of a $\gamma = 13/3$ process at this amplitude, with a strong Savage–Dickey Bayes factor, $\log_{10}\mathcal{B} = 6.8$. By reducing the prior range to our nominal $-18 \leq \log_{10}A_{13/3} \leq -11$, we find that the evidence decreases to $\log_{10}\mathcal{B} = 4.8$, indicating that this statistic alone is a poor metric for determining if a common-noise process is present in the array. We stress that, while the evidence for a common-noise process varies under different assumed prior ranges, the recovered characteristics do not, because the likelihood support under a common-noise model is encompassed for all pulsars for all prior ranges considered.




Figure 1. Factorized-likelihood analysis on the $\log_{10} A_{13/3}$ from the PPTA-DR3 pulsars. We highlight the three pulsars showing the highest (PSRs J1909–3744, J0437–4715, and J2145–0750) and lowest (PSRs J1744–1134, J1603–7202, and J1713+0747) consistency with the inferred common process with colored histograms. Other pulsars are presented as light-gray histograms. The prior density is shown with the green dashed line, and the total factorized-likelihood constraint for $\log_{10} A_{13/3}$ (the product of all other histograms) is shown by the black line.

3.2. Noise Consistency Metrics

Next, we assessed the consistency of each of the pulsars with the common red process. Using the posterior distributions in Figure 1, we define metrics to assess the consistency of individual pulsars with the common process. We compute posterior probability density ratios, which compare the probability taken at the measured common amplitude $p(\log_{10} A_{13/3} = -14.69)$, against both the prior probability density $\pi(\log_{10} A_{13/3}^{\text{CRN}})$ and the mean probability at low (effectively zero) amplitude $p(\log_{10} A_{13/3} < -16.5)$. The probability ratios, shown in the top panel of Figure 2, are intended to portray similar information to the “dropout factors” used in previous works (e.g., Arzoumanian et al. 2020; Goncharov et al. 2021b). We also compute the Savage–Dickey Bayes factor, $\mathcal{B} = \pi(\log_{10} A_{13/3})/p(\log_{10} A_{13/3} < -16.5)$, which provides the support from each pulsar for a $-13/3$ process at any amplitude (bottom panel in Figure 2). The three pulsars with the highest $\mathcal{B}$ and the three with the lowest $\mathcal{B}$ are highlighted with unique colors in Figure 1. The three pulsars that show the highest support show evidence for red noise in single-pulsar noise analyses and also are likely to dominate the factorized likelihood. The three pulsars with the lowest $\mathcal{B}$ show no evidence for red noise in single-pulsar analyses. The posterior distributions for these pulsars show negative support for the CRN at the inferred amplitude. This is discussed further below.

3.3. Common Uncorrelated Process

We next search for a single uncorrelated common process with a variable spectral index. In contrast to the factorized-likelihood analysis discussed previously, this process is included for the PTA as a whole and sampled simultaneously

with all red components of the single-pulsar noise models. The characteristics of this common process are considered for multiple SSEs from the Jet Propulsion Laboratory (JPL; DE421, DE438, and DE440), with our fiducial results generated with the most recent DE440 ephemeris (Park et al. 2021). The one and two-dimensional marginal posterior probability distributions for the (log) amplitude and spectral index for the recovered process are shown in Figure 3. All ephemerides return consistent constraints for the process, with DE440 providing $\gamma = 3.9 \pm 0.4$ and $\log_{10} A^{\text{CRN}} = -14.50^{+0.14}_{-0.16}$. While DE421 is an obsolete SSE that is inaccurate for the Jovian system, the recovered common noise is consistent, suggesting that the spectral characteristics are dominated by another source. Under the assumption of a fixed spectral index of $\gamma = 13/3$ we recover the same amplitude constraint as the factorized-likelihood analysis, as expected.

We also searched for a common uncorrelated process using commonly used broad priors on single-pulsar noise terms ($-20 \leq \log_{10} A \leq -11$, $0 \leq \gamma \leq 7$) under the DE440 ephemeris. We compare this with our distribution recovered under tailored priors for single-pulsar noise terms in Figure 4. The distributions are consistent, and for the broader priors we recover $\gamma = 4.0^{+0.3}_{-0.3}$ and $\log_{10} A^{\text{CRN}} = -14.52^{+0.14}_{-0.15}$. The differences are comparable to those induced by SSE choice and much smaller than those induced by weakness in single-pulsar noise models (Reardon et al. 2023).

3.4. Time-dependence of the Common Process Amplitude

Next we search for any time evolution of the CRN process. With the longest multiwavelength timing baseline and high timing precision on many pulsars, the PPTA data provides the best opportunity to assess the stationarity of any purported




Figure 2. Statistics assessing the support for a noise process with $\gamma = 13/3$ for each pulsar in PPTA-DR3. Top: posterior probability density ratio for a $\gamma = 13/3$ process, at $\log_{10} A_{13/3} = -14.69$, relative to $\log_{10} A_{13/3} < -16.5$ (where the data are insensitive to a $\gamma = 13/3$ process) in blue, and the prior density in orange. Bottom: Savage–Dickey Bayes factor $\mathcal{B}$ for a $\gamma = 13/3$ process at any amplitude, which is also the ratio of the orange to blue points in the top panel. Filled and hollow circles correspond to prior range lower bounds on $\log_{10} A_{13/3}$ of -20 and -18 , respectively.

common red signal. Our motivation is to investigate the reasons why the previous upper limits on the GWB from all PTAs were lower than the amplitude of the purported common noise being detected in recent and current data sets. This includes a limit of $A_{13/3}^{\text{CRN}} < 1.2 \times 10^{-15}$ (95% confidence), which we place using the first half of our data set (described below). This limit can be considered a replacement of the limit placed using a subarray of four PPTA pulsars with data extending to the beginning of 2015 (Shannon et al. 2015), which used the now-obsolete DE421 SSE (Folkner et al. 2009).

We estimate the amplitude of the common noise in 6 and 9 yr sliding windows (*slices*) assuming $\gamma = 13/3$ and present the results in Figure 5. In the early time windows, marked in orange, a common-noise process with the $\gamma = 13/3$ spectrum was not detected, and we have shown the 95%-confidence upper limit for the amplitude. In the later windows, marked in blue, a detection was made, and the full 95% credible interval was given. It is clear that these bounds on the common-noise increase with time, and in the earlier data, they are broadly consistent with earlier upper limits (e.g., Shannon et al. 2015; Arzoumanian et al. 2018). The observational systems improved significantly with time, but it is hard to avoid the conclusion that, if the common noise represents a single distinct physical

process such as a GWB, then the process is not time stationary. However, we cannot at this point rule out a nonastrophysical origin for this effect.

The earlier parts of the data used less powerful signal processing equipment, resulting in reduced sensitivity (and higher upper limits) due to lower received observing bandwidth and greater levels of digital artifacts in the data (Manchester et al. 2013; Kerr et al. 2020). In intermediate windows (dates centering between 2010 and 2012), which included data with the more sensitive last-generation narrowband (pre-UWL) digital receiving systems, we also find no evidence of a common red process. The 95% upper limits were computed by reweighting the posterior samples to a prior that is uniform in linear amplitude. We infer upper limits from these intermediate windows ranging from $A_{13/3}^{\text{CRN}} < 1.2 \times 10^{-15}$ to $A_{13/3}^{\text{CRN}} < 1.5 \times 10^{-15}$, which are below the measured $A_{13/3}^{\text{CRN}} = 2.0 \pm 0.2 \times 10^{-15}$ from the full data set. The upper limit from the first 9 yr ($A_{13/3}^{\text{CRN}} < 1.2 \times 10^{-15}$) slice corresponds with approximately the same data span and time range considered in Shannon et al. (2015). The similarly tight bound that we derive consolidates this result and other similar upper limits (e.g., Arzoumanian et al. 2018). The amplitude inferred from the most recent 6 yr slice is $A_{13/3}^{\text{CRN}} = 2.7^{+0.8}_{-0.7} \times 10^{-15}$. This is higher than, but still consistent




Figure 3. Marginal posterior probability distributions for the measured logarithmic amplitude and spectral index (γ) of a common uncorrelated process assuming different solar system ephemerides (SSE). Using DE440 (green) we measure $\gamma = 3.87 \pm 0.36$ and $\log_{10} A^{\text{CRN}} = -14.50^{+0.14}_{-0.16}$ (median and 68% credible interval; shaded regions in one-dimensional histograms). The contours on the two-dimensional marginal distribution show the 1σ , 2σ , and 3σ credible intervals for each of the SSEs.




Figure 4. Marginal posterior probability distributions for the logarithmic amplitude and spectral index (γ) of a common uncorrelated process, under our tailored single-pulsar priors (blue) and broader equal priors (orange).

with, the $A_{13/3}^{\text{CRN}}$ measurement from the full data span. Johnson et al. (2022) found that the use of a small number of pulsars (as in Shannon et al. 2015) can increase the chance of

underestimating upper limits on the amplitude of a common process. However, this is not the origin for the apparent time variability that we observe, as we use all available pulsars in each slice (i.e., 23 pulsars for the first slice, and 30 pulsars in the last) and the same priors and noise models across the slices.

Using the first three 9 yr windows, the upper limits on the amplitude of the signal are inconsistent with the amplitude measured in the entire data set, and in later subsets. The measured uncorrelated common-spectrum amplitude $\log_{10} A_{13/3}^{\text{CRN}} = -14.69$ lies at the 99.8 percentile of the reweighted samples from the first 9 yr window. While some variability would be expected because of the stochasticity of the background (Hazboun et al. 2020b), this level of variation is extreme, and the implications are discussed below.

3.5. Common Free Spectrum

To assess whether or not the apparent common red signal is consistent with a power-law process, we formed the free spectrum. As with power-law processes, the free spectrum is modeled as a Fourier series but the variances of each Fourier coefficient are independent (i.e., not constrained to follow a power-law process).

We formed the free spectrum using multiple methods. First we modeled it as a common process both with zero correlations and with Hellings–Downs correlations. In addition, we formed the spectrum using a factorized-likelihood approach. The




Figure 5. Posterior probability density *violins* for the uncorrelated common-spectrum noise amplitude as a function of time, using a 6 yr (top) and 9 yr (bottom) sliding window (*slices*) over the data set. The spectral index is fixed at $-13/3$. Slices with unconstrained (log) amplitude posteriors have been reweighted to have linear priors in $A_{13/3}^{\text{CRN}}$ (orange), and the 95% confidence upper limits are denoted with vertical arrows. Slices with constrained measurements of $\log_{10} A_{13/3}^{\text{CRN}}$ are colored in blue. For reference, the dashed horizontal line indicates the 1.2×10^{-15} upper limit set by our first 9 yr slice. The solid horizontal line and gray band indicate the measured $\log_{10} A_{13/3}^{\text{CRN}}$ of -14.69 and its 68% credible interval from our full PPTA-DR3 analysis. Here and elsewhere, the violins represent the probability density of a parameter, with broader segments of a violin corresponding to higher probability density.

probability density for the PSD of individual Fourier components can be factorized in the same way as the power-law amplitude. We determine, for each pulsar, the probability density of the (achromatic) noise PSD for each Fourier component. The frequencies for the components are the same for each pulsar, and are the frequencies that the total data set is sensitive to, from $1/T_{\text{span}}$ to $60/T_{\text{span}}$, where T_{span} is the total length of the PPTA-DR3 data set.

The different estimates of the free spectrum can be seen in Figure 6. The results are consistent, showing that the common noise is well described as the probability-weighted mean of all achromatic noise processes in individual pulsars. Additionally, it is consistent with a power law, particularly at the lowest frequencies, which suggests that a single process may be dominating. The factorized-likelihood spectrum shows excess power at a frequency of 14.0 nHz , which corresponds to a period of 2.26 yr . The common free-spectra show reduced sensitivity for periods near 3 yr , which corresponds to the approximate time span of the new pulsars added to the array

(with only UWL data). Including Hellings–Downs correlations does not significantly change the nature of the spectrum. We interpret this as a result of the spectral characteristics of the common noise being dominated by the autocorrelations (Spiewak et al. 2022).

3.6. Isotropic Stochastic Gravitational-wave Background

We searched for evidence of the correlations of an isotropic GWB by considering models that include the Hellings–Downs spatial correlations in the PTA likelihood. We first searched for a power-law process with an unknown spectral index and the DE440 SSE, recovering $\gamma = 3.87 \pm 0.47$ and $\log_{10} A^{\text{HD}} = -14.51^{+0.18}_{-0.20}$, which is consistent with, but less precise than the measurement using the uncorrelated component alone. The agreement is unsurprising as the sensitivity of the PPTA data set is dominated by the autocorrelations. At a fixed spectral index for the assumed GWB from SMBHBs, we measure $\log_{10} A^{\text{GWB}} = -14.68 \pm 0.06$, which is also consistent with the uncorrelated noise. We have also measured the common-noise free spectrum assuming Hellings–Downs correlations and find that it is consistent, at all Fourier frequencies considered, with the uncorrelated common noise (Figure 6, gold). Again this indicates that the cross correlations are dominated by the autocorrelated component.

To quantify whether the data supports the inclusion of these correlations, we estimate the Bayesian odds ratio for using product-space sampling (Arzoumanian et al. 2018, 2020). Relative to a common uncorrelated process, we find no additional Bayesian support, for or against, the Hellings–Downs using this approach, with $\mathcal{B}_{\text{CRN}}^{\text{HD}} \sim 1.5$ for a varied spectral index, and $\mathcal{B}_{\text{CRN}}^{\text{HD}} \sim 2$ for $\gamma = 13/3$. These values are not significant.

3.7. Search for the Spatial Correlations

We directly search for the spatial correlations using our technique described in Section 2.2.2, instead of assuming an ORF and testing for model support through a common-noise model. The measurement of interest here is the correlation coefficient Γ of a $\gamma = 13/3$ process, and we obtain one measurement (the posterior probability density) for each of the 435 unique pulsar pairs in our array. These measurements are informed by the cross correlations only and assume a common distribution for $\log_{10} A_{13/3}$. The parameters in the single-pulsar noise models have also been accounted for in each pair. The numerically marginalized single-pulsar noise properties are self-consistent, meaning that the observed noise characteristics of a pulsar are independent of which pulsar it is paired with.

To visualize the correlations, we place our 435 pulsar-pair PDFs, $p(\Gamma)$, into eight independent and equally spaced bins in a sky separation angle ζ . The PDFs in each bin are multiplied to give the total probability density for that bin. The resulting PDFs are shown in Figure 7, with a gray histogram representing the number of pulsar pairs in each angular bin. It should be noted that because the cross-correlation measurements are not completely independent, the apparent uncertainties will be underestimated (emphasizing the need for statistical validation using bootstrap noise realizations).

We show both the measurements assuming a common distribution of $\log_{10} A_{13/3} = -14.69 \pm 0.05$, and assuming a fixed amplitude of $\log_{10} A_{13/3} = -14.69$. The measurements




Figure 6. Free-spectrum common-noise inference for the PPTA-DR3. The *violins* show the probability density of the power spectral density for the free spectrum of a common-noise process (the width of the violin represents the probability density, with a linear scale). Solid green violins are from a factorized-likelihood analysis using the achromatic noise in single pulsars. Black violins are from a common free-spectrum inference assuming no cross correlations. Gold violins show the free spectrum of a common-noise process with Hellings–Downs cross correlations.




Figure 7. Measured spatial correlations as a function of the angular separation angle, ζ . The width of each *violin* as a function of Γ is proportional to the inferred probability density, $p(\Gamma)$. Binned correlations are shown assuming a common-noise distribution of $\log_{10} A_{13/3}^{\text{CRN}} = -14.69 \pm 0.05$ (green, filled) and with amplitude fixed at the median (black, hollow). The dashed black line is the Hellings–Downs ORF. The gray histogram shows the number of pulsar pairs in each angular bin.

were produced with independent posterior samples and the resulting correlations are nearly identical, which indicates that the effect of KDE uncertainty due to a finite number of samples is negligible. It can be clearly seen that the Hellings–Downs curve (solid black line, Equation (5)) is a better description of the resulting PDFs than zero correlations (which we quantify below). It can also be seen that a monopole (with correlation coefficient, $\Gamma = 1$) and dipolar correlations (with $\Gamma = \cos(\theta)$, i.e., $\Gamma = 1$, at $\zeta = 0^\circ$, and $\Gamma = -1$, at $\zeta = 180^\circ$) are not the dominate source of the common noise at this assumed amplitude

and spectral index. However, these alternative correlations must be present at some level in all PTA data sets.

We compute the model likelihoods for the correlations associated with the common noise and find a likelihood ratio $\log_{10} \Delta \mathcal{L}_{\text{CRN}}^{\text{HD}} = 1.1$ in support of Hellings–Downs over zero correlations. To interpret the significance of this result, we calibrate the false-alarm probability by generating noise from the $p(\Gamma)$ distributions themselves via sky scrambling. We generate 10^4 random pulsar sky distributions with the criteria for near-independence described in Section 2.2.2. As the sky




Figure 8. False-alarm probability calculations from 10^4 quasi-independent randomized pulsar sky distributions (sky scrambles). As with Figure 7, the green (filled) histogram corresponds to the measurements assuming $\log_{10} A_{13/3}^{\text{CRN}} = -14.69 \pm 0.05$, and the black (hollow) histogram assumes the median amplitude. The measured likelihood ratios $\Delta \mathcal{L}_{\text{CRN}}^{\text{HD}}$ from the data under the two assumptions about the amplitude are marked by dashed lines and correspond to a one-sided false-alarm probability of $p \lesssim 0.014$.

distributions are not truly independent, the false-alarm probability is only an estimate (for discussion see Di Marco et al. 2023). A histogram of the $\log_{10} \Delta \mathcal{L}_{\text{CRN}}^{\text{HD}}$ for Hellings–Downs over uncorrelated noise from these randomized skies is shown in Figure 8, from which we derive one-sided p -values of $p \lesssim 0.014$ and $p \lesssim 0.015$ for our observations from the varied and fixed-amplitude correlations, respectively. Using Silverman’s rule (Silverman 1986) to set the kernel bandwidth for each pair, we observe consistent values of $p \lesssim 0.018$ and $p \lesssim 0.012$, respectively. We note that, while sky scrambling with this method is simple and efficient, other false-alarm probability calculations will be required for validating any future GWB detection claims (e.g., phase shifting and accounting for the different sensitivities of the pulsars; Taylor et al. 2017).

The support for Hellings–Downs correlations can be efficiently quantified on a per-pulsar basis with our scheme. Figure 9 shows, for each pulsar, the likelihood ratio (Bayes factor estimator) $\Delta \mathcal{L}_{\text{CRN}}^{\text{HD}}$ computed using only pairs involving the given pulsar. The pulsars are ordered by their support for the uncorrelated common noise in Figure 2. The pairs involving PSR J1744–1134 give the most discrepant $\Delta \mathcal{L}_{\text{CRN}}^{\text{HD}}$ between the two amplitude assumptions we consider. This pulsar has fewer independent posterior samples for the varied amplitude analysis than most others, because of its low likelihood support in this region of the parameter space (PSR J1713+0747, with lower likelihood support, was analyzed including samples from additional MCMC chains to compensate). Therefore the difference likely reflects some small KDE uncertainty due to a finite number of samples.

We compute the number of effective pulsar pairs using σ_k from each pulsar pair k using Equation (9), and find $n_{\text{eff}} = 355$.

3.8. Marginalizing Over Solar System Ephemeris Errors

An accurate model for the position of the center of the solar system is paramount for the timing precision required by PTA experiments. Systematic errors in the positions or masses of the

planets can add additional signals to pulsar timing data sets that could manifest as both spatial correlations, but also as a common-noise process. The most recent SSE from JPL, DE440, used data from the Juno mission and updated very-long-baseline interferometry measurements of the outer planets, and is described as having reduced systematic uncertainties (Park et al. 2021).

We investigated if the amplitude of the apparent CRN could be affected by systematic uncertainties by perturbing the masses of outer planets and the orbital elements of the Jovian system with BAYESEPEHEM models (Vallisneri et al. 2020), with respect to the ephemerides DE421, DE438, and DE440. We find that in all ephemerides considered, our data recovers a nonzero (95% confidence) perturbation to at least one orbital element of the Jovian system. We note that there was no Bayesian evidence for perturbations in the DE438 ephemeris using PPTA-DR2 (Goncharov et al. 2021b).

If we marginalize over these potential perturbations, we recover a steeper spectral index for the common-noise process, $\gamma = 4.02^{+0.81}_{-0.41}$, with a log amplitude $\log_{10} A^{\text{CRN}} = -14.56^{+0.16}_{-0.39}$ with DE440. All ephemerides considered give consistent constraints for this process, as shown by the posterior densities in Figure 10. While we observe weak evidence for SSE errors in the form of nonzero perturbations, it is assumed that if such errors were a significant source of common noise, they would manifest as a dipole-like feature in the inter-pulsar correlations (Tiburzi et al. 2016). Our data shows that the measured common-noise process is not dominated by dipolar correlations, and we have therefore assumed that the DE440 SSE is sufficiently accurate for the purpose of our GWB search.

4. Discussion

The properties of the common noise are consistent, within reported uncertainties, with measurements from recent works by the PPTA and the constituent members of the IPTA. A previous analysis by the PPTA (Goncharov et al. 2021b) found a common uncorrelated process with $\gamma = 4.11^{+0.52}_{-0.41}$ and $\log_{10} A^{\text{CRN}} = -14.55^{+0.10}_{-0.23}$. The EPTA (Chen et al. 2021) have reported a similarly consistent result of $\gamma = 3.78^{+0.69}_{-0.59}$ and $\log_{10} A^{\text{CRN}} = -14.29^{+0.25}_{-0.33}$. NANOGrav (Arzoumanian et al. 2020) identified a common-spectrum process with a steeper median spectral index, remaining consistent through relatively larger uncertainties reported with the value $\gamma = 5.52^{+1.26}_{-1.76}$ and $\log_{10} A^{\text{CRN}} = -14.71^{+0.14}_{-0.15}$. This work used only the five lowest-frequency components, and a shallower process was recovered when more components were considered. This may be due to high-fluctuation-frequency processes present in the data set, such as unmodeled pulse-shape variations (Shannon et al. 2016; Arzoumanian et al. 2020; Goncharov et al. 2021a).

The search for a common-spectrum process through the combination of a previous generation of these data sets by the IPTA (Antoniadis et al. 2022) also finds a consistent signal most closely resembling that identified by the EPTA, $\gamma = 3.90^{+0.90}_{-0.90}$ and $\log_{10} A^{\text{CRN}} = -14.29^{+0.14}_{-0.36}$. The PPTA used the DE436 SSE for the previous analysis, while other PTAs used DE438. The spectral index marginally steepens when we marginalize over the SSE errors modeled by Bayesephem, and becomes more consistent with the prediction for the background predicted by SMBHBs.

This common spectral noise may be caused by the GWB, or it may represent one or more other signals masquerading as a




Figure 9. Likelihood ratio for Hellings–Downs spatial correlations over zero correlations for a $\gamma = 13/3$ process, at an amplitude of $\log_{10} A_{13/3} = -14.69$. Each point is calculated using the 29 pulsar pairs involving the labeled pulsar, and therefore, the points are not independent.

background (Tiburzi et al. 2016; Zic et al. 2022). If the purported common-spectrum noise is a genuine signature of an isotropic, stochastic GWB, then we can expect the IPTA to make a detection using the current-generation data sets (Siemens et al. 2013). In this scenario we need to convincingly explain (1) why three of our pulsars offer negative likelihood support for the existence of the GWB, (2) why the current detected amplitude is higher than previous upper bounds, and (3) why the apparent amplitude of the GWB in our data is measured to be increasing with time.

There are several reasons why individual pulsars may not exhibit the signature of the GWB at the same amplitude. This includes interaction between the GWB signal and intrinsic pulsar timing noise (although the probability of this decreases as data spans increase), misspecification of the intrinsic noise parameters in the modeling, errors in the pulsar timing model, or natural variance in the SMBHB source distribution and anisotropy.

In the PPTA-DR3 data set, the pulsars most discrepant with the inferred common noise are PSRs J1744–1134, J1603–7202, and J1713+0747. Two of these (PSRs J1744–1134 and J1713+0747) are observed by all of the constituent member PTAs of the IPTA, and hence it is key to determine whether they show positive or negative support for the GWB in other data sets, and to explain any discrepancies. Previous analyses have revealed similar effects, including PSR J1713+0747, which exhibited low dropout factors in the previous NANOGrav GWB search analysis (Arzoumanian et al. 2020). This may be attributed in part to imperfect modeling of the known timing events observed in this pulsar (Lam et al. 2018). PSR J1603–7202 is observed as part of the MeerKAT PTA (Miles et al. 2023), which has shorter, but significantly more sensitive measurements of this pulsar at greater cadence. The MeerKAT PTA has fortnightly observations and should be sensitive to the noise processes at higher fluctuation frequencies.

The time-dependence of the signal amplitude is puzzling under an isotropic GWB model. Limits placed on the amplitude of the signal in earlier subsets of the data are inconsistent with high probability ($\log_{10} A_{13/3} < -14.69$ with 99.8% confidence)

with the amplitudes of the signal measured in later data, and the data set in the entirety. As shown in Figure 5, the amplitude obtained with the entire data set is consistent with that from the most recent 9 yr data span. This result is heavily influenced by the most dominant pulsar, PSR J1909–3744. The noise characteristics of this pulsar are unusual, and reflect the recovered properties of the common-spectrum noise. When considering the entire data set, there is only modest evidence for a deviation in the common-spectrum process from a pure power law.

If our results do not represent the signal from a GWB, then we must consider the implications. The three pulsars that offer negative likelihood support for the existence of a $\gamma = 13/3$ process at the measured common amplitude could therefore be representing an approximate upper bound on the amplitude of the background. The amplitude and spectral slope of the GWB is poorly constrained, and so lower amplitudes are still plausible. We reiterate the findings of the previous PPTA analysis (Goncharov et al. 2021b) that the characterization of the common noise can be affected by misspecification, and it is not necessarily a noise floor. That is, the sensitive pulsars with zero red noise can be outweighed in the likelihood by a population of pulsars with similar (but not necessarily identical) noise characteristics (Zic et al. 2022). As it is currently defined, the measured common noise is consistent with the probability-weighted mean of the achromatic noise in our sample of pulsars, across all frequency components.

In the PPTA-DR3 data set, the number of pulsar pairs is not highly sensitive to the correlated component of the GWB at the amplitude of the purported CRN. Using autocorrelations and cross correlations, we find no Bayesian support, for or against, the presence of a GWB. However, using our technique that hierarchically analyses the correlations of the common noise alone, we find a $\log_{10} \Delta \mathcal{L}_{\text{CRN}}^{\text{HD}} = 1.1$ in support of the Hellings–Downs over zero correlations. We computed the significance of this likelihood ratio via sky scrambling, which provides a reliable false-alarm probability through noise generated with the data itself. It is unclear why the latter method provides greater support than the former. It could be that the former method




Figure 10. As with Figure 3, but with BAYESEPHMEM used to marginalize over potential solar system ephemeris errors. We measure $\gamma = 4.02^{+0.81}_{-0.41}$, and $\log_{10} A^{\text{CRN}} = -14.56^{+0.16}_{-0.39}$ for DE440 (green).

requires the autocorrelated and cross-correlated noise to have the same amplitude. If the pulsar noise is misspecified, a mismatch in the amplitude of these two terms could degrade the sensitivity of the statistic. From this result, we conclude that future GWB searches should use (Bayesian) statistical methods that search for the GWB in the cross correlations alone (e.g., Arzoumanian et al. 2020). Confirmation of consistency of the amplitude in the cross correlations and autocorrelations will be necessary to build confidence in using the entire PTA signal for astrophysical interpretation.

5. Conclusions

We have presented a search for an isotropic stochastic gravitational-wave background using the third data release of the PPTA. We have measured the characteristics of a common-spectrum process in the array. Using the pulsar autocorrelations only, we measure $\log_{10} A^{\text{CRN}} = -14.5$ and $\gamma = 3.87$, when using the DE440 SSE. By marginalizing over potential errors in the SSE, we find that the inferred process steepens with $\log_{10} A^{\text{CRN}} = -14.56$ and $\gamma = 4.02$, although there is only weak evidence for these errors. At a fixed spectral index of $\gamma = 13/3$ we find $\log_{10} A_{13/3}^{\text{CRN}} = -14.69 \pm 0.05$, which translates to a constraint on the linear strain amplitude of $A = 2.04^{+0.25}_{-0.22} \times 10^{-15}$ for the model of an isotropic GWB. By analyzing only the spatial correlations using a new technique, we show that the process at this amplitude is consistent with an isotropic GWB at a level that may occur by chance with a probability of $p \lesssim 0.014$. This false-alarm

probability corresponds to a one-tailed test significance of approximately 2σ .

However, some of the apparent characteristics of this common-spectrum process are surprising under the model of an isotropic GWB, including the following:

1. Three pulsars disfavor the presence of a $\gamma = 13/3$ process at the measured *common* amplitude. The measured common-spectrum process, as it is defined, is not necessarily a *noise floor* as would be expected of a GWB. We have demonstrated that the spectral characteristics of the common noise are nearly identical to the probability-weighted mean of all achromatic pulsar noise, which may have significant contributions from intrinsic spin noise. However, these same pulsars appear to contribute positively to the Hellings–Downs correlations when the measured common-noise process is introduced, which may suggest misspecification in the noise models of the pulsars.
2. The apparent amplitude of the common noise changes as a function of time. This observation cannot be explained simply by the growing sensitivity of the data. For approximately the first half of our data set, we appear to be in a nondetection regime, where only upper limits can be placed on the amplitude of a common-spectrum process at $\gamma = 13/3$. The 95% confidence upper limits derived from these early parts of our data set are in tension with the current measurements. For example, in the first 9 yr slice, we find $\log_{10} A_{13/3}^{\text{CRN}} < -14.92$ at 95% confidence

and $\log_{10} A_{13/3}^{\text{CRN}} = -14.69$ is ruled out at 99.8% confidence.

Similar results have been observed by NANOGrav and the PPTA during previous GWB searches. However, the significance of this tension has been debated, owing to differences in the data sets (e.g., differing numbers of pulsars) and analysis procedures (e.g., choice of SSEs, priors, and single-pulsar noise models). For this analysis, we used identical models, priors, and inference techniques on all subsets of the data. We cannot determine the origin of this apparent nonstationarity as the recovered common-noise amplitude (measured as a function of time) is not necessarily determined by a distinct physical process.

3. There is no Bayesian support, for or against, the Hellings–Downs spatial correlations using both the autocorrelations and cross correlations. Yet by analyzing only the cross correlations, using our computationally efficient method, we observe some model support $\log_{10} \Delta \mathcal{L}_{\text{CRN}}^{\text{HD}} = 1.1$ for the Hellings–Downs ORF, from which we derive the significance p -value above. If the signal is genuine, this result may indicate excess noise affecting the autocorrelations, or other weaknesses of the model.

Some of these findings may be resolved by invoking anisotropy in the SMBHB source distribution, as simulations suggest a large variance is possible in the spectrum of a GWB (Rosado et al. 2015). The recovered amplitude of the common noise is close to the maximum expected one of a GWB from a population of SMBHBs (Zhu et al. 2019). The eccentricity of a nearby source or an unresolved local population may result in the apparent nonstationarity of the common noise. An analysis with a larger number of pulsars, for example under the IPTA, should help answer these questions. It will also be interesting to investigate if other PTA observations of our *quiet* pulsars (e.g., J1713+0747, J1744–1134) yield similar results under consistent noise models, and whether the temporal properties of the common noise and its upper limits are observed with independent sets of pulsars. Most importantly, it will be interesting to see if the global IPTA combination of the constituent PTA data sets strengthens the significance of any spatial correlations, resulting in an unambiguous detection.

Acknowledgments

Murriyang, the Parkes 64 m radio telescope is part of the Australia Telescope National Facility (<https://ror.org/05qajvd42>), which is funded by the Australian Government for operation as a National Facility managed by CSIRO. We acknowledge the Wiradjuri People as the Traditional Owners of the Observatory site. We acknowledge the Wurundjeri People of the Kulin Nation and the Wallumedegal People of the Darug Nation as the Traditional Owners of the land where this work was primarily carried out. We thank the Parkes Observatory staff for their support of the project for nearly two decades. We thank the International Pulsar Timing Array Detection Committee for their contributions to the IPTA GW search papers. We also thank CSIRO Information Management and Technology High Performance Computing group for access and support with the petrichor cluster. We acknowledge the use of the Python packages NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), CHAINCONSUMER (Hinton 2016), CORNER (Foreman-Mackey 2016), and KDEPY (Odland 2018) for parts of this work. Part of this research was undertaken as

part of the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) under grant No. CE170100004. R.M.S. acknowledges support through ARC Future Fellowship FT190100155. Work at Naval Research Laboratory (NRL) is supported by NASA. S.D. is the recipient of an Australian Research Council Discovery Early Career Award (DE210101738) funded by the Australian Government. Y.L. acknowledges support of the Simons Investigator grant No. 827103.

ORCID iDs

Daniel J. Reardon  <https://orcid.org/0000-0002-2035-4688>
 Andrew Zic  <https://orcid.org/0000-0002-9583-2947>
 Ryan M. Shannon  <https://orcid.org/0000-0002-7285-6348>
 George B. Hobbs  <https://orcid.org/0000-0003-1502-100X>
 Matthew Bailes  <https://orcid.org/0000-0003-3294-3081>
 Valentina Di Marco  <https://orcid.org/0000-0003-3432-0494>
 Agastya Kapur  <https://orcid.org/0009-0001-5071-0962>
 Eric Thrane  <https://orcid.org/0000-0002-4418-3895>
 Jacob Askew  <https://orcid.org/0009-0002-9845-5443>
 N. D. Ramesh Bhat  <https://orcid.org/0000-0002-8383-5059>
 Andrew Cameron  <https://orcid.org/0000-0002-2037-4216>
 Małgorzata Curyło  <https://orcid.org/0000-0002-7031-4828>
 William A. Coles  <https://orcid.org/0000-0002-5714-7471>
 Shi Dai  <https://orcid.org/0000-0002-9618-2499>
 Boris Goncharov  <https://orcid.org/0000-0003-3189-5807>
 Matthew Kerr  <https://orcid.org/0000-0002-0893-4073>
 Atharva Kulkarni  <https://orcid.org/0000-0003-4847-4427>
 Yuri Levin  <https://orcid.org/0000-0002-6987-1299>
 Marcus E. Lower  <https://orcid.org/0000-0001-9208-0009>
 Richard N. Manchester  <https://orcid.org/0000-0001-9445-5732>
 Rami Mandow  <https://orcid.org/0000-0001-5131-522X>
 Matthew T. Miles  <https://orcid.org/0000-0002-5455-3474>
 Rowina S. Nathan  <https://orcid.org/0000-0002-3922-2773>
 Stefan Osłowski  <https://orcid.org/0000-0003-0289-0732>
 Christopher J. Russell  <https://orcid.org/0000-0002-1942-7296>
 Renée Spiewak  <https://orcid.org/0000-0002-6730-3298>
 Songbo Zhang  <https://orcid.org/0000-0001-7049-6468>
 Xing-Jiang Zhu  <https://orcid.org/0000-0001-7049-6468>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *PhRvL*, **116**, 061102
 Antoniadis, J., Arzoumanian, Z., Babak, S., et al. 2022, *MNRAS*, **510**, 4873
 Antoniadis, J., Babak, S., Bak Nielsen, A. S., et al. 2023, *A&A*, in press
 Arzoumanian, Z., Baker, P. T., Blecha, L., et al. 2023, *ApJL*, **951**
 Arzoumanian, Z., Baker, P. T., Blumer, H., et al. 2020, *ApJL*, **905**, L34
 Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018, *ApJ*, **859**, 47
 Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2016, *ApJ*, **821**, 13
 Burke-Spolaor, S., Taylor, S. R., Charisi, M., et al. 2019, *A&ARv*, **27**, 5
 Caprini, C., Durrer, R., & Siemens, X. 2010, *PhRvD*, **82**, 063511
 Carlin, B. P., & Chib, S. 2018, *J. R. Stat. Soc., B: Stat. Methodol.*, **57**, 473
 Chen, S., Caballero, R. N., Guo, Y. J., et al. 2021, *MNRAS*, **508**, 4970
 Corbin, V., & Cornish, N. J. 2010, *arXiv:1008.1782*
 Cordes, J. M., & Shannon, R. M. 2010, *arXiv:1010.3785*
 Cornish, N. J., & Sampson, L. 2016, *PhRvD*, **93**, 104047
 Detweiler, S. 1979, *ApJ*, **234**, 1100
 Di Marco, V., Miles, T. M., Zic, A., et al. 2023, *ApJ*, in press
 Dickey, J. M. 1971, *Ann. Math. Statist.*, **42**, 204
 Ellis, J., & van Haasteren, R. 2019, PTMCMCSampler: Parallel tempering MCMC sampler package written in Python, Astrophysics Source Code Library, ascl:1912.017
 Ellis, J. A. 2013, *CQGrA*, **30**, 224004

- Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2019, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, *Astrophysics Source Code Library*, ascl:1912.015
- Folkner, W. M., Williams, J. G., & Boggs, D. H. 2009, *IPNPR*, 42-178, 1
- Foreman-Mackey, D. 2016, *JOSS*, 1, 24
- Foster, R. S., & Backer, D. C. 1990, *ApJ*, 361, 300
- Gelman, A. B., Carlin, J. B., Stern, H. S., & Rubin, D. B. 1995, *Bayesian Data Analysis* (3rd ed.; Boca Raton, FL: Chapman and Hall/CRC Press)
- Goncharov, B., Reardon, D. J., Shannon, R. M., et al. 2021a, *MNRAS*, 502, 478
- Goncharov, B., Shannon, R. M., Reardon, D. J., et al. 2021b, *ApJL*, 917, L19
- Goncharov, B., Thrane, E., Shannon, R. M., et al. 2022, *ApJL*, 932, L22
- Grishchuk, L. P. 2005, *PhyU*, 48, 1235
- Guo, Y. J., Li, G. Y., Lee, K. J., & Caballero, R. N. 2019, *MNRAS*, 489, 5573
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, 585, 357
- Hazboun, J. S., Simon, J., Siemens, X., & Romano, J. D. 2020a, *ApJL*, 905, L6
- Hazboun, J. S., Simon, J., Taylor, S. R., et al. 2020b, *ApJ*, 890, 108
- Hellings, R. W., & Downs, G. S. 1983, *ApJL*, 265, L39
- Hemberger, D. A., & Stinebring, D. R. 2008, *ApJL*, 674, L37
- Hinton, S. R. 2016, *JOSS*, 1, 00045
- Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, *CQGra*, 27, 084013
- Hobbs, G., Manchester, R. N., Dunning, A., et al. 2020, *PASA*, 37, e012
- Hourihane, S., Meyers, P., Johnson, A., Chatzioannou, K., & Vallisneri, M. 2023, *PhRvD*, 107, 084045
- Jenet, F. A., Lommen, A., Larson, S. L., & Wen, L. 2004, *ApJ*, 606, 799
- Johnson, A. D., Vigeland, S. J., Siemens, X., & Taylor, S. R. 2022, *ApJ*, 932, 105
- Joshi, B. C., Arumugasamy, P., Bagchi, M., et al. 2018, *JApA*, 39, 51
- Keith, M. J., Coles, W., Shannon, R. M., et al. 2013, *MNRAS*, 429, 2161
- Kerr, M., Reardon, D. J., Hobbs, G., et al. 2020, *PASA*, 37, e020
- Lam, M. T., Ellis, J. A., Grillo, G., et al. 2018, *ApJ*, 861, 132
- Lasky, P. D., Mingarelli, C. M. F., Smith, T. L., et al. 2016, *PhRvX*, 6, 011035
- Lee, K. J. 2016, in *ASP Conf. Ser. 502, Frontiers in Radio Astronomy and FAST Early Sciences Symposium 2015*, ed. L. Qian & D. Li (San Francisco, CA: ASP), 19
- Lentati, L., Alexander, P., Hobson, M. P., et al. 2013, *PhRvD*, 87, 104021
- Manchester, R. N., Hobbs, G., Bailes, M., et al. 2013, *PASA*, 30, e017
- Miles, M. T., Shannon, R. M., Bailes, M., et al. 2023, *MNRAS*, 519, 3976
- Mingarelli, C. M. F., Lazio, T. J. W., Sesana, A., et al. 2017, *NatAs*, 1, 886
- Odland, T. 2018, *tommyod/KDEpy: Kernel Density Estimation in Python v0.9.10*, Zenodo, doi:10.5281/zenodo.2392268
- Ölmez, S., Mandic, V., & Siemens, X. 2010, *PhRvD*, 81, 104028
- Park, R. S., Folkner, W. M., Williams, J. G., & Boggs, D. H. 2021, *AJ*, 161, 105
- Phinney, E. S. 2001, arXiv:astro-ph/0108028
- Pol, N. S., Taylor, S. R., Kelley, L. Z., et al. 2021, *ApJL*, 911, L34
- Rajagopal, M., & Romani, R. W. 1995, *ApJ*, 446, 543
- Ravi, V., Wyithe, J. S. B., Shannon, R. M., & Hobbs, G. 2015, *MNRAS*, 447, 2772
- Reardon, D. J., Zic, A., Shannon, R. M., et al. 2023, *ApJL*, 951, L7
- Romano, J. D., Hazboun, J. S., Siemens, X., & Archibald, A. M. 2021, *PhRvD*, 103, 063027
- Rosado, P. A., Sesana, A., & Gair, J. 2015, *MNRAS*, 451, 2417
- Sazhin, M. V. 1978, *SvA*, 22, 36
- Schuster, E. F. 1985, *Commun. Stat. Theory Methods*, 14, 1123
- Sesana, A. 2013, *CQGra*, 30, 244009
- Shannon, R. M., & Cordes, J. M. 2010, *ApJ*, 725, 1607
- Shannon, R. M., Lentati, L. T., Kerr, M., et al. 2016, *ApJL*, 828, L1
- Shannon, R. M., Ravi, V., Lentati, L. T., et al. 2015, *Sci*, 349, 1522
- Siemens, X., Ellis, J., Jenet, F., & Romano, J. D. 2013, *CQGra*, 30, 224015
- Siemens, X., Mandic, V., & Creighton, J. 2007, *PhRvL*, 98, 111101
- Silverman, B. W. 1986, *Density Estimation for Statistics and Data Analysis* (London: Chapman and Hall)
- Spiewak, R., Bailes, M., Miles, M. T., et al. 2022, *PASA*, 39, e027
- Taylor, S. R. 2021, arXiv:2105.13270
- Taylor, S. R., & Gair, J. R. 2013, *PhRvD*, 88, 084001
- Taylor, S. R., Lentati, L., Babak, S., et al. 2017, *PhRvD*, 95, 042002
- Taylor, S. R., Simon, J., Schult, L., Pol, N., & Lamb, W. G. 2022, *PhRvD*, 105, 084049
- Taylor, S. R., Vallisneri, M., Ellis, J. A., et al. 2016, *ApJL*, 819, L6
- Thrane, E., & Talbot, C. 2019, *PASA*, 36, e010
- Tiburzi, C., Hobbs, G., Kerr, M., et al. 2016, *MNRAS*, 455, 4339
- The LIGO-Virgo-KAGRA Collaboration, The LIGO Scientific Collaboration, the Virgo Collaboration, et al. 2021, arXiv:2111.03606
- Vallisneri, M., Taylor, S. R., Simon, J., et al. 2020, *ApJ*, 893, 112
- van Haasteren, R., Levin, Y., McDonald, P., & Lu, T. 2009, *MNRAS*, 395, 1005
- van Haasteren, R., & Vallisneri, M. 2014, *PhRvD*, 90, 104012
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, 17, 261
- Wyithe, J. S. B., & Loeb, A. 2003, *ApJ*, 590, 691
- Xu, H., Chen, S., Guo, Y., et al. 2023, *RAA*, in press
- Xue, X., Bian, L., Shu, J., et al. 2021, *PhRvL*, 127, 251303
- You, X. P., Hobbs, G., Coles, W. A., et al. 2007, *MNRAS*, 378, 493
- Zhu, X.-J., Cui, W., & Thrane, E. 2019, *MNRAS*, 482, 2588
- Zhu, X. J., Wen, L., Hobbs, G., et al. 2015, *MNRAS*, 449, 1650
- Zic, A., Hobbs, G., Shannon, R. M., et al. 2022, *MNRAS*, 516, 410
- Zic, A., Reardon, D. J., Kapur, A., et al. 2023, *PASA*, in press