The Astrophysical Journal, 976:212 (13pp), 2024 December 1 https://doi.org/10.3847/1538-4357/ad8a60

Characterizing the Nanohertz Gravitational Wave Background Using a *t*-process Power Spectral Density

Shashwat C. Sardesai[1], Joseph Simon[2,3], and Sarah J. Vigeland[1]
[1] Center for Gravitation, Cosmology and Astrophysics, University of Wisconsin–Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA
[2] Department of Astrophysical and Planetary Sciences, University of Colorado, Boulder, CO 80309, USA

Received 2024 August 15; revised 2024 October 18; accepted 2024 October 21; published 2024 November 25

Abstract

Pulsar timing arrays (PTAs) are sensitive to low-frequency gravitational waves (GWs), which induce correlated changes in millisecond pulsars' timing residuals. PTA collaborations around the world have recently announced evidence of a nanohertz gravitational wave background (GWB), which may be produced by a population of supermassive black hole binaries (SMBHBs). The GWB is often modeled as following a power-law power spectral density (PSD); however, a GWB produced by a cosmological population of SMBHBs is expected to have a more complex power spectrum due to the discrete nature of the sources. In this paper, we investigate using a *t*-process PSD to model the GWB, which allows us to fit for both the underlying power-law amplitude and spectral index as well as deviations from that power law, which may be produced by individual nearby binaries. We create simulated data sets based on the properties of the NANOGrav 15 yr data set, and we demonstrate that the *t*-process PSD can accurately recover the PSD when deviations from a power law are present. With longer timed data sets and more pulsars, we expect the sensitivity of our PTAs to improve, which will allow us to precisely measure the PSD of the GWB and study the sources producing it.

Unified Astronomy Thesaurus concepts: Gravitation (661); Computational methods (1965); Markov chain Monte Carlo (1889); Gravitational wave astronomy (675)

1. Introduction

Massive objects orbiting each other radiate away energy in the form of gravitational waves (GWs). Supermassive black hole binaries (SMBHBs), which form following massive galaxy mergers, emit GWs with frequencies ranging from 10^{-9} to 10^{-7} Hz when they are at $\sim 10^{-1}$–10^{-2} pc separations (A. Sesana 2013; S. Burke-Spolaor et al. 2019). Pulsar timing arrays (PTAs) detect GWs at nanohertz frequencies by looking for the time delay induced in the pulses of a collection of rapidly spinning millisecond pulsars (MSPs; M. V. Sazhin 1978; S. Detweiler 1979; R. W. Hellings & G. S. Downs 1983). MSPs are recycled neutron stars with extremely low spin-down rates, and their rotational stability makes them ideal for GW detection (R. S. Foster & D.C. Backer 1990).

The superposition of GWs emitted by a cosmological population of SMBHBs is expected to give rise to a stochastic gravitational wave background (GWB; M. Rajagopal & R. W. Romani 1995; A. H. Jaffe & D.C. Backer 2003; J. S. B. Wyithe & A. Loeb 2003; P. A. Rosado et al. 2015). Recently, the North American Nanohertz Observatory for Gravitational Waves (NANOGrav), the European Pulsar Timing Array (EPTA), the Indian Pulsar Timing Array (InPTA), the Parkes Pulsar Timing Array (PPTA), and the Chinese Pulsar Timing Array announced the first evidence for a GWB at nanohertz frequencies (G. Agazie et al. 2023a; EPTA Collaboration et al. 2023; D. J. Reardon et al. 2023; H. Xu et al. 2023). The observed signal is consistent with a GWB produced by SMBHBs (G. Agazie et al. 2023b; EPTA Collaboration et al. 2024), but it is also consistent with a GWB produced by other sources such as inflation, phase transitions, and cosmic strings (e.g., A. Afzal et al. 2023). An analysis by the International Pulsar Timing Array (IPTA) found that the results from NANOGrav, the EPTA, the InPTA, and the PPTA are broadly consistent with one another (G. Agazie et al. 2024a), and work is underway preparing the third IPTA data set that will combine these data with data from the MeerKat Pulsar Timing Array (M. T. Miles et al. 2023), the Canadian Hydrogen Intensity Mapping Experiment (CHIME/Pulsar Collaboration et al. 2021), and the Low-Frequency Array (M. P. van Haarlem et al. 2013). With longer timing baselines and more pulsars, the sensitivity of PTAs will increase, allowing for more precise measurements of the GWB and the study of its astrophysical or cosmological sources (X. Siemens et al. 2013).

For a population of circular SMBHBs evolving only due to GW emission, the GWB power spectral density (PSD) is expected to follow $P(f) \propto f^{-13/3}$ (E. S. Phinney 2001). However, the PSD for an astrophysically realistic GWB is expected to be more complicated. Environmental effects such as stellar hardening and alpha-disk interactions accelerate the orbital evolution at large separations, resulting in different power-law spectral indices at the low-frequency and high-frequency ends of the spectrum (M. Milosavljević & D. Merritt 2003; B. Kocsis & A. Sesana 2011; L. Sampson et al. 2015), while the presence of a significant fraction of SMBHBs with nonnegligible eccentricity will result in a flattening of the GWB spectrum (A. Susobhanan et al. 2020; F. Fastidio et al. 2024). Because of the discrete nature of the sources, we also expect to find deviations from a power-law PSD at higher frequencies where fewer sources are contributing to the GWB. There may also be excess power at individual frequencies caused by a small number of loud,

---

[3] NSF Astronomy and Astrophysics Postdoctoral Fellow.

nearby sources (B. Bécsy et al. 2022; G. Agazie et al. 2024b). The finiteness of the population of SMBHBs can lead to spectral variation in the GWB, as shown in W. G. Lamb & S. R. Taylor (2024).

In this paper, we explore the validity of detecting deviations from a power-law PSD. We create simulations of the NANOGrav 15 yr data sets, consisting of 67 pulsars, and employ a Bayesian analysis to fit multiple PSD models. We introduce a "*t*-process" model, which consists of an underlying power law convolved with multiplicative variables, each described by an inverse gamma distribution prior, at each frequency and compare the findings to other standard PSD models used in PTA analyses.

This paper is organized as follows. In Section 2 we give an overview of PTAs, we discuss the *t*-process model as well as the other models used to describe the GWB PSD, and we describe our analysis methods. In Section 3 we present the results of our analysis, comparing the results of the *t*-process to other PSD models across different simulations containing a pure power-law GWB, a GWB with excessive noise at a single frequency component, and a simulated astrophysical background from SMBHBs. Finally, in Section 4 we summarize our results and look to the future.

2. Model and Methods

2.1. PTA Overview

Pulsar timing residuals are the difference between the observed times of arrival and the expected, deterministic times of arrival. It can also be represented as a combination of noise processes:

$$\delta t = M\epsilon + Fa + n, \tag{1}$$

where $M\epsilon$ are linear perturbations to the timing model, Fa is the the contribution from intrinsic red noise as well as the common process (GWB), and n is the contribution of white noise. The intrinsic red noise can be modeled as a power law with an amplitude and a spectral index:

$$P_{\rm rn}(f) = A_{\rm rn}^2\left(\frac{f}{f_{\rm yr}}\right)^{-\gamma}{\rm yr}^3. \tag{2}$$

Each pulsar is expected to be affected by intrinsic red noise processes (e.g., spin noise (R. M. Shannon & J. M. Cordes 2010; M. T. Lam et al. 2017; M. E. Lower et al. 2020), variations in the interstellar medium (M. L. Jones et al. 2017; M. T. Lam et al. 2017)), white noise processes (e.g., radiometer noise, pulse jitter; J. M. Cordes & G. S. Downs 1985; M. T. Lam et al. 2016, 2019; A. Parthasarathy et al. 2021), and the GW background itself. It has been shown that in some cases that intrinsic red noise can be mistaken for a common process (B. Goncharov et al. 2021; A. Zic et al. 2022; B. Goncharov et al. 2022; R. van Haasteren 2024) and is very dependent on the Bayesian priors chosen. However, we can distinguish between intrinsic red noise and the GWB based on the correlations between pairs of pulsars—the GWB induces interpulsar correlations that follow the Hellings–Downs curve (R. W. Hellings & G. S. Downs 1983), while noise processes will be uncorrelated or induce other types of correlations (C. Tiburzi et al. 2016).

2.2. Common Red Noise Signal

The incoherent superposition of GWs produced by a population of SMBHBs leads to a GW background that can be described in terms of its characteristic strain $h_c(f)$. For a population of circular SMBHBs evolving only due to GW emission, this is predicted to follow a power law (E. S. Phinney 2001),

$$h_c = A_{\rm gw}\left(\frac{f}{f_{\rm yr}}\right)^{-2/3}, \tag{3}$$

where f is the GW frequency and A is the amplitude at a reference frequency of $f_{\rm yr} = 1\ {\rm yr}^{-1} = 31.7$ nHz. It is common to use a reference frequency of $1\ {\rm yr}^{-1}$, but other reference frequencies such as $0.1\ {\rm yr}^{-1}$ may be used.

PTAs detect GWs by measuring the redshift induced in the pulses of a collection of MSPs. The response in a PTA is a time delay induced in the pulse arrival times of all the pulsars. The measured residual PSD is related to the characteristic strain spectrum according to

$$\begin{aligned} P_{\rm pl}(f) &= \frac{h_c^2(f)}{12\pi^2 f^3}df \\ &= \frac{A_{\rm gw}^2}{12\pi^2}\left(\frac{f}{f_{\rm yr}}\right)^{-\gamma_{\rm gw}} df, \end{aligned} \tag{4}$$

where $A_{\rm gw}$ is the amplitude of the characteristic strain from Equation (3), $df = 1/T_{\rm span}$, the total time span of the PTA, and $\gamma_{\rm gw}$ is the spectral index in the PSD. For purely circular SMBHBs, $\gamma_{\rm gw} = 13/3$.

On average, we expect the GW background to produce a residual power spectrum according to Equation (4). However, an astrophysically realistic spectrum will deviate from a power-law PSD due to the finite number of binaries contributing to the spectrum (A. Sesana et al. 2008). There also may be a turnover in the spectrum at low frequencies if interactions with the environment extend into the PTA band (L. Sampson et al. 2015). At high frequencies, the power spectrum of the residuals flattens due to white noise.

We model the red noise processes using a Fourier series with linearly spaced frequencies $f_i = i/T_{\rm span}$, where $T_{\rm span}$ is the total span of the data and i is an integer. For our analyses we use $i = 1, 2,\ldots,30$, which corresponds to a frequency range of 1.98–59.3 nHz. To capture the subtleties at each frequency component, one can make use of a free spectral approach, which treats the power at each frequency as an independent parameter,

$$P_{\rm fs}(f_i) = \frac{h_c^2(f_i)}{12\pi^2 f_i^3}df = \delta t_{\rm delay}^2(f_i). \tag{5}$$

By design, there is no relation between the power at different frequencies: the free spectral is completely model agnostic and purely measures the power at each frequency individually.

The *t*-process model is a modification to the power law where every frequency component has a multiplicative factor

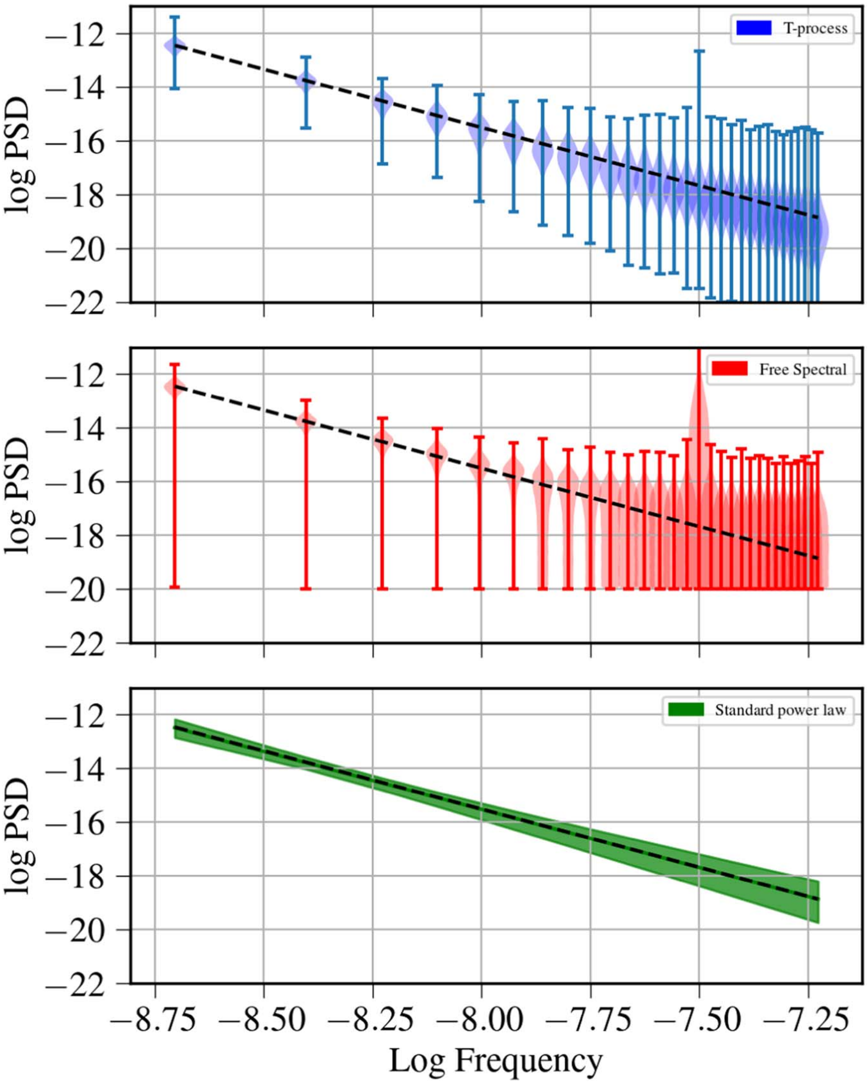

Figure 1. Reconstructed GWB PSD for power-law simulations using a *t*-process model (blue), free spectral model (red), and power-law model (green). The violins and shaded region show the 90% confidence interval. The injected PSD is shown as the black dashed line. The standard power law recovers the injected PSD the best, although all three models recover the injected PSD accurately at low frequencies.

governed by an inverse gamma distribution prior,

$$P_{tp}(f_i) = \alpha_i P_{\mathrm{pl}}(f_i)$$
$$\alpha_i \sim \mathrm{invgam}(1,\ 1). \tag{6}$$

The inverse gamma function follows a probability of

$$invgam(\alpha_{\mathbf{i}},\ 1,\ 1) = \frac{1}{\Gamma(1)} \alpha_{\mathbf{i}}^{-2} \exp\left(-\frac{1}{\alpha_{\mathbf{i}}}\right), \tag{7}$$

where $P_{\mathrm{pl}}(f_i)$ is the standard power-law model from Equation (4) and α_i is a multiplicative factor that adjusts to power at a frequency f_i compared to the underlying power-law PSD. This model is a good compromise between the power-law model, which assumes a relationship between the power at different frequencies, and the agnostic free spectral model since it allows for deviations in the PSD from a power law. This model was previously used to describe the intrinsic red noise of pulsar J0613−0200 in Z. Arzoumanian et al. (2020) since that pulsar showed both low-frequency red noise and excess noise at $f_{\mathrm{gw}} = 15$ nHz. In this work, we use it to model the GWB because it is able to model the underlying power-law PSD of the GWB as well as capture and quantify deviations from the power law.

2.3. Bayesian Methods and Software

We create simulations of the NANOGrav 15 yr data set using the `enterprise` (J. A. Ellis et al. 2019) and `enterprise_extensions`(S. R. Taylor et al. 2021) packages. Our likelihood is modeled as a Gaussian,

$$p(\vec{\delta t}|\phi) = \frac{\exp\left(-\frac{1}{2}\vec{\delta t}^{T}\Sigma^{-1}\vec{\delta t}\right)}{\sqrt{2\pi(\det\Sigma)}}, \tag{8}$$

where $\vec{\delta t}$ is the vector of residuals of all pulsars, ϕ is the vector of parameters drawn from the priors, and Σ is the covariance block matrix of all pulsars. We use `PTMCMCSampler` (J. Ellis & R. van Haasteren 2019) to draw values from the posteriors for intrinsic red noise and common red noise parameters using a Markov Chain Monte Carlo approach. For

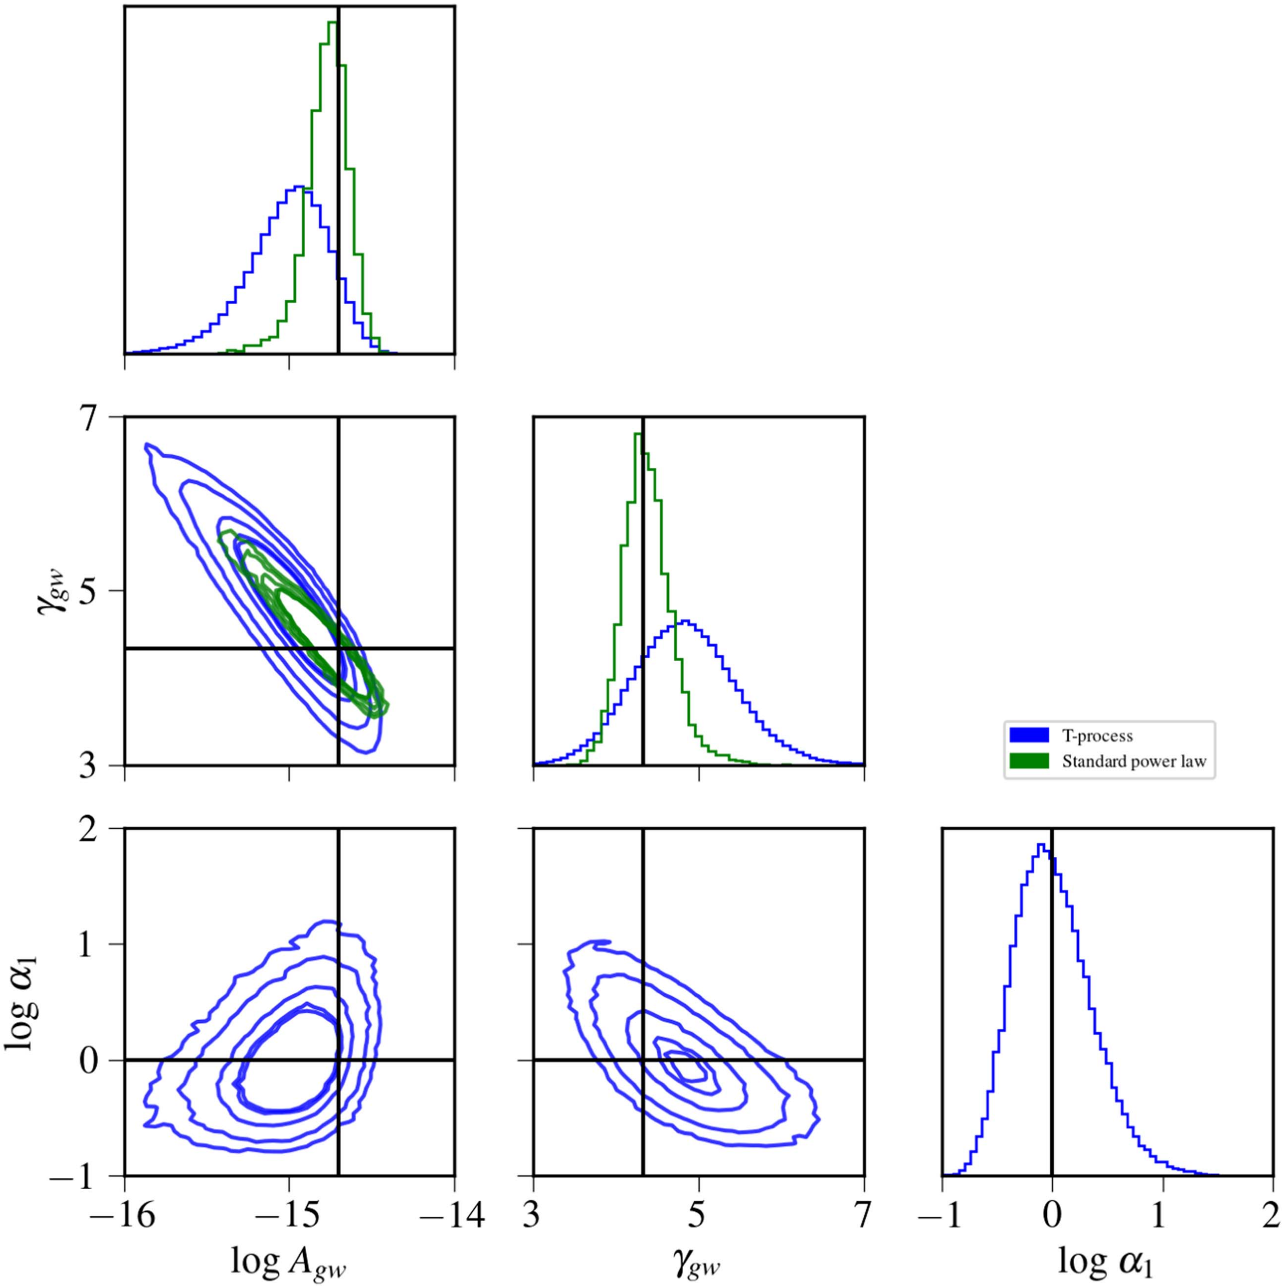

Figure 2. Combined posterior plot for all 50 simulations of $\log A_{\mathrm{gw}}$, γ_{gw}, and $\log \alpha_1$ for the power-law recovery (blue) and *t*-process recovery (green). We see that the standard power law recovers the injected amplitude and spectral index value very accurately. The *t*-process is also able to recover the injected values but underestimates the amplitude and overestimates the spectral index. We also see that $\log \alpha_1$ is centered at 0, implying there is no deviation from a power-law PSD at that frequency. The other α_i values not shown have similar distributions, with increasing variance as we go to higher frequencies.

the purposes of this paper, interstellar variations/dispersion medium, pulse jitter, and correlated white noise are not considered.

2.4. Simulations

Our simulated PTA is based on the NANOGrav 15 yr data set G. Agazie et al. (2023a), which contains 67 pulsars and a baseline of 16.03 yr. In each of our simulations, the injected uncertainty in time-of-arrivals (TOAs) and intrinsic pulsar red noise are equal to the measured quantities in the pulsars in the 15 yr data set (the injected values are given in Table 1). This is done by running a Bayesian analysis with intrinsic red noise included for each pulsar as well as a common red noise process across the entire array. We then only consider the maximum likelihood intrinsic red noise values (this is to prevent contamination of the pulsar-dependent red noise with the common red noise signal present in the data set). We then inject a new common red noise spectra using `enterprise` and analyze the results.

To test the *t*-process, we use three different types of injected common signal. We use (1) a pure power law, (2) a power law with a known amount of excess noise at 5.93 nHz (the third frequency bin), and finally (3) an astrophysically realistic background from a simulated population of SMBHBs created by the `holodeck`[4] (G. Agazie et al. 2023b)

---

[4] https://github.com/nanograv/holodeck

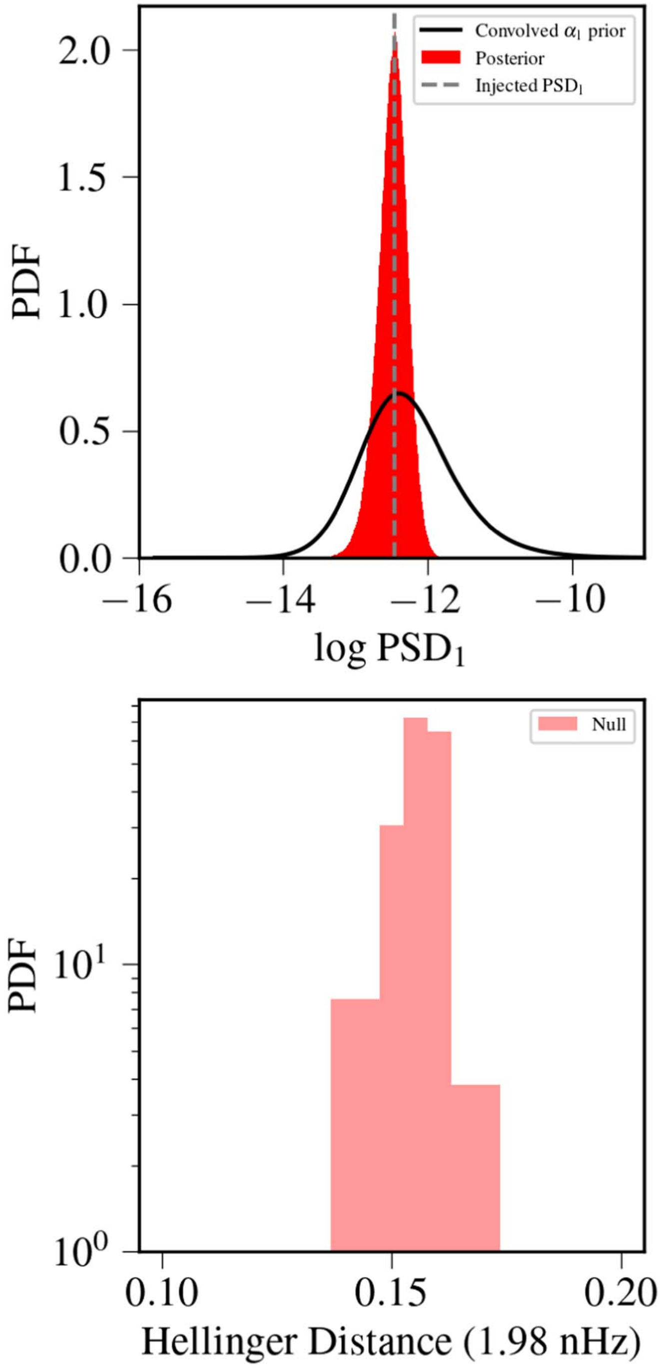

Figure 3. The log-PSD values (top) and the distribution of Hellinger distances for the first frequency (1.98 nHz) of the pure power-law injection (bottom) of all 50 simulations. The black curve indicates the convolution of the log $A_{\rm gw}$, $\gamma_{\rm gw}$ posteriors and the inverse gamma prior. We see that the PSD is accurately recovered at the maximum PDF of the convoluted prior, implying that there is no adjustment of the α_1 parameter. This is the null distribution of Hellinger distances (no deviations from the power law), and at 1.98 nHz the 90% confidence interval is between 0.146 and 0.163.

simulation package. To create a realistic simulation of the GWB, we must take into account the evolution and merging of galaxies that host the SMBHs. The `holodeck` package uses a semianalytical model, which includes hyperparameters (the injected values are listed in Table 2) for galactic merger time, galaxy stellar mass function, galaxy close pair fraction, and SMBH mass—host galaxy relation to get a differential number density of SMBHBs (J. J. Simon 2017). The characteristic strain from a population of SMBHBs is then calculated by E. S. Phinney (2001), A. Sesana et al. (2008), S. T. McWilliams et al. (2012), and L. Sampson et al. (2015):

$$h_c^2(f) = \int dz \int dM \int dq \frac{d^4N}{dM\, dq\, dz\, d\ln f} h_s^2(f). \tag{9}$$

Here, N is the total number of SMBHBs present within each bin of total mass M, mass ratio q, redshift z, and log frequency $\ln f$. Since we are dealing with discrete sources, we must modify our integral to be a sum. The `holodeck` package uses a Poisson distribution to estimate the number of binaries within a range of dz, dM, and dq at each $d\ln f$ to calculate the overall characteristic strain from individual sources as shown in J. J. Simon (2017) and G. Agazie et al. (2023b):

$$h_c^2(f) = \sum_{M,q,z,f} \mathcal{P}\left(\frac{\partial^4 N}{\partial M \partial q \partial z \partial \ln f} \Delta M \Delta q \Delta z \Delta \ln f\right) \frac{h_s^2(f)}{\Delta \ln f}, \tag{10}$$

where $\mathcal{P}$ indicates a random number drawn from a Poisson distribution with the mean value given by the number of binaries in each bin of M, q, z, and $\ln f$. For a large population of sources, h_c^2 closely follows a power law, as described in Equation (3). However, when the signal is dominated by a small number of discrete sources, there is a breakdown of stochasticity and a subsequent deviation from power-law behavior: this can be due to the signal being dominated by a single loud binary or because the binaries are evolving quickly at high GW frequency.

3. Results

For all of our simulations, we use three models to reconstruct the PSD: a power law, a t-process, and a free spectrum. We show reconstructions of the PSDs using all three models and compare to the injected GWB. For the power-law and t-process models, we also show corner plots for the model parameter posteriors. This allows us to compare the recovered parameter posteriors to the injected values, as well as explore the covariances between model parameters.

For the t-process model, the α parameters provide a way to measure deviations from a power-law PSD. If there are no deviations present, the posteriors for α will peak at 1. In order to assess the significance of deviations, we use the Hellinger distance (E. Hellinger 1909) to compare the posteriors and priors. The Hellinger distance is a measure of the similarity between two distributions: a Hellinger distance of 0 means the two distributions are the same and 1 means the two are completely different. Because there is a large covariance between the log $A_{\rm gw}$, $\gamma_{\rm gw}$, and the α_i parameters, instead of comparing the prior and posterior distributions for the α_i, we compare distributions for the power at a given frequency f_i. We construct convolved priors using the posteriors for log $A_{\rm gw}$ and γ and the prior on α_i, then compute the Hellinger distance between these distributions and the posteriors on the power at f_i, according to Equation (6).

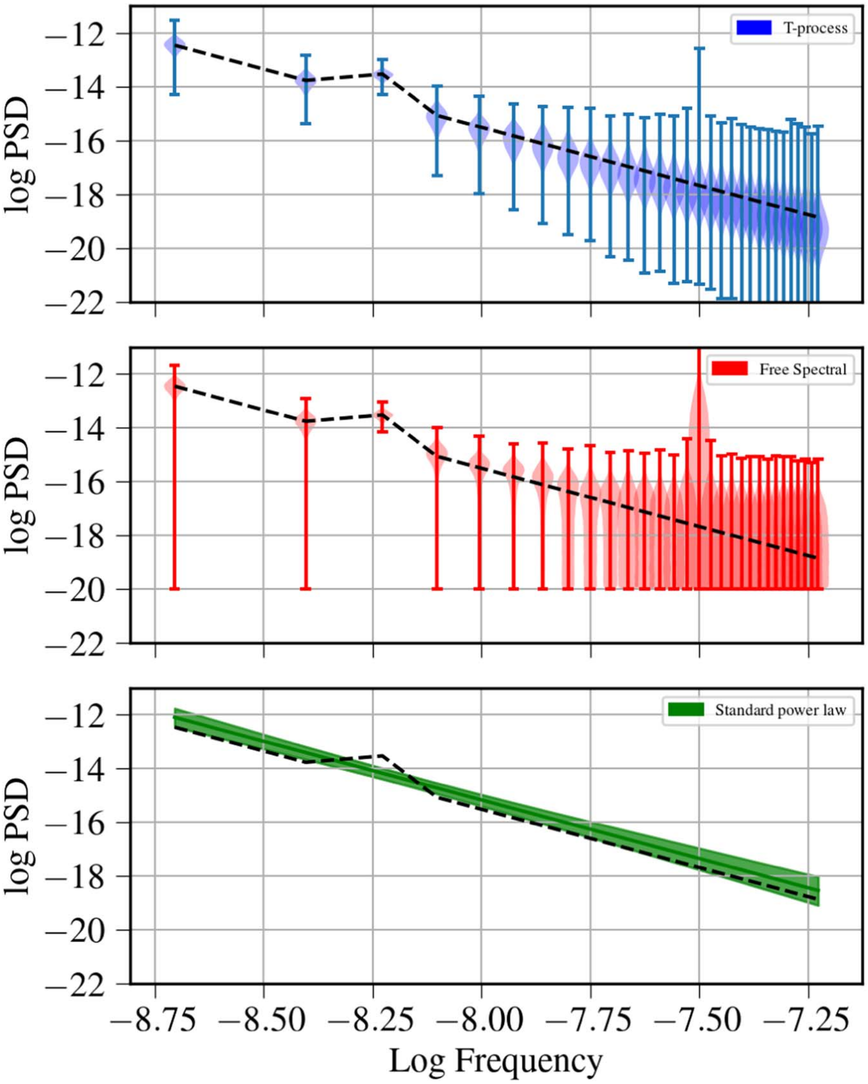

Figure 4. Reconstructed GWB PSD for excess power simulation using a *t*-process model (blue), free spectral model (red), and power-law model (green). The violins and shaded region show the 90% confidence interval. The injected is shown in the black dashed line. For these simulations, the PSD follows a power-law PSD with excess power at 5.93 nHz. The *t*-process and the free spectral are able to recover the bump as well as the rest of the injection, whereas the standard power law overestimates the amplitude of the entire PSD.

3.1. Pure Power-law Background

We generate simulated data sets containing a GWB described by a power-law PSD with an amplitude of $A_{\mathrm{gw}} = 2 \times 10^{-15}$ and a spectral index of $\gamma_{\mathrm{gw}} = 13/3$. This simulation can be thought of as a litmus test to see how the *t*-process model recovers the PSD when there are no deviations present. For an ideal recovery, the *t*-process would find the injected amplitude and spectral index, and we expect the $\log \alpha_i$ posteriors to have maximum PDF values at 0 $\forall f_i$.

In Figure 1 we present the reconstructed PSD at each frequency component for all three models. We find good agreement between the reconstructed PSDs using all three models at low frequencies: at higher frequencies, the reconstructed PSD using the free spectral model has much larger variances compared to the other two because the reconstruction is dominated by pulsar white noise. By design, the free spectral model does not assume any relationship between power at different frequencies, whereas the power-law and *t*-process models do. In Figure 2, we present the posteriors of recovered parameters for the power-law and *t*-process models. We find that the power-law model best recovers the injected values of $\log A_{\mathrm{gw}}$ and γ, while the *t*-process slightly underestimates the amplitude and slightly overestimates the spectral index. We also see that the posteriors of $\log \alpha_1$ are recovered with maximum posterior values at 0, which is consistent with a pure power-law PSD injection. For the other components, we also see a $\log \alpha_i$ distribution that is centered close to 0, and as we go to higher frequencies, the variance of the distribution increases.

In Figure 3 we present the posterior of the log PSD (top) and the distribution of Hellinger distances (bottom) of the first frequency component (1.98 nHz) for all 50 simulations. Across these simulations, the Hellinger distances range from 0.146 to 0.163 (90% confidence interval), indicating that the posteriors are very similar to the priors.

Since this is the pure power-law simulation, the Hellinger distances we calculate can be treated as a “null distribution,”

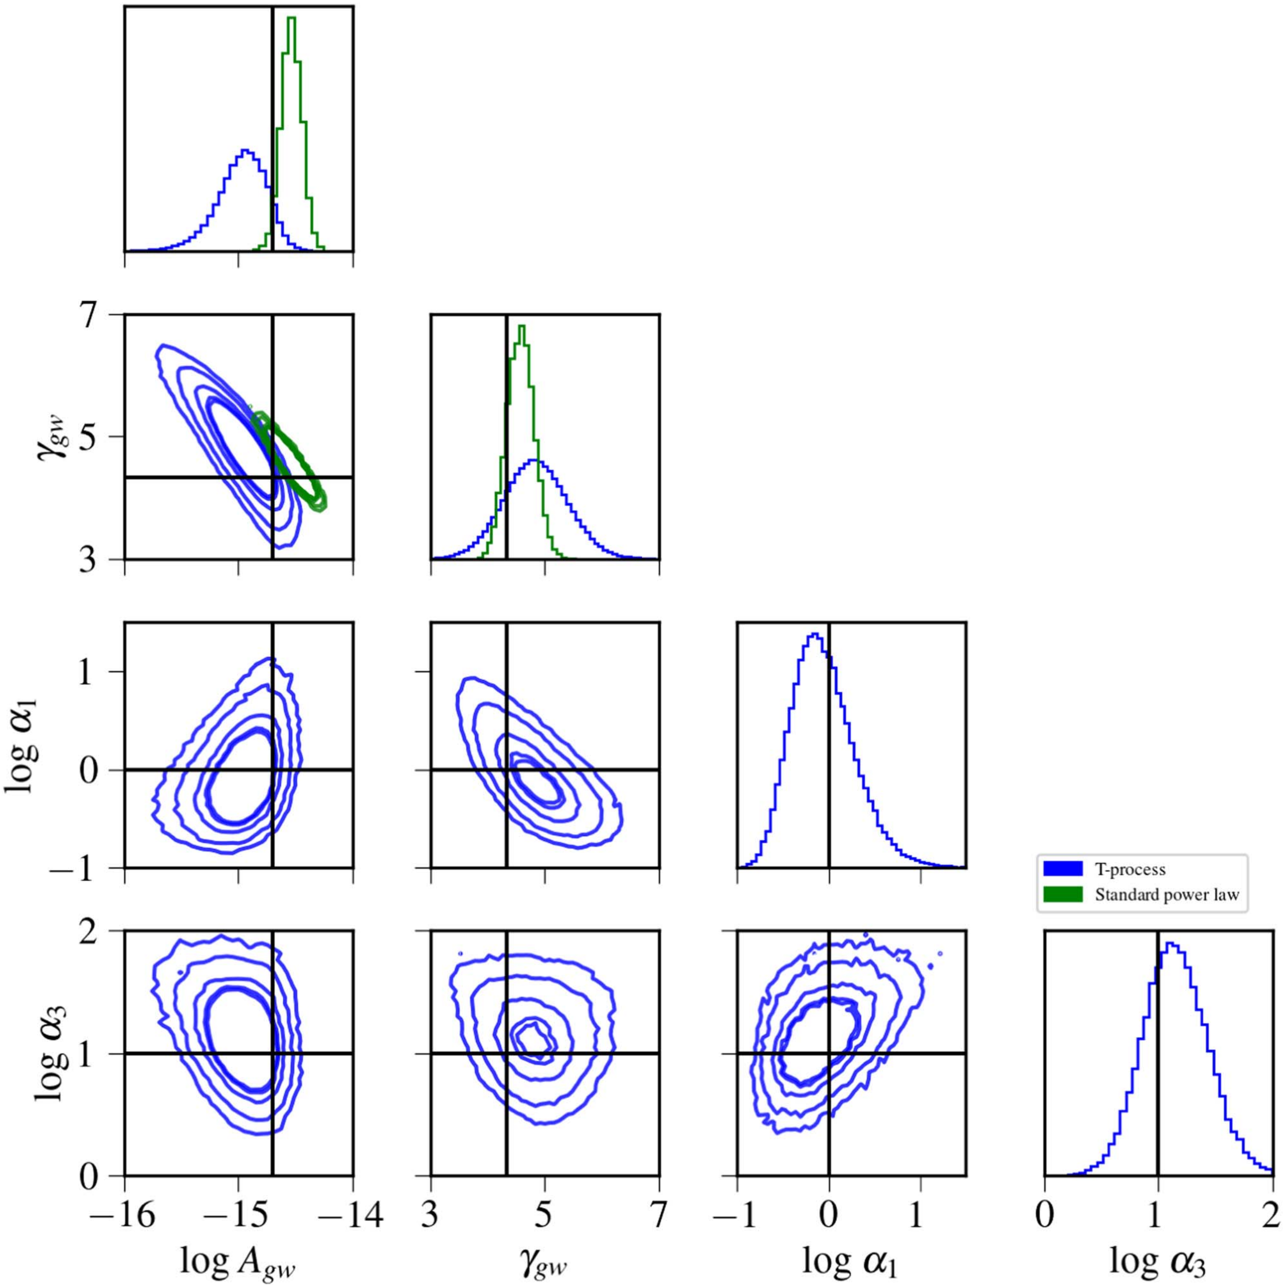

Figure 5. Corner plot of all 50 simulations of the 3rd frequency bump simulations. The blue contours and histograms are for the *t*-process, while the green histograms and contours are for the standard power-law recovery. We can clearly see that the standard power-law model is completely skewed toward higher $\log A$ values, whereas the *t*-process captures the injected values, as well as recovers the deviation at 5.93 nHz. We see that since $\log \alpha_3 = 1.0$, the *t*-process detects the presence of excess power. The other $\log \alpha_i$ distributions not shown here all peak around 0, and as we increase frequency, the variance increases as well.

i.e., the distribution of the Hellinger distances in the absence of deviations from a power law. We use these results as a metric for determining the significance of deviations from a power law when analyzing our other simulations.

Our sensitivity to deviations varies as a function of frequency. The range of Hellinger distances shown in Figure 3 are typical for low frequencies up to ∼13.8 nHz. As we move to higher frequencies, the signal is contaminated by pulsar white noise, resulting in smaller Hellinger distances since the recovered posteriors are more similar to the priors. The lowest Hellinger distance occurs at $16/T_{\rm span} \sim 1{\rm yr}^{-1}$ where we are least sensitive to GWs due to the Earth's orbital period.

3.2. Excess Power at Single Frequency

In order to test how well we can characterize deviations from a power-law PSD, we generate simulations where the GWB follows a power-law PSD with a known value of excess power at a single frequency. Such a feature could be produced by a single nearby circular individual SMBHB, although the detection of such a source would require performing a search using the appropriate deterministic signal model.

A search of the NANOGrav 15 yr data set found no evidence of GWs from individual SMBHBs in G. Agazie et al. (2023c), but with longer baselines and more pulsars, we expect the sensitivity of PTAs to improve to the point where a

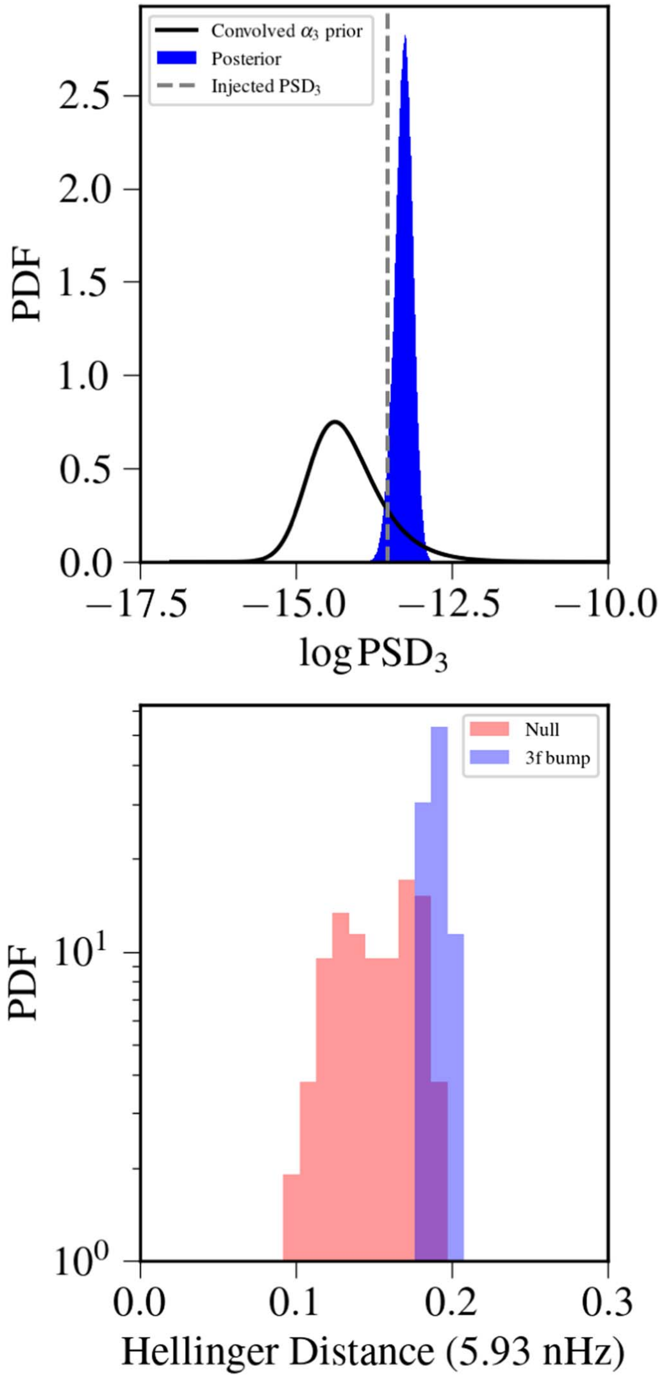

Figure 6. The recovered log-PSD values (top) and the distribution of Hellinger distances (bottom) comparing the 3f bump simulations in blue to the pure power-law simulations (null hypothesis) in red of all 50 simulations. The black curve indicates the convolution of the $\log A_{\mathrm{gw}}$, γ_{gw} posteriors and the inverse gamma prior. The posterior is overestimated from the injected value but still presents a clear distinction from the prior, implying that the α_3 parameter is adjusted for the excess noise at 5.93 nHz. The null distribution has a 90% confidence interval between 0.114 and 0.184, while the excess noise simulations have a 90% confidence interval between 0.182 and 0.201 at 5.93 nHz.

handful of sources will be detectable (P. A. Rosado et al. 2015; C. M. F. Mingarelli et al. 2017; B. Bécsy et al. 2022). Excess correlated noise could also be produced by common sources, such as error in the solar system ephemeris (G. Hobbs et al. 2012; E. Roebber 2019; M. Vallisneri et al. 2020), as well as contributions from solar wind (C. Tiburzi et al. 2021).

We generate the simulated data sets with a common process with an amplitude of 2×10^{-15} and a spectral index of $13/3$, then multiply the power at 5.93 nHz (the third frequency component in our Fourier basis) by a factor of 10. In terms of the t-process model, this corresponds to $\log \alpha_3 = 1$ and $\log \alpha_i = 0$ for $i \neq 3$. This excess noise could be caused by a feature of the background or an unmodeled red noise process. An important note here is the fact that the source of excess noise will likely be at a frequency different than the bin center. PTAs have a finite frequency resolution in Fourier space ($\Delta f = T_{\mathrm{span}}^{-1}$). Any excess power within the frequency bin edges will be recovered at the bin centers our PTA is sensitive to. When excess power is in between two Fourier bins, both bins see an excess of power. In this toy-model case, we have chosen to inject the excess power at the bin center so that the only one α parameter sees the deviation. We consider the case where multiple Fourier bins are affected in Section 3.3.

In Figure 4 we present the reconstructed PSDs for each of the three models for all 50 simulations. We find that the t-process model accurately recovers the injected PSD. The free spectral model accurately recovers the PSD at low frequencies, but at high frequencies the spread in the recovered power is significant due to the presence of pulsar white noise. The power-law model is not able to accurately recover the injected PSD because it cannot recover the deviation at 5.93 nHz.

In Figure 5 we present the posterior distributions for the parameters of the t-process and power-law models. For the t-process model, we find that the amplitude is slightly underestimated while the spectral index is slightly overestimated, as was the case for the power-law simulation (Figure 2). We also find that the posterior for $\log \alpha_3$ has a maximum PDF value at 1, indicating that the magnitude of the deviation is accurately recovered, while the posterior at $\log \alpha_1$, where there is no deviation, has a maximum PDF value at 0. In contrast, the power-law model overestimates both the amplitude and spectral index in an attempt to fit the excess power at 5.98 nHz. This shows that the t-process can recover deviations without significantly biasing the recovery of $\log A_{\mathrm{gw}}$ and γ_{gw}, unlike the recovery with the power-law model. For all other frequency components, $\log \alpha_i$ peaks close to 0 and the variance of each distribution increases as we go to higher frequencies.

In Figure 6 we present the posterior log PSD (top) and the distribution of Hellinger distances (bottom) at the third frequency component (5.93 nHz) for all 50 simulations (blue histogram). We compare these to the distribution of Hellinger distances from the power-law simulations (red histogram). For the excess power simulations, the Hellinger distances range from 0.182 to 0.201 (90% confidence interval), whereas for the power-law simulations, the Hellinger distances range from 0.114 to 0.184 (90% confidence interval). In 82% of the excess power simulations, the Hellinger distance falls outside the 90% confidence interval for the power-law simulations. From this, we conclude that we can use the Hellinger distances to determine the significance of deviations from a power-law PSD. We also show the posterior on the recovered power at 5.93 nHz (blue histogram) to the convolved prior (black histogram). We see the posterior differs significantly

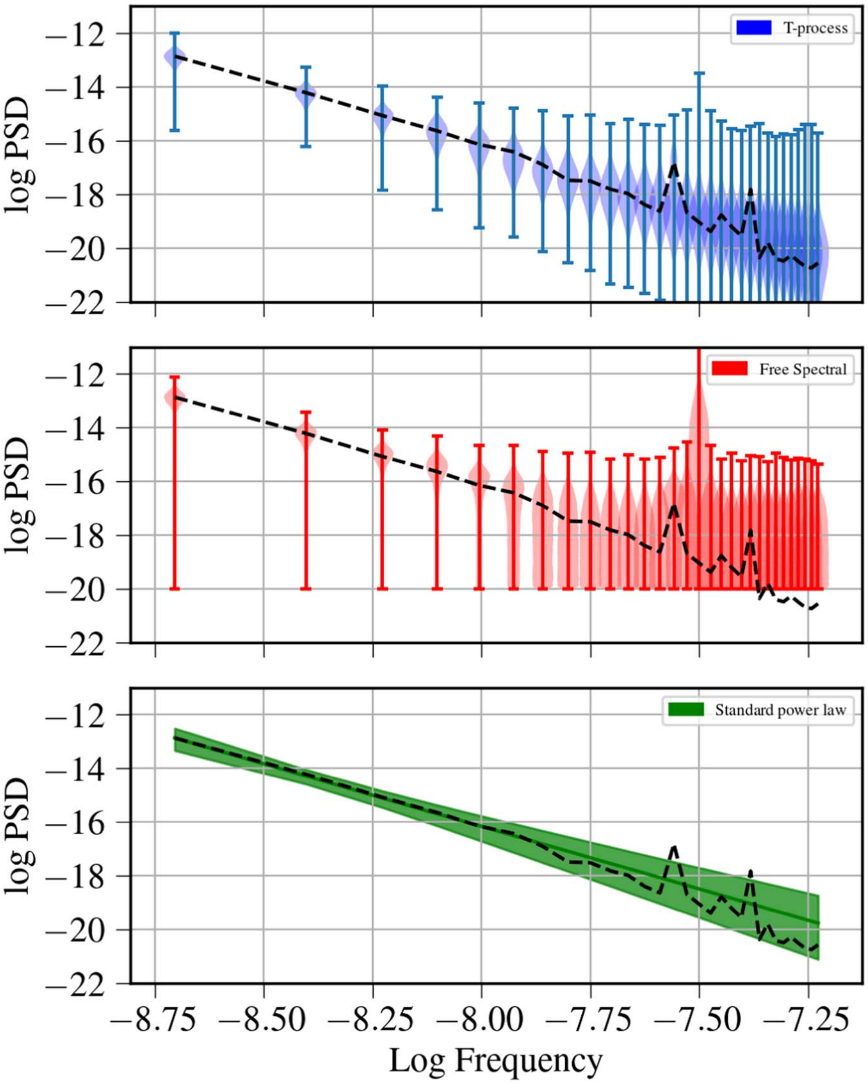

Figure 7. Reconstructed GWB PSD for the `holodeck` simulations using a *t*-process model (blue), free spectral model (red), and power-law model (green). The violins and shaded region show the 90% confidence interval. The injected PSD is shown as the black dashed line. The *t*-process model is best able to recover the injected PSD, while the free spectral model and power-law model recover the injected PSD well at low frequencies but not at high frequencies.

from the prior, and has a maximum posterior value of −12.674, which is consistent with the injected PSD.

3.3. Astrophysically Realistic GWB from SMBHB Populations

In the previous subsections, we have shown how the *t*-process PSD can be used to recover the GWB PSD for two types of simulations: one where the GWB PSD follows a power-law PSD, and one where the PSD follows a power-law PSD with excess power at a single frequency. While these tests show the effectiveness of the *t*-process, it is also important to consider how the *t*-process would analyze realistic PTA data. In this section, we present results from analyzing simulations where we generate the GWB using simulated populations of SMBHBs. We use `holodeck` to generate the GWB PSD, using hyperparameters derived from astrophysical priors as were shown in G. Agazie et al. (2023b). The values used to generate these simulations are listed in Appendix. Unlike the previous set of simulations, the GWB measured here is generated from many individual sources that are at frequencies between the fourier bin edges, this leads to a more realistic case where binaries may lie between two different frequency bins contributing to both of them.

In Figure 7 we present the reconstructed PSDs using the *t*-process, free spectral, and power-law models. We see that the *t*-process accurately reconstructs the PSD across all frequencies, including some of the higher frequencies that show deviations from a power-law PSD. The free spectral model is able to recover the PSD at low frequencies but not at high frequencies where pulsar white noise dominates. The standard power-law is able to recover the PSD at lower frequencies, but fails to capture some of the deviations present at higher frequencies.

In Figure 8 we present the posterior distributions of the parameters for both the *t*-process and power-law models for all 50 simulations. Both the *t*-process and power-law models recover consistent posteriors for $\log A_{\rm gw}$ and $\gamma_{\rm gw}$, although the posteriors from the *t*-process model are significantly broader. For the *t*-process model, we also show the posteriors of $\log \alpha_1$ and $\log \alpha_{14}$. The $\log \alpha_1$ distribution peaks at 0, implying no deviations, while the $\log \alpha_{14}$ posterior shows a tail in the posterior that extends to high values, indicating that there is

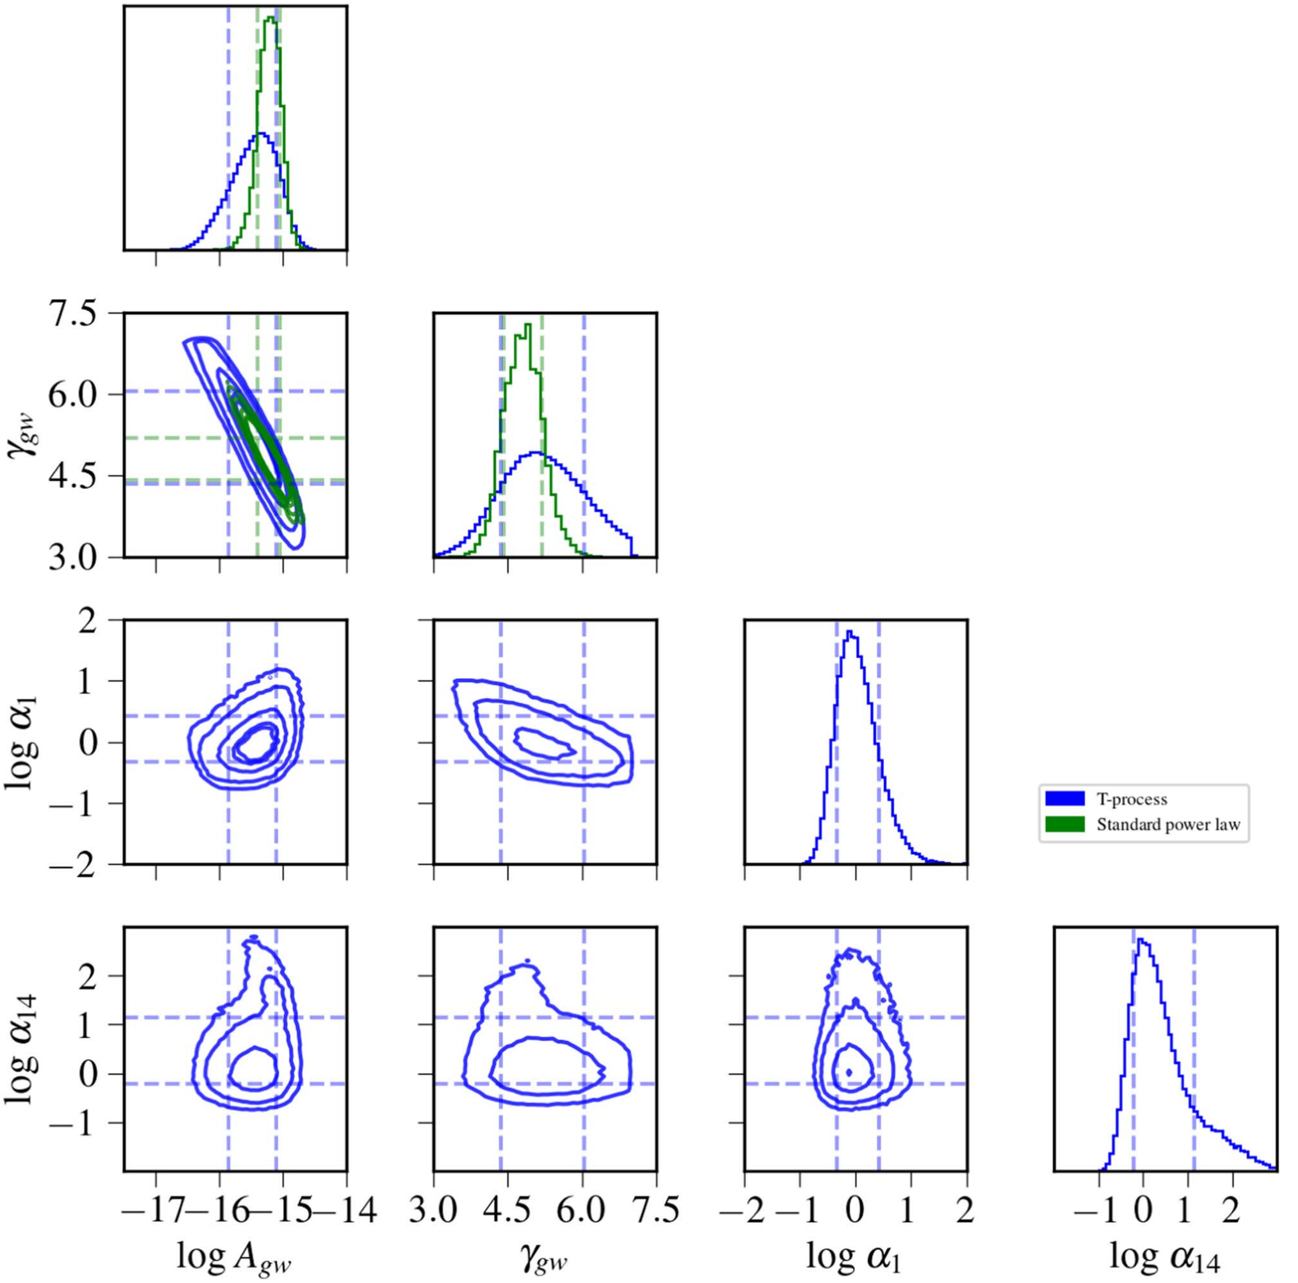

Figure 8. Corner plots of all 50 holodeck simulations. The blue contours and histograms are for the t-process, while the green histograms and contours are for the standard power-law recovery. The green and blue dashed lines indicate the 1σ confidence interval. We see that both models recover the same $\log A$ and gamma values, but the t-process can also detect the presence of a deviation at 27.7 nHz, which was seen in Figure 7, whereas the t-process cannot detect such a deviation. We also see that α_{14} component searches over more of the inverse gamma prior. The $\log \alpha_i$ distributions not shown here peak between 0.13 and 0.3, with increasing variance as we go to higher frequencies.

some evidence of excess power at a frequency of $14/T_{\rm span} = 27.7$ nHz. Across all other $\log \alpha_i$, we observe a peak close between 0.13 and 0.3 and a larger spread as we move to higher frequencies.

In Figure 9 we present the posterior for log PSD (top) and the distribution of Hellinger distances (bottom) at 27.7 nHz. While the posterior (blue histogram) and prior (black histogram) are very similar, we note that there is a second peak in the posterior that extends beyond the prior. This comes from the fact that for some simulations the excess power at this frequency can be recovered, leading to a prior that is peaked at higher values, but for others it cannot. We also show the distribution of Hellinger distances compared to the null hypothesis at 27.7 nHz. For the `holodeck` simulations, the Hellinger distances range from 0.022 to 0.457 (90% confidence interval), while for the power-law PSD simulations, the Hellinger distances range from 0.027 to 0.104 (90% confidence interval). For 24% of the `holodeck` simulations, the Hellinger distance is outside of the 90% confidence interval for the power-law simulations, indicating that the deviation at 27.7 nHz from a power law is significant. The `holodeck` simulations show additional features in the PSD at higher frequencies; however, the distributions of Hellinger distances at these frequencies are very similar to

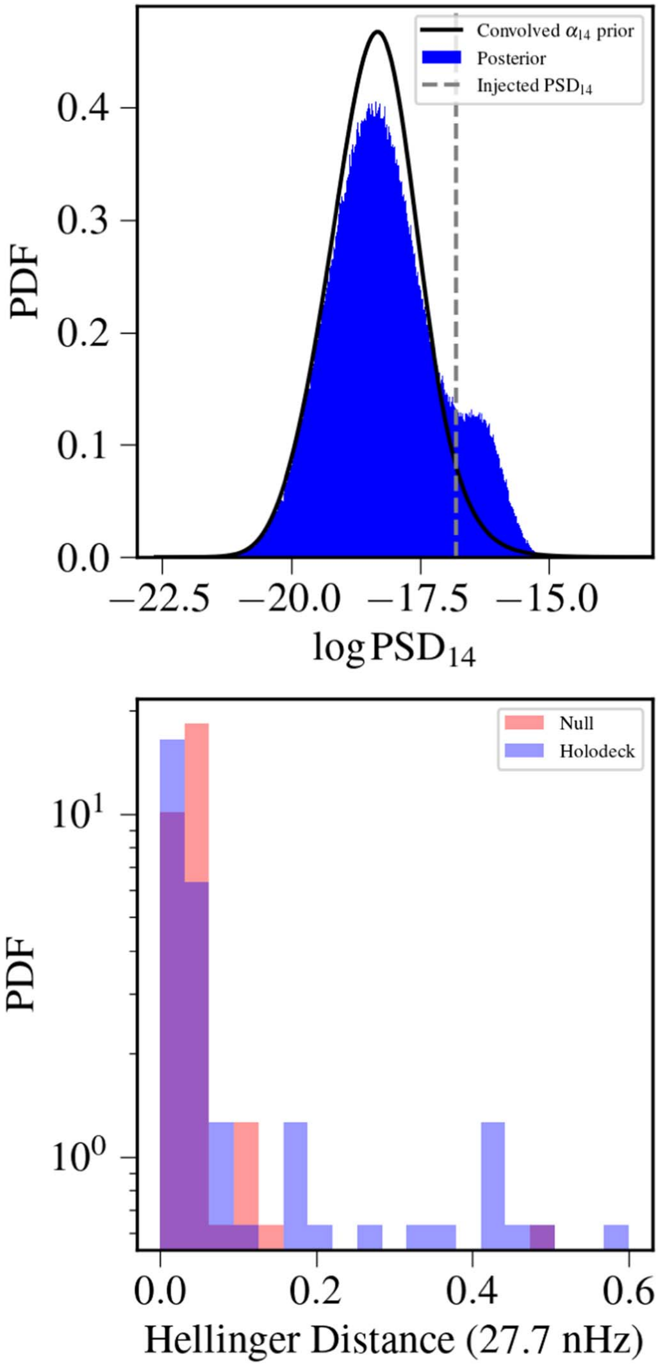

Figure 9. The recovered log PSD (top) and the distribution of Hellinger distances (bottom) comparing all 50 of the holodeck simulations in blue to the pure power-law simulations (null hypothesis) in red for the 14th frequency component (27.7 nHz), where a major deviation from the power law occurs. The astrophysically realistic simulations have a 90% confidence interval between 0.022 and 0.457, whereas the null distribution has a 90% confidence interval between 0.027 and 0.104. The bottom panel shows the log-PSD posteriors (blue), and the black curve indicates the convolution of the $\log A_{\mathrm{gw}}$, γ_{gw} posteriors and the inverse gamma prior. At the this frequency component, the posterior tends to search over the entire prior space, but we see a deviation for some simulations, which are able to detect the presence of excess noise at 27.7 nHz. Simulations with $\mathrm{HD} < 0.1$ are the closest to the prior, whereas simulations with $\mathrm{HD} > 0.1$ display the bimodality of the posterior in the bottom panel.

the distribution of Hellinger distances for the power-law simulations, indicating that these deviations are not significant.

4. Conclusions

PTA experiments around the world have published the first evidence of a nanohertz GWB, which is broadly consistent with that predicted to be produced by a cosmological population of SMBHBs. By measuring the GWB PSD, we can learn about the astrophysical or cosmological sources producing it. In this paper, we used simulated PTA data to study how a *t*-process PSD can recover the GWB. We considered three types of simulations: (1) a pure power-law PSD, (2) a power-law PSD with excess power at a single frequency component (which could be the result of the background or an unmodeled red noise process), and (3) an astrophysically motivated background from a simulated population of SMBHBs. We have shown that the *t*-process is able to accurately recover the PSD for all three simulation types. In addition, the *t*-process model provides a way to determine the significance of any deviations from a simple power-law PSD.

We show that in all three different simulation types, the *t*-process is accurately able to reconstruct the PSD of the GWB. In simulation set (1), it was able to recover the PSD but has some spread in its recovery due to the 30 extra parameters (α_i), which complicate the model. In simulation set (2), the *t*-process was able to recover not only the amplitude and spectral index but the excess of power at the third frequency component as well. In simulation set (3), the *t*-process was able to recover the PSD accurately at lower frequencies and was even able to detect deviations present at the higher-frequency spectrum in a few realizations. In contrast, the standard power-law model was only able to accurately recover the PSD for all frequencies in simulation set (1). In simulation set (2), it overestimated the PSD across all frequencies and was not able to measure some of the subtleties at the higher frequencies in simulation set (3). Across all three simulation types, the free spectral model was able to accurately recover the PSD at lower frequencies, but due to its model-agnostic nature, it was contaminated by the white noise floor, which significantly reduces the information gained at higher frequencies (as is seen in Figures 1, 4, and 7). Furthermore, we were able to measure the significance of deviations measured from the power law by computing the Hellinger distances and comparing to the distribution of Hellinger distances for the power-law simulations.

While the observed nanohertz GWB is consistent with predictions for a GWB generated by SMBHBs, there are many other possible cosmological sources. An important difference between a GWB produced by SMBHBs and cosmological sources is that a GWB produced by SMBHBs is being generated by a discrete number of individual sources: this causes some degree of anisotropy in the GWB (G. Agazie et al. 2023c) as well as affecting the GWB PSD. In this paper, we have shown how we can use the *t*-process model to reconstruct the GWB PSD and detect deviations from a power law. With longer timing baselines and better telescopes, PTA sensitivity will improve, allowing for more robust spectral characterization of the common process spectrum. If or when deviations from a power-law PSD arise, the *t*-process model will play a critical role in assessing their significance and allow the community to ultimately assess whether the GWB source is astrophysical or cosmological.

Acknowledgments

We thank Aaron Johnson for help with generating the simulated data sets with `enterprise`. We thank Luke Kelly and Emiko Gardiner for help in generating the astrophysical realizations using `holodeck`. We also thank Gabriel Freedman for useful discussions. The authors are members of the NANOGrav collaboration, and this work was supported by NSF Physics Frontiers Center award number 2020265. J.S. is supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship under award AST-2202388.

Appendix
Injected Parameters

We present the injected values for the intrinsic red noise parameters (Table 1) and the injected hyperparameters for the astrophysical simulation (Table 2).

Table 1
The Injected Intrinsic Red Noise Parameters (Spectral Index and Log Amplitude) for All the Simulations, as well as the Common Process Power Law that was Used in the First Two Types of Simulations

Pulsar	γ	$\log_{10} A$	Pulsar	γ	$\log_{10} A$
B1855+09	3.718	−14.752	J1730−2304	3.007	−15.696
B1937+21	4.002	−13.563	J1738+0333	3.589	−16.711
B1953+29	2.283	−12.918	J1741+1351	3.021	−16.609
J0023+0923	3.095	−16.546	J1744−1134	3.090	−16.546
J0030+0451	3.834	−16.228	J1745+1017	2.370	−11.838
J0340+4130	3.300	−16.929	J1747−4036	2.712	−12.612
J0406+3039	3.508	−16.359	J1751−2857	3.394	−16.617
J0437−4715	3.225	−16.721	J1802−2124	1.826	−12.284
J0509+0856	3.370	−15.807	J1811−2405	3.341	−16.867
J0557+1551	3.325	−16.268	J1832−0836	3.286	−16.902
J0605+3757	3.411	−16.251	J1843−1113	3.251	−16.553
J0610−2100	3.386	−12.815	J1853+1303	1.890	−14.393
J0613−0200	2.782	−15.311	J1903+0327	1.603	−12.191
J0636+5128	3.258	−16.860	J1909−3744	3.282	−17.226
J0645+5158	1.872	−14.376	J1910+1256	3.112	−16.736
J0709+0458	3.616	−15.617	J1911+1347	3.107	−16.757
J0740+6620	3.232	−17.087	J1918−0642	3.137	−16.886
J0931−1902	3.243	−16.879	J1923+2515	3.247	−17.103
J1012+5307	0.627	−12.641	J1944+0907	2.550	−14.442
J1012−4235	3.513	−15.997	J1946+3417	1.046	−12.462
J1022+1001	2.694	−15.176	J2010−1323	3.134	−16.934
J1024−0719	3.197	−16.755	J2017+0603	3.269	−17.160
J1125+7819	3.193	−16.863	J2033+1734	3.371	−16.801
J1312+0051	3.402	−16.603	J2043+1711	3.238	−17.166
J1453+1902	3.360	−16.773	J2124−3358	3.444	−16.565
J1455−3330	3.231	−16.936	J2145−0750	0.781	−12.930
J1600−3053	2.806	−16.507	J2214+3000	3.094	−16.913
J1614−2230	3.292	−17.095	J2229+2643	3.344	−17.017
J1630+3734	3.386	−16.285	J2234+0611	3.325	−16.138
J1640+2224	3.265	−17.170	J2234+0944	3.269	−16.826
J1643−1224	1.103	−12.302	J2302+4442	3.349	−16.910
J1705−1903	0.946	−11.929	J2317+1439	3.106	−17.075
J1713+0747	2.469	−15.950	J2322+2057	3.240	−17.014
J1719−1438	3.380	−16.506	GWB (sims. 1 & 2)	4.333	−14.699

Table 2
The Astrophysical Parameters Injected Into the PTA for the Holodeck Simulations

Parameters	Injected Value	Prior
Hard τ	3.133 Gyr	$\mathcal{U}(0.1, 11)$
Hard r_0	1000 pc	⋯
Hard $r_{\rm char}$	100 pc	⋯
Hard $\gamma_{\rm inner}$	−0.467	$\mathcal{U}(-1.5, 0.5)$
Hard $\gamma_{\rm outter}$	2.5	⋯
GSMF $\log_{10} \phi_0$	−2.140	$\mathcal{N}(-2.56, 0.4)$
GSMF ϕ_z	−0.6	⋯
GSMF $\log_{10} M_{\rm char0}$	11.370	$\mathcal{N}(10.9, 0.4)$
GSMF $M_{{\rm char},z}$	0.11	⋯
GSMF alpha0	−1.21	⋯
GSMF alphaz	−0.03	⋯
GPF frac norm all q	0.033	⋯
GPF m_α	0	⋯
GPF q_γ	0	⋯
GPF z_β	1	⋯
GPF max frac	1	⋯
GMT norm	1.514	$\mathcal{U}(0.2, 5)$
GMT m_α	0	⋯
GMT q_γ	−1	⋯
GMT z_β	−.5	⋯
MMB $\log_{10} m_{\rm amp}$	8.498	$\mathcal{N}(8.6, .2)$
MMB plaw	1.1	⋯
MMB scatter dex	0.292	$\mathcal{N}(0.32, 0.15)$
$f_{\rm bulge}$	0.615	⋯

Note. The parameters are taken directly from the NANOGrav 15 yr constraints paper or are randomly sampled from astrophysical priors (if given) given in the same paper. The MMBulge relation is from Kormendy & Ho (J. Kormendy & L. C. Ho 2013).

ORCID iDs

Shashwat C. Sardesai https://orcid.org/0009-0006-5476-3603
Joseph Simon https://orcid.org/0000-0003-1407-6607
Sarah J. Vigeland https://orcid.org/0000-0003-4700-9072

References

Afzal, A., Agazie, G., Anumarlapudi, A., et al. 2023, ApJL, 951, L11
Agazie, G., Antoniadis, J., Anumarlapudi, A., et al. 2024a, ApJ, 966, 105
Agazie, G., Anumarlapudi, A., Archibald, A. M., et al. 2023a, ApJL, 951, L8
Agazie, G., Anumarlapudi, A., Archibald, A. M., et al. 2023b, ApJL, 952, L37
Agazie, G., Anumarlapudi, A., Archibald, A. M., et al. 2023c, ApJL, 956, L3
Agazie, G., Baker, P. T., Bécsy, B., et al. 2024b, arXiv:2404.07020
Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2020, ApJ, 900, 102
Bécsy, B., Cornish, N. J., & Kelley, L. Z. 2022, ApJ, 941, 119
Burke-Spolaor, S., Taylor, S. R., Charisi, M., et al. 2019, A&ARv, 27, 5
CHIME/Pulsar Collaboration, Amiri, M., Bandura, K. M., et al. 2021, ApJS, 255, 5
Cordes, J. M., & Downs, G. S. 1985, ApJS, 59, 343
Detweiler, S. 1979, ApJ, 234, 1100
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., 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
EPTA CollaborationInPTA Collaboration, Antoniadis, J., et al. 2023, A&A, 678, A50
EPTA CollaborationInPTA Collaboration, Antoniadis, J., et al. 2024, A&A, 685, A94
Fastidio, F., Gualandris, A., Sesana, A., Bortolas, E., & Dehnen, W. 2024, MNRAS, 532, 295
Foster, R. S., & Backer, D. C. 1990, ApJ, 361, 300
Goncharov, B., Shannon, R. M., Reardon, D. J., et al. 2021, ApJL, 917, L19
Goncharov, B., Thrane, E., Shannon, R. M., et al. 2022, ApJL, 932, L22
Hellinger, E. 1909, JRAM, 1909, 210
Hellings, R. W., & Downs, G. S. 1983, ApJL, 265, L39
Hobbs, G., Coles, W., Manchester, R. N., et al. 2012, MNRAS, 427, 2780
Jaffe, A. H., & Backer, D. C. 2003, ApJ, 583, 616
Jones, M. L., McLaughlin, M. A., Lam, M. T., et al. 2017, ApJ, 841, 125
Kocsis, B., & Sesana, A. 2011, MNRAS, 411, 1467
Kormendy, J., & Ho, L. C. 2013, ARA&A, 51, 511
Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2016, ApJ, 819, 155
Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2017, ApJ, 834, 35
Lam, M. T., McLaughlin, M. A., Arzoumanian, Z., et al. 2019, ApJ, 872, 193
Lamb, W. G., & Taylor, S. R. 2024, ApJL, 971, L10
Lower, M. E., Bailes, M., Shannon, R. M., et al. 2020, MNRAS, 494, 228
McWilliams, S. T., Ostriker, J. P., & Pretorius, F. 2012, arXiv:1211.4590
Miles, M. T., Shannon, R. M., Bailes, M., et al. 2023, MNRAS, 519, 3976
Milosavljević, M., & Merritt, D. 2003, in AIP Conf. Ser. 686, The Astrophysics of Gravitational Wave Sources, ed. J. M. Centrella (Melville, NY: AIP), 201
Mingarelli, C. M. F., Lazio, T. J. W., Sesana, A., et al. 2017, NatAs, 1, 886
Parthasarathy, A., Bailes, M., Shannon, R. M., et al. 2021, MNRAS, 502, 407
Phinney, E. S. 2001, arXiv:astro-ph/0108028
Rajagopal, M., & Romani, R. W. 1995, ApJ, 446, 543
Reardon, D. J., Zic, A., Shannon, R. M., et al. 2023, ApJL, 951, L6
Roebber, E. 2019, ApJ, 876, 55
Rosado, P. A., Sesana, A., & Gair, J. 2015, MNRAS, 451, 2417
Sampson, L., Cornish, N. J., & McWilliams, S. T. 2015, PhRvD, 91, 084055
Sazhin, M. V. 1978, SvA, 22, 36
Sesana, A. 2013, CQGra, 30, 244009
Sesana, A., Vecchio, A., & Colacino, C. N. 2008, MNRAS, 390, 192
Shannon, R. M., & Cordes, J. M. 2010, ApJ, 725, 1607
Siemens, X., Ellis, J., Jenet, F., & Romano, J. D. 2013, CQGra, 30, 224015
Simon, J. J. 2017, PhD thesis, Univ. of Wisconsin, Milwaukee
Susobhanan, A., Gopakumar, A., Hobbs, G., & Taylor, S. R. 2020, PhRvD, 101, 043022
Taylor, S. R., Baker, P. T., Hazboun, J. S., Simon, J., & Vigeland, S. J., 2021 enterprise_extensions v2.4.3, https://github.com/nanograv/enterprise_extensions
Tiburzi, C., Hobbs, G., Kerr, M., et al. 2016, MNRAS, 455, 4339
Tiburzi, C., Shaifullah, G. M., Bassa, C. G., et al. 2021, A&A, 647, A84
Vallisneri, M., Taylor, S. R., Simon, J., et al. 2020, ApJ, 893, 112
van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2
van Haasteren, R. 2024, ApJS, 273, 23
Wyithe, J. S. B., & Loeb, A. 2003, ApJ, 590, 691
Xu, H., Chen, S., Guo, Y., et al. 2023, RAA, 23, 075024
Zic, A., Hobbs, G., Shannon, R. M., et al. 2022, MNRAS, 516, 410