



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Constraining the solution to the last parsec problem with pulsar timing

Laura Sampson, Neil J. Cornish, and Sean T. McWilliams

Phys. Rev. D **91**, 084055 — Published 24 April 2015

DOI: [10.1103/PhysRevD.91.084055](https://doi.org/10.1103/PhysRevD.91.084055)

Constraining the Solution to the Last Parsec Problem with Pulsar Timing

Laura Sampson,¹ Neil J. Cornish,¹ and Sean T. McWilliams²

¹*Department of Physics, Montana State University, Bozeman, MT 59717, USA.*

²*Department of Physics and Astronomy, West Virginia University, Morgantown, WV 26506, USA.*

(Dated: April 7, 2015)

The detection of a stochastic gravitational-wave signal from the superposition of many inspiraling supermassive black holes with pulsar timing arrays (PTAs) is likely to occur within the next decade. With this detection will come the opportunity to learn about the processes that drive black-hole-binary systems toward merger through their effects on the gravitational-wave spectrum. We use Bayesian methods to investigate the extent to which effects other than gravitational-wave emission can be distinguished using PTA observations. We show that, even in the absence of a detection, it is possible to place interesting constraints on these dynamical effects for conservative predictions of the population of tightly bound supermassive black-hole binaries. For instance, if we assume a relatively weak signal consistent with a low number of bound binaries and a low black-hole-mass to galaxy-mass correlation, we still find that a non-detection by a simulated array, with a sensitivity that should be reached in practice within a few years, disfavors gravitational-wave-dominated evolution with an odds ratio of $\sim 30:1$. Such a finding would suggest either that all existing astrophysical models for the population of tightly bound binaries are overly optimistic, or else that some dynamical effect other than gravitational-wave emission is actually dominating binary evolution even at the relatively high frequencies/small orbital separations probed by PTAs.

PACS numbers: 04.30.-w, 04.30.Tv, 97.60.Lf

I. INTRODUCTION

The history of structure formation in our universe is understood to be one of hierarchical building, in which small galaxies merge to form larger galaxies, and large-scale structure emerges in the course of these interactions. It is commonly assumed that this history of galactic mergers implies a companion history of black-hole mergers [1], which suggests the existence of a large population of supermassive-black-hole (SMBH) binaries that emit gravitational radiation. There is some observational evidence that SMBH binaries do, in fact, exist [2, 3]. If the SMBHs are at separations where they emit appreciable energy in gravitational waves (GWs), these GWs will add together to form a stochastic background of gravitational radiation, which can be searched for using pulsar timing arrays (PTAs) [4–7].

To search for GWs, PTAs take advantage of the fact that millisecond pulsars are the best natural clocks in the universe. These rapidly spinning neutron stars emit radio waves that sweep past the Earth with stunning regularity - for the best systems, we can predict the time of arrival (TOA) of the radio pulses with accuracies of tens of nanoseconds. These predictions are made using timing models that include the orbits of the pulsars and the Earth, spin-down of the individual pulsars, changes in the instruments, and many other effects [8–10]. Once the parameters of this timing model are fitted, the model is subtracted from the TOA data, leaving behind timing residuals. If there is a stochastic background of GW radiation permeating the space between the Earth and the observed pulsars, these timing residuals will include the GW signal, manifesting as a red “noise” source.

In an individual pulsar, the effect of this stochastic

background would be impossible to detect, as it could be modeled away as one of the other known red noise processes. However, Hellings and Downs [11] showed that the GW signal due to a stochastic background of binaries will be correlated between different pulsars, with a correlation function that depends on the angle between two pulsars on the sky. It is by searching for red power content that possesses this correlation that PTAs can detect a stochastic GW signal [12–15].

It is certainly true that at some point in the inspiral of SMBH binaries, the emission of GWs and the concordant loss of energy and angular momentum becomes the primary driver toward merger. In order for these systems to merge within a time period shorter than the current age of the universe, however, it is not possible for GWs to be solely responsible for the evolution of the SMBH system after the initial galactic merger. There must be other mechanisms that drive the black holes near enough to each other for GW emission to dominate. These mechanisms include phenomena such as stellar scattering and the accretion of gas, which are thought to become less efficient as the black holes approach merger [16–22]. For the elliptical galaxies that host BHs massive enough to emit GWs in the PTA band, gas is thought to be scarce, and is therefore unlikely to drive SMBH binaries close enough to become GW dominated. Major mergers of comparable mass galaxies have long been known to result in gas poor elliptical galaxies [23–25]. The assumption that PTA sources in particular are exceedingly gas poor is strongly supported by the large drop in the observed quasar luminosity function for $z < 2$ [26], which implies a decrease in the gas-mass accretion rate of PTA candidate sources at $z < 1$ by 3-4 orders of magnitude compared to the peak rates found at $z > 2$ [27]. Furthermore, once

the SMBH binaries reach orbital separations of $\mathcal{O}(1 \text{ pc})$, simple estimates imply that there are not enough stars available in the region of phase space capable of removing the necessary amount of energy and angular momentum for the SMBH binary to merge via GW emission within a Hubble time. This region of phase space for stars is referred to as the loss cone, and the depletion of the loss cone is often referred to as the last parsec problem [28]. There are many proposed resolutions to the last parsec problem, but they primarily involve mechanisms that can refill the loss cone (see [29] and references therein).

It has commonly been assumed in PTA analyses that the last parsec problem is solved by some unspecified mechanism, but that GWs are the dominant driver of binary mergers throughout the frequency range ($\sim 10^{-9} - 10^{-6} \text{ Hz}$) to which pulsar timing is sensitive. This is not, however, guaranteed to be the case, and it is important to remember that the SMBH binaries of interest emit GWs throughout their evolution, even when the driving evolutionary force is something other than GWs. This means that the stochastic background of GWs to which PTAs will be sensitive could have a spectral shape that is significantly different than that predicted by assuming purely GW-driven binary dynamics [30–36].

This reality can have both negative and positive consequences. The negative effect of gas- or stellar-driven binary evolution is that it causes the amplitude of the stochastic GW background within the PTA sensitivity band to be lower than what it would be if all mergers were GW-driven throughout the band. However, the presence of more complicated spectra, which could be shaped by many astrophysical processes, would give us the opportunity to learn about the galactic environments in which SMBH binaries evolve and merge.

In this paper we explore both of these possibilities, using the techniques of Bayesian inference and model selection. First, in Section II we discuss some of the processes that can affect the shape of the stochastic spectrum, and how to determine what that shape will be. In Section III, we describe our simulations and the analysis techniques used, and present the results of our analysis. We find that the detection of a stochastic GW background is not hindered, and can in fact be significantly helped, by using a more complicated model for the spectrum that includes the possibility of non-GW inspiral mechanisms. We also show the extent to which the parameters that describe this spectrum can be measured in the event of a detection, and constrained in the event of a non-detection. Throughout this section we analyze the effects that the choice of priors has on model selection and parameter estimation, and discuss which prior choice is appropriate for a particular goal. Finally, we summarize and conclude in Section IV. Unless otherwise noted, we will work in geometrized units where Newton’s constant G and the speed of light c satisfy $G = c = 1$.

II. SPECTRAL MODELS

The characteristic amplitude of the stochastic GW signal generated by a population of inspiralling SMBH binaries on circular orbits can be calculated via [33, 37, 38]

$$h_c^2(f) = \int_0^\infty dz \int_0^\infty d\mathcal{M} \frac{d^3 N}{dz d\mathcal{M} dt} \frac{dt}{d \ln f} h^2(f). \quad (1)$$

In this expression, $h(f)$ is the instantaneous GW strain amplitude emitted by a single circular binary with a Keplerian rest frame orbital frequency of $f/2$, since circular binaries emit gravitational waves at twice the orbital frequency, and is given by [39]

$$h(f) = \frac{8}{\sqrt{10}} \frac{\mathcal{M}^{5/3}}{D_L} (\pi f)^{2/3} \quad (2)$$

where D_L is the luminosity distance to the source. For a binary with component masses M_1 and M_2 , $\mathcal{M} \equiv (M_1 M_2)^{3/5} (M_1 + M_2)^{1/5}$ is a combination of masses known as the chirp mass, z is the redshift, and t is time measured in the binary rest frame. We note that the amplitude $h(f)$ applies regardless of whether the source is evolving primarily under the influence of gravitational waves, or is dominated by some other dynamical effect. The term

$$\frac{d^3 N}{dz d\mathcal{M} dt} \frac{dt}{d \ln f} = \frac{d^3 N}{dz d\mathcal{M} d \ln f} \quad (3)$$

is the differential number of inspiraling binaries per unit \mathcal{M} , z , and $\ln f$. The term $dt/d \ln f$ encodes the frequency evolution of the binary. For a population of circular binaries that are driven purely by GW emission, this frequency evolution is given by the standard quadrupole formula:

$$\frac{dt}{d \ln f} = \frac{5}{64\pi^{8/3}} \mathcal{M}^{-5/3} f^{-8/3}. \quad (4)$$

With this substitution, Eq. 1 becomes [37]

$$h_c^2(f) = \frac{4f^{-4/3}}{3\pi^{1/3}} \int \int dz d\mathcal{M} \frac{d^2 n}{dz d\mathcal{M}} \frac{\mathcal{M}^{5/3}}{(1+z)^{1/3}}, \quad (5)$$

where $d^2 n/(dz d\mathcal{M}) \equiv d^3 N/(dV_c d\mathcal{M} dz)$ is the comoving number density of merged remnants per unit \mathcal{M} and z , with V_c being the comoving volume. Finally, dV_c intervals are related to dz intervals via [37]

$$dV_c = \frac{4\pi D_L^2}{H_o(1+z)^2 \sqrt{\Omega_M(1+z)^3 + \Omega_\Lambda}} dz, \quad (6)$$

where H_o , Ω_M , and Ω_Λ are the Hubble constant, total (baryon and dark matter) mass-energy density, and dark energy density, respectively, and D_L is related to z in a

flat cosmology via

$$D_L(z) = \frac{1+z}{H_0} \int_0^z \frac{dz'}{\sqrt{\Omega_M(1+z')^3 + \Omega_\Lambda}} \quad (7)$$

(see [40] and references therein).

Eq. (5) is often re-parameterized as $h_c = A(f/\text{yr}^{-1})^{-2/3}$, where A is the amplitude at $f = 1\text{yr}^{-1}$. A is therefore the standard metric that has been used throughout the literature for characterizing the sensitivity of PTAs, although this only makes sense if we assume all binaries have circular orbits and are gravitational-wave dominated.

GWs cannot, however, be the only driver of SMBH coalescence - not if these binaries are to merge within a Hubble time. There are many other mechanisms, such as interactions between the SMBH binary and surrounding stars and gas, that draw the binaries together after their galaxies merge. These mechanisms will alter the calculation above exclusively via a change in $dt/d\ln f$; generically, $dt/d\ln f$ can be no smaller than the prediction of Eq. 4, since any other dominating mechanism will cause binaries to evolve more rapidly through frequency. This will result in a different frequency dependence in the final power spectrum, such that the spectral slope will be smaller at frequencies dominated by other mechanisms.

To see how other dynamical effects will impact the shape of the spectrum, we can recast Eq. (1) as

$$h_c^2(f) \propto f \frac{dN}{df} h^2(f) \propto \frac{f^{7/3}}{df/dt}. \quad (8)$$

This re-writing is made under the assumption that dN/dt is constant, i.e. that the merger rate is not evolving much with time. We can see from Eq. (8) that the spectrum depends on the emitted gravitational waves, $h(f)$, scaled by an amplitude, dN/df , that encodes the number of binaries in a given frequency bin. The final step emphasizes that, given a constant dN/dt , the number of binaries per frequency bin is determined solely by the frequency evolution of these binaries, df/dt .

Then we can write

$$\frac{df}{dt} = \sum_i \left(\frac{df}{dt} \right)_i, \quad (9)$$

where the sum extends over all the physical processes that drive the black holes to inspiral. For instance, for GW-driven evolution

$$\left(\frac{df}{dt} \right)_{\text{GW}} \propto f^{11/3}. \quad (10)$$

Given the strong frequency dependence of GW-driven inspiral, we will assume that at least the final stage of binary evolution is dominated by GW emission, and pull this term out of the sum in Eq. 9. We can then recast

the expression in Eq. (8) in the form

$$h_c(f) = A \frac{(f/f_{\text{year}})^{-2/3}}{(1 + \sum_i b_i (f/f_{\text{year}})^{(\gamma_i - 11/3)})^{1/2}}, \quad (11)$$

where f_{year} is the fiducial frequency of 1/year, the γ_i 's encode the frequency dependence introduced by a particular astrophysical effect, the b_i are the dimensionless relative amplitudes of the astrophysical processes driving the binary toward merger scaled relative to the GW driven case at $f = f_{\text{year}}$, and A is a dimensionless amplitude. Note that unless the $a_i = 0$, A is *not* equal to the characteristic amplitude at a frequency of 1/year. On the other hand, for most scenarios we expect the $a_i \ll 1$, so to a good approximation, A is close to equaling $h_c(f = f_{\text{year}})$.

Before we consider specific alternative dynamical effects to GW-driven inspiral, it is useful to briefly review our understanding of the evolutionary histories of SMBH binaries. The early SMBH dynamics within merging galaxies is driven by dynamical friction [41] wherein each SMBH and its concomitant stellar bulge and dark matter subhalo gravitationally focus first the dark matter and then the baryonic content of the companion galaxy. This focusing of excess mass behind the SMBH induces a drag, thereby causing the SMBH to drop further into the gravitational potential well of the merging pair. Dynamical friction continues to operate until the SMBH orbital velocity exceeds the stellar velocity dispersion, at which point the binary is said to have hardened. Subsequent evolution of the binary is driven by stellar scattering, in which background stars on centrophilic orbits (i. e. stars outside the SMBH binary that are scattered onto low angular-momentum orbits toward the center of the gravitational potential) interact with the SMBH binary in a three-body interaction, and remove energy and angular momentum from the SMBH binary as a result [21]. The last parsec problem occurs when there are insufficient stars on the necessary orbits to continue driving the SMBH binary toward the GW-dominated regime. There are several proposed solutions to this problem, including any form of asymmetry in the gravitational potential (see [42] and references therein), but all of these solutions can be considered mechanisms to refill the loss cone, so that stellar-induced hardening via three-body interaction can continue until GWs are able to efficiently drive the binary onward to merging.

The semi-major axis a of a SMBH binary that interacts with a background of stars via three-body scattering evolves according to [21]

$$\frac{da}{dt} = \frac{a^2 \rho}{\sigma} H, \quad (12)$$

where ρ is the density of background stars and σ is their velocity dispersion, and H is approximately constant and was estimated in [21] via N-body simulations. Thus, for stellar scattering $df/dt \sim f^{1/3}$. If we assume that a binary is driven by stellar scattering immediately prior to

being driven by the emission of GWs, we can find the frequency at which GWs begin to dominate by equating $(da/dt)_{\text{stars}}$ from Eq. 12 to $(da/dt)_{\text{GW}}$ under the influence of gravitational-wave emission. This crossover frequency has been estimated as $f_b \approx 10^{-9} M_8^{-7/10} \eta^{-3/10}$ Hz, where M_8 is the total mass of the binary in hundreds of millions of solar masses and η is the symmetric mass ratio [43] (note that there is a typo in this reference, but we have confirmed that the expression we present here is correct). In what follows, we will call such transition frequencies “bend frequencies”, since the slope of $h_c(f)$ changes as one hardening process takes over from another. For a two part da/dt model (or, equivalently, df/dt), where some other dynamical mechanism gives way to GW-driven evolution, we have

$$h_c(f) = A \frac{(f/f_{\text{year}})^{-2/3}}{(1 + (f_b/f)^\kappa)^{1/2}}, \quad (13)$$

with $\kappa = 10/3$ being the appropriate value for the specific case of stellar scattering. We must mention that stellar scattering is not the only mechanism that could potentially influence the inspiral of SMBH binaries. Other possible effects include torquing of the binary from a circumbinary disks, where the details of the effect depend sensitively on the dissipative physics of the disk, and gas-driven evolution due to massive inflows of gas that can be triggered by dynamical instabilities during the galactic merger [16, 34–36, 44–48]. In general, a particular physical effect will lead to an expression for da/dt , which can be mapped to values of f_b and κ in the GW spectrum. Given the presumed absence of gas in the massive elliptical galaxies that host PTA sources, disk- or gas-driven inspiral is thought to be less probable than stellar scattering, so we will focus primarily on the latter as the dominant effect preceding GW-driven inspiral. We also mention that, even in the case of stellar-driven inspiral, it is possible that the balance of energy and angular momentum removal could drive SMBH binaries to nonzero eccentricities when they transition to being GW dominated [21]. Having said this, zero eccentricity appears to be an attractor solution for binaries in numerical experiments with plausible initial conditions of the stellar population [21], so we will focus on circular binaries hereafter.

Realistic astrophysical binaries will be driven by multiple mechanisms throughout their evolutionary history, which is reflected in the summation in Eq. (9). This will result in a GW power spectrum with several slopes and several transition frequencies. It is likely, however, that in order to measure the parameters of such a complicated power spectrum, it will be necessary to make a high signal to noise ratio (SNR) detection of the stochastic GW background. In the early days of PTA data analysis, a simple model that allows for a single transition between GW-dominated evolution and some other power law will most likely suffice.

Figure 1 shows examples of the different types of sim-




FIG. 1: Examples of the five different spectral shapes we will investigate. (I) has $f_b = 10^{-8}$ Hz and $\kappa = 29/6$, (II) has $f_b = 10^{-8}$ Hz and $\kappa = 10/3$, (III) has $f_b = 3 \times 10^{-8}$ Hz and $\kappa = 29/6$, (IV) has $f_b = 3 \times 10^{-8}$ Hz and $\kappa = 10/3$, and (V) (the dashed line in each panel) is the signal from binaries that are driven purely by GWs throughout the band. Here, we plot the power spectral density, $S_h(f) = h_c^2/(2f)$.

ulated signals we will analyze in Section III. Plotted in this figure are the four basic shapes of spectra which we will investigate, along with the baseline spectrum of a purely GW-driven binary population. The spectra include two different bend frequencies, $f_b = 1 \times 10^{-8}$ Hz and $f_b = 3 \times 10^{-8}$ Hz (which are quite arbitrary choices - the theoretical uncertainty of what realistic values are for f_b is very large, and so a theoretically-motivated value is difficult to define), and two different slopes for the low-frequency part of the data, $\kappa = 10/3$, which corresponds to stellar scattering, and $\kappa = 26/9$, which is chosen to be a somewhat more dramatic bend. These different shapes are labeled with roman numerals. Further details of the simulated data will be discussed in the next section.

The final astrophysical function that enters into Eq. (13), is the merger rate, and it is this rate that determines the overall amplitude of the GW background. When performing a Bayesian analysis with the spectral model (13), it is necessary to provide priors on all the parameters. A prior on a particular parameter encodes any information that is known, or believed, *a priori* about that parameter. In the analysis presented in the next section, we will focus on two different astrophysically motivated priors on A , as well as priors that are uniform in A and in $\ln A$. The two astrophysical priors are motivated by the work of McWilliams *et al.* [32, 33] (Model A), and Sesana and Ravi *et al.* [34, 36] (Model B). Both models can be approximated as Gaussian distributions in $\ln A$. The prior for Model A is centered at $A = 4.1 \times 10^{-15}$, with an uncertainty in $\ln A$ of $\sigma = 0.6$ [33]. The Model B prior is centered at $A = 10^{-15}$, with an uncertainty in $\ln A$ of $\sigma = 0.5$. Since these prior distributions are Gaussian, nei-

ther requires the specification of an upper or lower limit on A . The third prior that we consider, which is uniform in A , has the natural lower limit of $A = 0$. To calculate the upper limit for this prior, we impose the (rather conservative) constraint that the energy density in GWs in the PTA observing band cannot close the universe. An alternative would be to demand that the energy density in GWs from PTA sources cannot close the universe, which would extend the integration over the energy density to include the out-of-band merger and ringdown of the black hole binaries. We chose to use the frequency range covered by the PTA observing band since those are the frequencies that are directly constrained by the data. To use this requirement to impose an upper limit on A we relate h_c to the energy density per logarithmic frequency interval via

$$\Omega_{\text{GW}}(f) = \frac{2\pi^2}{3H_0^2} f^2 h_c^2(f), \quad (14)$$

then impose that

$$\int_{f_{\text{min}}}^{f_{\text{max}}} \Omega_{\text{GW}}(f) d \ln f < 1. \quad (15)$$

The lower frequency limit of this integral is unimportant since the integral is dominated by the high frequency behavior. Performing the integral results in a limit on A of the form

$$A^2 < f_y^{-4/3} f_{\text{max}}^{-2/3} \frac{H_0^2}{\pi^2}. \quad (16)$$

For $f_{\text{max}} = 5.5 \times 10^{-7}$ Hz, this implies that $A < 4.4 \times 10^{-12}$ (assuming a Hubble parameter of $H_0 = 67.8 \text{ km/s/Mpc}$, which is the upper limit we use when employing a uniform prior in A throughout our analyses. We take the uniform prior on the logarithm of the amplitude to share the same upper limit derived for the uniform amplitude prior, but there is no natural value for the lower limit. This prevents us from using this prior to set upper limits, but by choosing a lower limit several orders of magnitude below the noise level, it is possible to use the uniform logarithmic prior for parameter estimation once a detection has been made.

The other parameters that enter into our model for the GW spectrum are f_b and κ , as discussed above. The prior we choose to employ on f_b is uniform in $\ln(f_b)$, indicating our lack of knowledge about even the order of magnitude of the frequency at which the GW slope may bend, and has a lower limit of $f_b = 2.5 \times 10^{-9}$ Hz (which is below the lower edge of the sensitivity band for a 10 year data span), and an upper limit of $f_b = 10^{-7}$ Hz. This upper limit is somewhat arbitrary, but it does indicate our belief that GWs must, at some high enough frequency, dominate the evolution of SMBH binaries. For κ , we use a prior that is uniform in κ , with a range of 0 to $23/3$ (the non-integer value for the upper end of the range is due to us changing our conventions for the

parameterization of the spectrum late in the project).

We note that we have chosen to study our ability to measure the low-frequency slope, bend frequency, and amplitude for a set of simulated data, rather than trying to assess how efficiently or frequently the last parsec problem is solved, or what the correct correlation is between the black hole mass and a particular property of the host galaxy. Our principle motivation for this choice is that the astrophysically-motivated distributions only differ by a factor of $\sim 2-3$ in predicted signal amplitude between the most conservative and the most optimistic estimates, despite making different assumptions about the solution to the last parsec problem and the correlation between black-hole mass and host-galaxy properties. Furthermore, there are multiple elements that contribute to this level of amplitude uncertainty, including the overall galaxy merger rate and its dependence on mass and environment, in addition to the last parsec solution and the choice of black-hole mass - host-galaxy-property correlation. Given our comparative ignorance of the low-frequency signal slope and the transition frequency between the dynamical process dominating the final parsec and GW-driven evolution, we choose to focus our study on constraining these parameters for a given amplitude.

Another approach we could have taken with the amplitude prior is to separate the uncertainties into an observational part (from factors such as the rate of galaxy mergers and the M - σ relation) and a theoretical part (from factors such as the efficiency of dynamical friction in hardening the binary and the fraction of systems where the last parsec problem is overcome). This could be done by writing the prior as a Gaussian distribution in $\ln A$, centered at some value $\ln \bar{A}$, with width σ_{obs} to account for the observational uncertainties. The central value, $\ln \bar{A}$, would then be a hyper-parameter to be determined by the data. If we quantify the uncertainty on the central value \bar{A} as $\bar{A} = \eta A_*$, where A_* is some reference value and η encodes the theoretical uncertainty in the merger efficiency, then assigning a Gaussian prior on the hyper-parameter $\ln \eta$ of width σ_{th} centered on η_* , and marginalizing over the hyper parameter, leads to Gaussian distributions of the form used for Models A and B with width $\sigma = \sqrt{\sigma_{\text{obs}}^2 + \sigma_{\text{th}}^2}$. Alternatively, a uniform prior on $\ln \eta$ over some range leads to a roughly uniform prior on $\ln A$ between $\ln(\eta_{\text{min}} A_*)$ and $\ln(\eta_{\text{max}} A_*)$ with rounded edges of width σ_{ons} . Thus, the analyses we consider are equivalent to a hierarchical Bayesian analysis that separates out the theoretical and observational uncertainties for certain choices of the hyper-prior.

III. SIMULATED DATA AND ANALYSIS

Our simulated data set consists of the timing residuals from 20 pulsars, randomly distributed on the sky, and observed for 10 years with a two-week cadence. The data from each pulsar is generated including white noise at a level of 200 ns, and no red noise. We then recover this

simulated signal using a model that includes the three parameters that describe the GW spectrum, A, f_b, κ , and the noise parameters for each pulsar. These include a white noise level, red noise level, and red noise slope. The white noise is fully described by the amplitude of its power spectral density (PSD), which we label S_n . The prior on S_n was uniform in $\ln(S_n)$, and ranged from $S_n = 4 \times 10^{-18} \text{ Hz}^{-1/2}$ to $S_n = 10^{-2} \text{ Hz}^{-1/2}$. The red noise is parameterized by its PSD amplitude, S_r , and by a slope, r , as $S_r(f) = S_r(f/f_{\text{year}})^r$. The prior on the red noise amplitude was also uniform in $\ln(S_r)$, with the same range as S_n , and the prior on the slope r was uniform from -2 to -6 .

Finally, we include two parameters for each pulsar that encode the effects of the timing model on the timing residuals. As discussed briefly in the introduction, the timing model used to predict the TOA of radio pulses from a given pulsar includes a large set of parameters [8–10]. In this analysis we only consider the two timing model parameters that have the greatest effect on the low-frequency sensitivity. These are the quadratic and linear terms in the spin down model for each pulsar, which take the form

$$P(t) = P_0 + \dot{P}_0 t + \ddot{P}_0 t^2 + \dots \quad (17)$$

Here, P_0 is the initial spin period of the pulsar, and \dot{P}_0 and \ddot{P}_0 encode the way that this spin evolves in time.

In order to speed up our analysis, we choose to perform our calculations in the Fourier domain. We therefore need to understand how the quadratic and linear terms in the timing model translate to effects on the timing residuals as a function of frequency. The Fourier transform of the timing model is given by

$$\tilde{P}(f_k) = \int_0^T P(t) e^{i2\pi f_k t} dt, \quad (18)$$

where $f_k = k/T$ for integer k , and T is the observation time. This integral evaluates to

$$\tilde{P}(f_k) = TP_0 \delta_{k0} - \frac{i\dot{P}_0 T^2}{2\pi k} - \frac{i\ddot{P}_0 T^3}{2\pi k} + \frac{\ddot{P}_0 T^3}{2\pi^2 k^2} + \dots \quad (19)$$

The $k = 0$ term is simply a constant offset that we can ignore. Writing $a = \ddot{P}_0 T^3 / (2\pi^2)$ and $b = -(\dot{P}_0 T^2 + \ddot{P}_0 T^3) / (2\pi)$, the timing model for each pulsar can be written as

$$\tilde{P}_k = \frac{a}{k^2} + \frac{ib}{k}. \quad (20)$$

A model of this form, with independent a and b for each pulsar, is then subtracted from the TOAs.

The full set of parameters in our model thus consists of the five noise/timing parameters for each pulsar, and three parameters to describe the GW background - A, f_b, κ . With this set of parameters, the likelihood is

defined by [49]

$$p(d|s, n) = \frac{1}{\sqrt{(2\pi)^L \det C}} \exp \left(-\frac{1}{2} \sum_{ab} \sum_{ij} r_{ai} C_{(ai)(bj)}^{-1} r_{bj} \right), \quad (21)$$

where C is the covariance matrix, which depends on both the noise in the individual pulsars and on the GW background, and $r = d - s$ denotes the timing residuals after the subtraction of the timing model s from the data d . The indices a and b label individual pulsars, and run from 1 to the number of pulsars, N_p . The indices i and j label the data samples, i.e. individual frequency bins. Since our simulated data is stationary, the correlation matrix is diagonal in i, j and $C_{(ai)(bj)} \rightarrow C_{ab}(f)$.

The timing model parameters for each pulsar enter the likelihood in the timing residuals; they are subtracted from the TOAs before the likelihood is evaluated. The red and white noise contributions for each pulsar enter along the diagonal of the covariance matrix. Finally, the GW signal enters via the Hellings and Downs [11] correlation matrix, which has the form

$$H_{ab} = \frac{1}{2} + \frac{3(1 - \cos \theta_{ab})}{4} \ln \left(\frac{1 - \cos \theta_{ab}}{2} \right) - \frac{1 - \cos \theta_{ab}}{8} + \frac{1}{2} \delta_{ab}. \quad (22)$$

The full covariance matrix is then given by

$$C_{ab}(f) = S_h(f) \frac{H_{ab}}{3} + \delta_{ab} \{S_{n_a} + S_{r_a}(f/f_y)^{r_a}\}, \quad (23)$$

where $S_h(f)$ is the PSD of the GW background, S_{n_a} is the PSD of the white noise, S_{r_a} is the amplitude of the PSD of the red noise, and r_a is the slope of the red noise.

A. Methods

Our analysis is carried out within the framework of Bayesian inference, using the technique of Markov chain Monte Carlo (MCMC). To calculate Bayes factors¹, we must calculate the evidence for each model, which necessitates performing an integral over the full parameter space. For this integral, we use the technique of thermodynamic integration [50–52]. This technique necessitates the use of parallel tempering, in which multiple chains are run at different ‘temperatures,’ which are defined by changing the likelihood to

$$p(d|s, n, \beta) = p(d|s, n)^\beta, \quad (24)$$

¹ Given equal prior belief in the validity of two models, A and B, the Bayes factor, B_{AB} , between models A and B, given the observed data, is the ‘betting odds’ that model A is the better theory, rather than model B.

Type	sub-Type	$A \times 10^{15}$	f_b (Hz)	κ
I	a	0.08	3×10^{-8}	29/6
	b	2.0	3×10^{-8}	29/6
	c	4.0	3×10^{-8}	29/6
II	a	0.08	3×10^{-8}	10/3
	b	2.0	3×10^{-8}	10/3
	c	4.0	3×10^{-8}	10/3
III	a	0.08	10^{-8}	29/6
	b	2.0	10^{-8}	29/6
	c	4.0	10^{-8}	29/6
IV	a	0.08	10^{-8}	10/3
	b	2.0	10^{-8}	10/3
	c	4.0	10^{-8}	10/3
V	a	0.08	0	N/A
	b	2.0	0	N/A
	c	4.0	0	N/A

TABLE I: The different types of simulated signals for Figures 2 and 3.

where β is analogous to an inverse temperature. This effectively ‘softens’ the likelihood, allowing the chains to effectively sample the full posterior. Chains with different temperatures are allowed to swap parameters with a probability given by the Hastings ratio

$$H_{i \leftrightarrow j} = \min \left(\frac{p(d|s_i, n_i, \beta_j) p(d|s_j, n_j, \beta_i)}{p(d|s_i, n_i, \beta_i) p(d|s_j, n_j, \beta_j)}, 1 \right). \quad (25)$$

The evidence for each model is then given by

$$\ln p(d) = \int_0^i E_\beta [\ln p(d|\vec{x})] d\beta, \quad (26)$$

where E_β denotes the expectation at inverse temperature β . Given equal prior belief in two models, the Bayes factor is then simply the evidence ratio between them. The technique of parallel tempering is useful not only as a means of calculating Bayes factors, but as a tool for improving the convergence of the MCMC runs.

In order to further speed convergence, we use a cocktail of jump proposals. These include small random-walk jumps of varying sizes, which we define as draws in each parameter from a Gaussian centered at the current location; parallel tempering swaps; and differential evolution proposals [53]. We find that the differential evolution jumps are particularly helpful in encouraging efficient exploration.

B. Detection and Model Selection

We first investigate the impact on detectability of the shape of the GW spectrum, both in the simulated signal and in the model used to extract it. We do this by simulating GW backgrounds of varying shapes and amplitudes, described in Table I, and recovering these signals

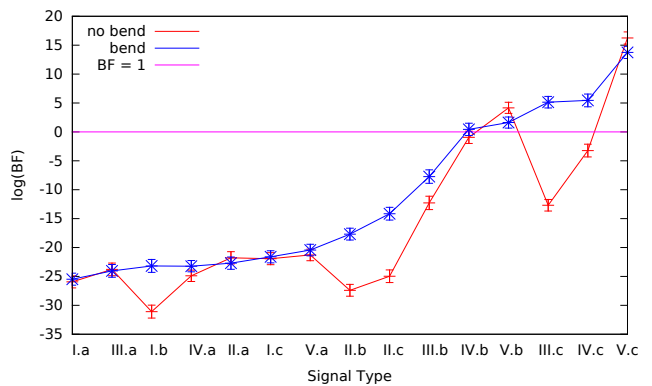


FIG. 2: The Bayes factors between noise and two GW signal models, both with uniform amplitude prior, one with a bend in the spectrum, the other assuming a purely GW driven evolution. A Bayes factor larger than unity indicates a detectable GW background. Unsurprisingly, the model that allows for a bend in the spectrum fares much better than the GW-driven model in cases where the simulated signal spectrum includes a bend (types I,II,III and IV), and performs only slightly worse when the simulated signal is purely GW-driven (type V).




FIG. 3: The Bayes factors between GW signal models with three different priors - uniform in amplitude and the Model A and B astrophysical amplitude priors - and noise. In this study, all the GW models allowed for a bend in the spectrum. A Bayes factor larger than unity indicates a detectable GW background. Here we see that both astrophysical priors outperform the uniform prior for detectability, and also that we are able to distinguish between the two astrophysical models.

using templates that either do or do not include a bend in the spectrum. The templates that do not include a bend are consistent with the belief that the SMBH binary population is driven purely by GW emission throughout the sensitivity band. Figure 2 shows the results of this study. The prior on the GW amplitude for the two signal models (with a bend and without) was uniform in amplitude. The horizontal axis is labeled by the simulated signals,

which have been arranged from left to right in ascending order of detectability for the model that includes a bend. The vertical axis shows the Bayes factor in favor of either signal model, where a Bayes factor larger than unity indicates detectability. There are two main features to notice in Figure 2. The first is that the model with no bend, i.e. for which the SMBH binaries are driven purely by GW emission throughout, does not perform appreciably better than the more complicated model that includes a bend for any simulated signal. This is true even when the signals themselves do not include a bend in the spectrum. This tells us that using the more complicated model for the purpose of detection will not significantly decrease our sensitivity to a GW background². The second, related feature is that the Bayes factor for the model that allows for a bend is consistently favored over the no-bend model when the simulated signal has a bend. From this we can see that we will be able to perform model selection between astrophysical models that predict a slope change in the spectrum, and those that predict a purely GW-driven evolution across the PTA frequency band.

The other main prediction of astrophysical models is reflected in the prior we impose on the amplitude of the GW background. In the previous section, we discussed the different amplitude priors from the merger models of McWilliams *et al.* [33] (Model A) and Sesana and Ravi *et al.* [34, 36] (Model B). In Figure 3, we demonstrate the ability to perform model selection between these two priors. This figure again shows the Bayes factors between a set of GW spectrum models and noise. We see that both of the astrophysical models outperform the uniform amplitude prior, and that the Model A and Model B priors are distinguishable from each other, with Model A preferred for large amplitude signals, and Model B preferred for small amplitude signals.

C. Parameter Estimation

In addition to selecting between models, we are also interested in determining the accuracy with which the GW spectrum can be recovered. As a first method for visualizing our ability to constrain the shape of the GW background, in Figure 4 we show density plots of the recovered spectra calculated by simulating a noise-only signal, and recovering with the full GW spectral model. This is the type of analysis that leads to an upper limit on the GW amplitude, as there is no detection. To generate this plot, we calculate the spectrum at a sample of points in the Markov chain, then make a histogram of this spectrum at every frequency. The color map shows

the weight of the histogram in each bin, with red corresponding to higher weight, and blue corresponding to lower weight. Also shown is the white noise level (as a white line), and the 5%, 50%, and 95% confidence regions for the shape of the spectra. The purely GW-driven signal model with a uniform amplitude prior, labeled “flat in A , no bend” in Figure 4, yields a 95% upper limit on the amplitude of $A_{95} < 7.3 \times 10^{-16}$. This number can be directly compared to the published upper limits from real PTAs, since these analyses also assume a purely GW-driven evolution. The current upper limits are: the Parkes Pulsar Timing Array - $A < 2.4 \times 10^{-15}$; the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) - $A < 7 \times 10^{-15}$; European Pulsar Timing Array - $A < 6 \times 10^{-15}$. Thus, our simulated array is roughly three times more sensitive than the reported state-of-the-art, but based on projections [15], we expect our simulated sensitivity to be reached in reality within the next two to three years. The upper limits on the amplitude from the models in Figure 4 that allow for a bend in the spectrum are, unsurprisingly, significantly weaker. For example, the model with a constant amplitude prior and a bend in the spectrum leads to an upper bound of $A_{95} < 6.8 \times 10^{-15}$ - a full order of magnitude worse than the limit for the purely GW driven model.

Figure 5 shows posterior distributions for the recovered power spectra, but this time for simulated data includes an astrophysically-driven GW background. The parameters for this background were $A = 4.0 \times 10^{-15}$, $f_b = 3 \times 10^{-8}$ Hz, and $\kappa = 10/3$. In this case, the simulated signal is plotted along with the white noise. It is again clear that the different choices in amplitude prior have a strong effect on the recovered spectrum. In this case, we have also included an example with a prior that is uniform in the logarithm of the amplitude. This choice of prior has been shown [54, 55] to be a poor choice for setting upper limits, but we suggest that it is a superior choice for parameter estimation.

Note that the recovered spectra in Figure 5 do not fall precisely along the simulated track. This is an effect of the marginalization that occurs when calculating the densities. We are not showing the shape of any particular spectrum that was recovered by the chain. Rather, we are showing the distribution of the spectrum at each frequency, which is not required to have the same shape as the underlying spectral model. It is, essentially, a projection of the joint posterior distribution of the signal parameters A , f_b , and κ . We do not, therefore, expect the median recovered spectrum to lie directly on the simulated spectrum.

We will now discuss in detail the recovery of A , f_b , and κ . To examine our ability to measure the values of these parameters, Figure 6 shows the prior and posterior distributions for the signal parameters, for the same simulated signal as was used to generate Figure 5, and for four choices of amplitude prior (uniform in A , uniform in $\ln A$, and the two astrophysical priors, Models A, B).

² The best approach would be to marginalize over the model dimension along with the model parameters using a Reversible Jump Markov Chain Monte Carlo algorithm so that the optimal model is selected by the data and not guessed in advance.




FIG. 4: Density plots of the recovered spectra from MCMC runs including the full GW signal model (with A , f_b , and κ all free), for three different choices of amplitude prior. The simulated data in this case contained only noise. The priors used are, clockwise from top left, the Model A prior, the Model B prior, and the prior uniform in A . Also shown are the simulated white noise (white line) and the 5%, 50%, and 95% confidence regions for the recovered spectra.

Along the diagonal, we show the one-dimensional prior and posterior distributions for each parameter. In the off-diagonals, we show two-dimensional posterior distributions, illustrating the correlations between the parameters. The biases that arise from a uniform amplitude prior are clear in this figure - in particular, the amplitude, A , is pulled to very high values by this choice in prior. While this can be good for setting conservative upper limits on A given the lack of a detection, it is clearly a poor choice for parameter estimation. In contrast, the prior uniform in $\ln A$, while bad for setting upper limits, is shown here to be a good choice for parameter estimation.

The maximum a posteriori (MAP) parameter values, as well as the 90% credible intervals, for A , f_b , and κ are shown in Table II for all four choices of amplitude prior. From this table we see that only the Model A prior and

the flat in $\ln A$ prior return credible intervals that contain the simulated value for A . Our model selection studies show that it would be very clear given a signal of this type that the Model A prior is favored, and consequently, is the model that should be used for parameter estimation in this case.

D. Astrophysical Inference

The measurement of κ is of particular interest, as it maps most easily to astrophysical predictions. There are two approaches we can take to use the recovered κ to distinguish between astrophysical models. The first is illustrated in Figure 6; we can simply produce posteriors of κ given the detection of a GW background, and assess whether the measured value is consistent with a




FIG. 5: Density plots of the recovered spectra from MCMC runs including the full GW signal model (with A , f_b , and κ all free), for four different choices of amplitude prior. The simulated data in this case contained a GW background of type IV.c. The priors used are, clockwise from top left, the Model A prior, the Model B prior, a prior uniform in $\ln A$, and a prior uniform in A . Also shown are the simulated white noise (white line), the simulated signal (gray line), and the 5%, 50%, and 95% confidence regions for the recovered spectra. The Model A prior and the prior uniform in $\ln A$ result in the best fit to the simulated spectrum.

prior	$A_{5\%}$	A_{MAP}	$A_{95\%}$	$(f_b)_{5\%}$	$(f_b)_{\text{MAP}}$ (Hz)	$(f_b)_{95\%}$	$\kappa_{5\%}$	κ_{MAP}	$\kappa_{95\%}$
Model A	2.05×10^{-15}	3.4×10^{-15}	7.10×10^{-15}	5.16×10^{-9}	9.4×10^{-9}	3.29×10^{-8}	1.65	3.3	5.17
Model B	1.35×10^{-15}	1.9×10^{-15}	2.83×10^{-15}	3.64×10^{-9}	5.0×10^{-9}	8.13×10^{-9}	2.6	5.0	7.01
flat in A	1.61×10^{-14}	5.2×10^{-14}	6.45×10^{-14}	4.53×10^{-8}	9.5×10^{-8}	9.83×10^{-8}	2.11	2.6	3.03
flat in $\ln A$	2.04×10^{-15}	8.8×10^{-15}	2.56×10^{-14}	6.16×10^{-9}	2.8×10^{-8}	7.96×10^{-8}	1.80	2.5	3.80

TABLE II: The maximum *a posteriori* values for A , f_b , and κ , along with the 90% credible intervals for each parameter, corresponding to the simulated signals and inferred distributions shown in Figure 6. Recall that the simulated values of these parameters were $A = 4.0 \times 10^{-15}$, $f_b = 1 \times 10^{-8}$ Hz, and $\kappa = 10/3$. Note that only the Model A and uniform in $\ln A$ cases include the simulated values within their 90% confidence intervals.

particular model. This presents a challenge, however, as the value of κ is not measured with high precision. That is, the posterior distributions of κ for both astrophysical amplitude priors and for the prior uniform in $\ln A$ have significant weight throughout the prior range.

Additionally, the shape and peak of these posteriors depend strongly upon the choice of prior on A . This can, however, be viewed more as a feature than a drawback, as the choice of prior on A encodes astrophysical information. This implies that the correct method for

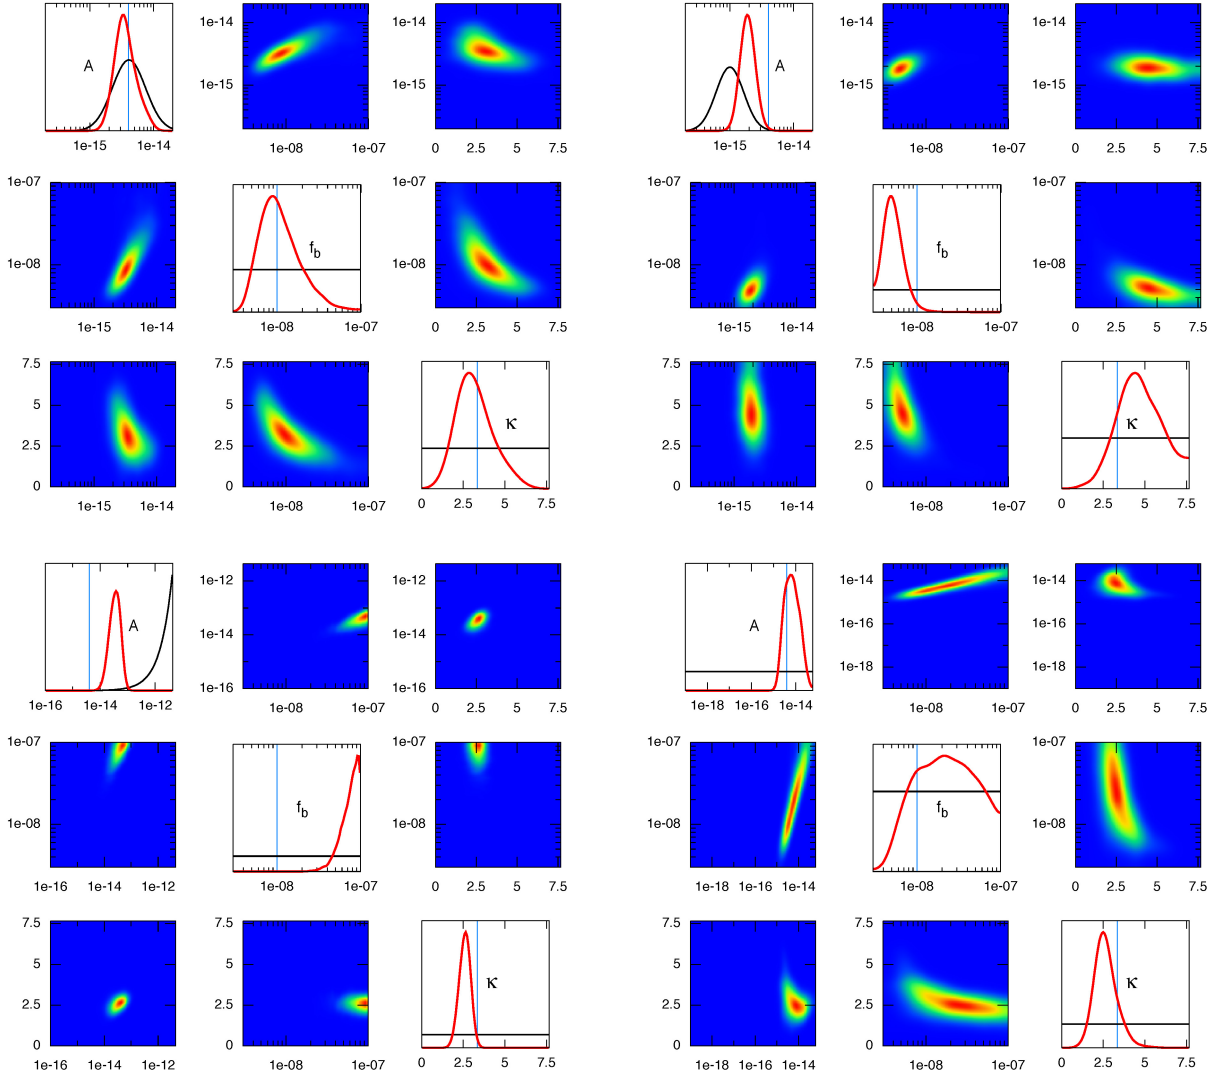


FIG. 6: Prior and posterior distributions for the three GW parameters, A , f_b , and κ , recovered using four different amplitude priors. Along the diagonal are plotted the one-dimensional prior and posterior distributions for (from top) A , f_b , and κ . The off-diagonals show the 2-d posteriors. The priors used are, clockwise from top left: Model A, Model B, uniform in $\ln A$, and uniform in A . The simulated signal had $A = 4.0 \times 10^{-15}$, $f_b = 10^{-8}$ Hz, and $\kappa = 10/3$. It is clear the prior that is uniform in A leads to strong biases in the recovered parameter values. Thus, although this prior is a good choice for producing conservative upper limits, it has limited utility in parameter estimation studies.

drawing astrophysical inferences from the measurement of the stochastic background is through model selection, where the models are defined both by κ and by the prior on A .

To investigate this type of model selection study, we simulated data of types II and IV, and recovered these simulated data sets using templates with f_b free, but with fixed values for κ , using the amplitude prior for Model A. We then calculated the Bayes factors between these two models. The spectral models used for recovery have fixed slopes $\kappa = 10/3$, $\kappa = 7/3$ or $\kappa = 29/6$. The simulated data had a spectrum with $\kappa = 10/3$, which corresponds to stellar slingshot hardening. The $\kappa = 7/3$ model corresponds to hardening due to a thin, cold circumbinary gas

disk, while the $\kappa = 29/6$ was randomly selected to provide a model with a shallower spectrum than was used to generate the data. Figure 7 shows the log Bayes factor of the $\kappa = 10/3$ model relative to the $\kappa = 29/6$ (red) and $\kappa = 7/3$ (blue). So long as the bend frequency is high enough (signal type IV), and the amplitude large enough (sub-types b,c), it is possible to distinguish between the different hardening mechanisms.

E. Learning from a Non-detection

Information is gained from any measurement that leads to a posterior distribution that is different from the prior




FIG. 7: The Bayes factors between a GW model with free parameters A and f_b and fixed $\kappa = 10/3$, and models with fixed $\kappa = 29/6$ (red) and $\kappa = 7/3$ (blue). The simulated signals (described in Table (I)) all had $\kappa = 10/3$. The amplitude prior used here was the Model A prior. The signals are arranged along the x-axis in order of ascending detectability. A Bayes factor larger than unity indicates a preference for the $\kappa = 10/3$ model. For the most detectable signals, there is a clear preference for the correct model.

distribution. In Figure 8, we show the prior and posterior distributions of A , f_b , and κ , as above, but now recovered from running on a signal that contains only noise. These distributions are shown for three different choices of amplitude prior - Model A, Model B, and uniform in A . It is clear that these distributions are significantly altered from the priors, showing that we can learn about astrophysical model space even lacking a GW detection.

The primary result from PTA searches for a stochastic GW background is typically quoted as an upper limit on the amplitude A . Unsurprisingly, for the signal analyzed in Figure 8, the 95% upper bound on A depends on the choice of prior. For Model A, $A_{95\%} = 9.95 \times 10^{-15}$; for Model B, $A_{95\%} = 1.99 \times 10^{-15}$; and finally, for a uniform prior in A , $A_{95\%} = 6.83 \times 10^{-15}$. We see, then, that Model A leads to the most conservative upper bounds. Note that these bounds are for spectral models that allow for a bend in the spectrum, the limits assuming purely GW driven evolution are almost an order of magnitude lower. We will now explore precisely how much information we gain by a non-detection for the two astrophysical models.

To quantify exactly how much we can learn in this case, we compute the information gain (in bits) as the Kullback-Leibler (KL) divergence [56, 57] between the posterior and prior distributions. If the posterior matches the prior, we have learned nothing from the data. The larger the difference, the more we have learned. The KL divergence between the posterior, $p(\vec{x}|d)$, and the prior, $p(\vec{x})$, is given by

$$\text{KL}(p(\vec{x}|d)||p(\vec{x})) = \int p(\vec{x}|d) \log \left(\frac{p(\vec{x}|d)}{p(\vec{x})} \right) d\vec{x}. \quad (27)$$

This quantity can be calculated via thermodynamic integration, and is equal to

$$\text{KL}(p(\vec{x}|d)||p(\vec{x})) = E_{\beta=1} [\ln p(d|\vec{x})] - \ln p(d), \quad (28)$$

which is the expectation value of the log likelihood minus the log evidence. For the case illustrated in Figure 8, the KL divergences between posterior and prior are $\text{KL}_{\text{Model A}} = 308.1 \pm 0.4$ bits and $\text{KL}_{\text{Model B}} = 304.6 \pm 0.4$ bits. These numbers encode the information we have learned about both the signal and the *noise* parameters. One way to get a rough estimate of what we learned about the signal model is to compare the information gain for the signal and noise model to the information gain for the noise model alone, which comes out to $\text{KL}_{\text{Noise Model}} = 301.2 \pm 0.4$ bits, suggesting an information gain of ~ 7 bits for Model A and ~ 3 bits for Model B. To properly isolate the information gained about the astrophysical models, we need to perform the integral in Eq. (27) over only the prior and posterior of the signal parameters A , f_b , and κ . The functional form of the prior distributions is known, so this presents no difficulty. For the posterior distributions, we perform a three-dimensional KDE smoothing of the posterior distributions derived from our Markov chains. The KDE is applied to chain samples that are mirrored at the prior boundaries to reduce edge effects. We then use these smoothed distributions to numerically integrate Eq. (27). We find that, for Model A the information gain for the spectral model is $\text{KL} = 1.5 \pm 0.08$ bits, while for Model B the information gain is $\text{KL} = 0.7 \pm 0.04$ bits. These numbers are significantly smaller than the crude estimate obtained by taking the difference between the noise and signal models, but this is not surprising - we know that there are significant correlations between the noise parameters and the signal parameters, and the information measure is not additive. While these information gains are not large, they show that the posterior and prior distributions are measurably different, and that we begin to learn about the astrophysical models driving the SMBH mergers even before a detection is made. As a point of comparison, the information gain for Model A is equal to the information gained about cosmological models in going from the 7 year Wilkinson Microwave Anisotropy Probe (WMAP) data set to the 9 year data set, but is significantly less than the 30 bits gained in going from the 9 year WMAP maps to the higher resolution Planck maps [58].

As the signal-to-noise ratio grows, the information gain grows. For the strong signal examples shown in Figure 6, the information gains are $\text{KL} = 3.1 \pm 0.2$ bits for Model A and $\text{KL} = 4.8 \pm 0.2$ bits for Model B. In this instance we learn more about Model B since the amplitude of the simulated signal is large compared to what is predicted by the model, so there is a greater difference between the prior and posterior distributions.

An alternative way of seeing that we learn something about the astrophysical models from a non-detection is to

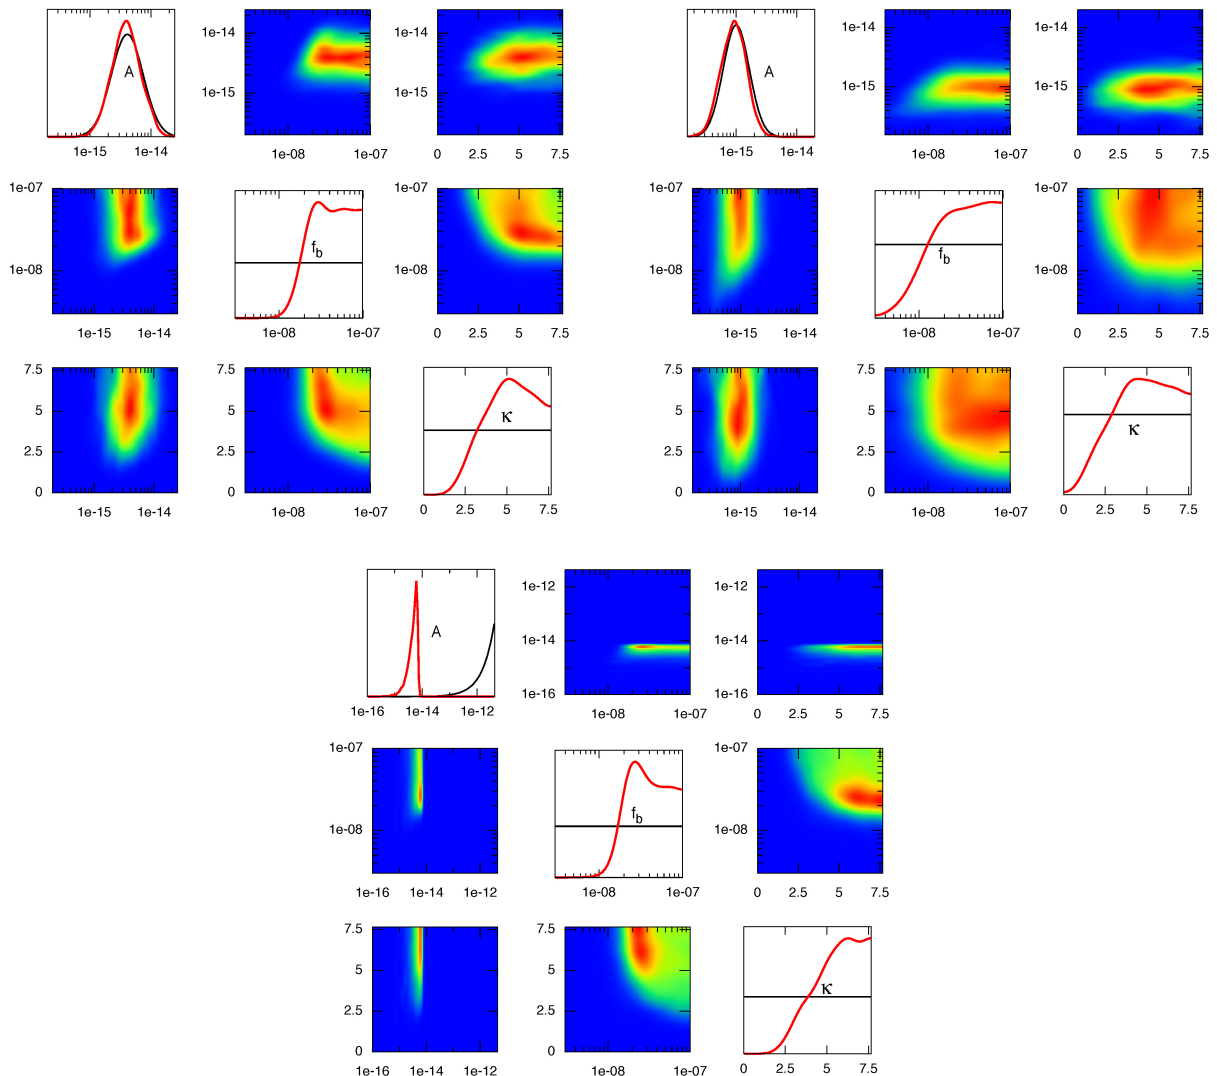


FIG. 8: Prior and posterior distributions for the three GW parameters, A , f_b , and κ , recovered by analyzing simulated data that contains only white noise. Along the diagonal are plotted the one-dimensional prior and posterior distributions for (from top) A , f_b , and κ . The off-diagonals show the 2-d posteriors for all combinations of these three parameters. The amplitude priors are, clockwise from top left: Model A, Model B, and uniform in A . Note that, even in the absence of a GW signal, the posterior distributions are substantially different from the priors. This indicates that we can learn about astrophysical models even if no detection is made.

compare the evidence for models that include a bend in the spectrum to those that assume a purely GW driven evolution. For the noise-only data sets used to generate Figure 8, the log Bayes factor in favor of there being a bend in the spectrum is $\ln \text{BF} = 18.7 \pm 0.5$ for the Model A amplitude prior and $\ln \text{BF} = 3.5 \pm 0.5$ for the Model B amplitude prior. These results say that a non-detection of GWs by a PTA with the sensitivity of our simulated array would rule out purely GW driven evolution of these merger models. Scaling back the sensitivity of the simulated array by a factor of two, (increasing the timing noise from 200 ns to 400 ns) to get something closer to the NANOGrav sensitivity in 2015, yields $\ln \text{BF} = 5.6 \pm 0.5$ for the Model A amplitude prior and $\ln \text{BF} = 1.4 \pm 0.4$ for

the Model B amplitude prior. At this lower sensitivity there would still be strong evidence for non-GW driven evolution for Model A, but not for Model B.

IV. SUMMARY

Pulsar Timing Arrays are likely to detect a stochastic GW background from binary SMBHBs before the end of the decade [15]. The astrophysical processes that drive SMBH binaries toward merger are not fully understood, and are largely unconstrained by observations. Processes such as stellar scattering can drive the binaries through the sensitive band of PTAs more quickly than GWs, lead-

ing to a diminution of the GW signal. This may make the detection of GWs more challenging, but also opens up a new avenue for learning about the astrophysics of SMBH mergers through measurement of the spectral shape.

We have shown that a simple model for the spectrum, described in Eqs. (8), (13), are useful for the detection of GW backgrounds that are generated by SMBH binaries which are driven by more than one type of mechanism within the pulsar timing band. This model can also be used in parameter estimation studies to characterize this GW background.

We found that the choice of prior on the amplitude can significantly impact parameter estimation, and that the commonly used uniform prior in amplitude leads to especially large biases. A prior that is uniform in the logarithm of the amplitude was found to be a far better choice.

We have shown that pulsar timing observations can be used to distinguish between models that are charac-

terized by different priors on the amplitude of the GW background. In future work, we plan to extend this study to more detailed models that make predictions about the slope parameters and the bend frequencies.

Finally, we have illustrated that information is gained about astrophysical models, even when no detection is made.

Acknowledgments

We have benefited from informative discussions with Justin Ellis and Xavier Siemens, and useful feedback from members of the NANOGrav collaboration. In particular, Chiara Mingarelli provided helpful insight into our theoretical calculations. LS was supported by Nicolás Yunes' NSF CAREER Award PHY-1250636. NJC was supported by NSF Award PHY-1306702

-
- [1] J. Kormendy and L. C. Ho, *Ann. Rev. Astron. Astrophys.* **51**, 511 (2013), 1304.7762.
 - [2] M. Eracleous, T. A. Boroson, J. P. Halpern, and J. Liu, *ArXiv e-prints* (2011), 1106.2952.
 - [3] P. Tsalmantza, R. Decarli, M. Dotti, and D. W. Hogg, *ApJ* **738**, 20 (2011), 1106.1180.
 - [4] M. A. McLaughlin, *Class. Quant. Grav.* **30**, 224008 (2013), 1310.0758.
 - [5] R. D. Ferdman, R. van Haasteren, C. G. Bassa, M. Burgay, I. Cognard, A. Corongiu, N. D'Amico, G. Desvignes, J. W. T. Hessels, G. H. Janssen, et al., *Class. Quant. Grav.* **27**, 084014 (2010), 1003.3405.
 - [6] R. N. Manchester, G. Hobbs, M. Bailes, W. A. Coles, W. van Straten, M. J. Keith, R. M. Shannon, N. D. R. Bhat, A. Brown, S. G. Burke-Spolaor, et al., *PASA* **30**, 17 (2013), 1210.6130.
 - [7] M. McLaughlin, *Gen. Rel. Grav.* **46**, 11 (2014).
 - [8] D. R. Lorimer and M. Kramer, *Handbook of pulsar astronomy*, vol. 4 (Cambridge University Press, 2005).
 - [9] G. Hobbs, W. Coles, R. N. Manchester, M. J. Keith, R. M. Shannon, D. Chen, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, D. Champion, et al., *MNRAS* **427**, 2780 (2012).
 - [10] D. J. Champion, G. B. Hobbs, R. N. Manchester, R. T. Edwards, D. C. Backer, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, W. Coles, P. B. Demorest, et al., *ApJ* **720**, L201 (2010), 1008.3607.
 - [11] R. W. Hellings and G. S. Downs, *ApJ* **265**, L39 (1983).
 - [12] Z. Arzumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, J. M. Cordes, P. B. Demorest, X. Deng, T. Dolch, J. A. Ellis, et al., *ApJ* **794**, 141 (2014), 1404.1267.
 - [13] X.-J. Zhu, G. Hobbs, L. Wen, W. A. Coles, J.-B. Wang, R. M. Shannon, R. N. Manchester, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, et al., *MNRAS* **444**, 3709 (2014).
 - [14] R. van Haasteren, Y. Levin, G. H. Janssen, K. Lazaridis, M. Kramer, B. W. Stappers, G. Desvignes, M. B. Purver, A. G. Lyne, R. D. Ferdman, et al., *MNRAS* **414**, 3117 (2011), 1103.0576.
 - [15] X. Siemens, J. Ellis, F. Jenet, and J. D. Romano, *Class. Quant. Grav.* **30**, 224015 (2013), 1305.3196.
 - [16] M. Dotti, A. Sesana, and R. Decarli, *Advances in Astronomy* **2012**, 940568 (2012), 1111.0664.
 - [17] A. Sesana, F. Haardt, and P. Madau, *ApJ* **651**, 392 (2006), astro-ph/0604299.
 - [18] P. J. Armitage and P. Natarajan, *ApJ* **634**, 921 (2005), astro-ph/0508493.
 - [19] P. Ivanov, J. Papaloizou, and A. Polnarev, *MNRAS* **307**, 79 (1999), astro-ph/9812198.
 - [20] F. M. Khan, A. Just, and D. Merritt, *ApJ* **732**, 89 (2011), 1103.0272.
 - [21] G. D. Quinlan, *New Astron.* **1**, 35 (1996), astro-ph/9601092.
 - [22] V. Ravi, J. S. B. Wyithe, R. M. Shannon, and G. Hobbs, *ArXiv e-prints* (2014), 1406.5297.
 - [23] A. Toomre and J. Toomre, *Astrophys. J.* **178**, 623 (1972).
 - [24] A. Toomre, *Annual Review of Astronomy and Astrophysics* **15**, 437 (1977).
 - [25] J. E. Barnes and L. Hernquist, *Annual Review of Astronomy and Astrophysics* **30**, 705 (1992).
 - [26] P. F. Hopkins, G. T. Richards, and L. Hernquist, *The Astrophysical Journal* **654**, 731 (2007), URL <http://stacks.iop.org/0004-637X/654/i=2/a=731>.
 - [27] M. Dotti, A. Merloni, and C. Montuori, *Mon. Not. Roy. Astron. Soc.* **448**, 3603 (2015), 1502.03101.
 - [28] M. Milosavljević and D. Merritt, in *The Astrophysics of Gravitational Wave Sources*, edited by J. M. Centrella (2003), vol. 686 of *American Institute of Physics Conference Series*, pp. 201–210, astro-ph/0212270.
 - [29] D. Merritt and M. Milosavljević, **8** (2005), <http://www.livingreviews.org/lrr-2005-8>, arXiv:astro-ph/0410364.
 - [30] A. Sesana, F. Haardt, P. Madau, and M. Volonteri, *Astrophys. J.* **611**, 623 (2004), astro-ph/0401543.
 - [31] B. Kocsis and A. Sesana, *MNRAS* **411**, 1467 (2011), 1002.0584.
 - [32] S. T. McWilliams, J. P. Ostriker, and F. Pretorius, *ArXiv*

- e-prints (2012), 1211.4590.
- [33] S. T. McWilliams, J. P. Ostriker, and F. Pretorius, *ApJ* **789**, 156 (2014), 1211.5377.
- [34] A. Sesana, *MNRAS* **433**, L1 (2013), 1211.5375.
- [35] R. Shannon, V. Ravi, W. Coles, G. Hobbs, M. Keith, et al., *Science* **342**, 334 (2013), 1310.4569.
- [36] V. Ravi, J. S. B. Wyithe, R. M. Shannon, G. Hobbs, and R. N. Manchester, *MNRAS* **442**, 56 (2014), 1404.5183.
- [37] E. Phinney (2001), astro-ph/0108028.
- [38] A. Sesana, A. Vecchio, and C. N. Colacino, *MNRAS* **390**, 192 (2008), 0804.4476.
- [39] K. S. Thorne, in *300 Years of Gravitation* (Cambridge University Press, Cambridge, 1987).
- [40] D. W. Hogg, ArXiv e-prints (1999), arXiv:astro-ph/9905116v4, 9905116v4.
- [41] S. Chandrasekhar, *ApJ* **97**, 255 (1943).
- [42] F. M. Khan, K. Holley-Bockelmann, P. Berczik, and A. Just, *ApJ* **773**, 100 (2013), 1302.1871.
- [43] A. Sesana, *Class. Quant. Grav.* **30**, 224014 (2013), 1307.2600.
- [44] A. Sesana, *Astrophys. Space Sci. Proc.* **40**, 147 (2015), 1407.5693.
- [45] A. Sesana, *Class. Quant. Grav.* **30**, 244009 (2013), 1307.4086.
- [46] A. H. Jaffe and D. C. Backer, *ApJ* **583**, 616 (2003), astro-ph/0210148.
- [47] A. Sesana, *ApJ* **719**, 851 (2010), 1006.0730.
- [48] M. Enoki and M. Nagashima, *Prog. Theor. Phys.* **117**, 241 (2007), astro-ph/0609377.
- [49] N. J. Cornish (2012), 1209.6428.
- [50] T. B. Littenberg and N. J. Cornish, *Phys. Rev. D* **80**, 063007 (2009).
- [51] N. J. Cornish and T. B. Littenberg (2014), 1410.3835.
- [52] T. B. Littenberg and N. J. Cornish (2014), 1410.3852.
- [53] R. Storn and K. Price, *Journal of Global Optimization* **11**, 341 (1997), ISSN 0925-5001.
- [54] S. Taylor, J. Ellis, and J. Gair, *Phys. Rev. D* **90**, 104028 (2014), 1406.5224.
- [55] J. A. Ellis, *Class. Quant. Grav.* **30**, 224004 (2013), 1305.0835.
- [56] G. Lefebvre, R. Steele, and A. C. Vandal, *Comp. Stat. Data Anal.* **54**, 1719 (2010).
- [57] S. Vitoratou and I. Ntzoufras, ArXiv e-prints (2013), 1308.6753.
- [58] S. Seehars, A. Amara, A. Refregier, A. Paranjape, and J. Akeret, *Phys. Rev. D* **90**, 023533 (2014), 1402.3593.