

University of Texas Rio Grande Valley

ScholarWorks @ UTRGV

Physics and Astronomy Faculty Publications
and Presentations

College of Sciences

4-10-2016

THE NANOGRV NINE-YEAR DATA SET: LIMITS ON THE ISOTROPIC STOCHASTIC GRAVITATIONAL WAVE BACKGROUND

Z. Arzoumanian

A. Brazier

S. Burke-Spolaor

S. J. Chamberlin

S. Chatterjee

See next page for additional authors

Follow this and additional works at: https://scholarworks.utrgv.edu/pa_fac



Part of the [Astrophysics and Astronomy Commons](#)

Recommended Citation

Z. Arzoumanian, et. al., (2016) THE NANOGRV NINE-YEAR DATA SET: LIMITS ON THE ISOTROPIC STOCHASTIC GRAVITATIONAL WAVE BACKGROUND. *Astrophysical Journal* 821:1. DOI: <http://doi.org/10.3847/0004-637X/821/1/13>

This Article is brought to you for free and open access by the College of Sciences at ScholarWorks @ UTRGV. It has been accepted for inclusion in Physics and Astronomy Faculty Publications and Presentations by an authorized administrator of ScholarWorks @ UTRGV. For more information, please contact justin.white@utrgv.edu, william.flores01@utrgv.edu.

Authors

Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, B. Christy, J. M. Cordes, N. J. Cornish, K. Crowter, P. B. Demorest, X. Deng, T. Dolch, J. A. Ellis, R. D. Ferdman, E. Fonseca, N. Garver-Daniels, M. E. Gonzalez, F. Jenet, G. Jones, M. L. Jones, V. M. Kaspi, M. Koop, M. T. Lam, T. J.W. Lazio, L. Levin, A. N. Lommen, D. R. Lorimer, J. Luo, R. S. Lynch, and D. R. Madison



THE NANOGRV NINE-YEAR DATA SET: LIMITS ON THE ISOTROPIC STOCHASTIC GRAVITATIONAL WAVE BACKGROUND

Z. ARZOUMANIAN¹, A. BRAZIER², S. BURKE-SPOLAOR^{3,28}, S. J. CHAMBERLIN⁴, S. CHATTERJEE², B. CHRISTY⁵, J. M. CORDES², N. J. CORNISH⁶, K. CROWTER⁷, P. B. DEMOREST³, X. DENG⁴, T. DOLCH^{2,8}, J. A. ELLIS^{9,29}, R. D. FERDMAN¹⁰, E. FONSECA⁷, N. GARVER-DANIELS¹¹, M. E. GONZALEZ^{7,12}, F. JENET¹³, G. JONES¹⁴, M. L. JONES¹¹, V. M. KASPI¹⁰, M. KOOP⁴, M. T. LAM², T. J. W. LAZIO⁹, L. LEVIN¹¹, A. N. LOMMEN⁵, D. R. LORIMER¹¹, J. LUO¹³, R. S. LYNCH¹⁵, D. R. MADISON^{2,16,28}, M. A. MCLAUGHLIN¹¹, S. T. MCWILLIAMS¹¹, C. M. F. MINGARELLI^{17,18,30}, D. J. NICE¹⁹, N. PALLIYAGURU¹¹, T. T. PENNUCCI²⁰, S. M. RANSOM¹⁶, L. SAMPSON⁶, S. A. SANIDAS^{21,22}, A. SESANA²³, X. SIEMENS²⁴, J. SIMON²⁴, I. H. STAIRS⁷, D. R. STINEBRING²⁵, K. STOVALL²⁶, J. SWIGGUM¹¹, S. R. TAYLOR⁹, M. VALLISNERI⁹, R. VAN HAASTEREN^{9,29}, Y. WANG²⁷, AND W. W. ZHU^{7,18}

(THE NANOGRV COLLABORATION)

¹ Center for Research and Exploration in Space Science and Technology and X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA

² Department of Astronomy, Cornell University, Ithaca, NY 14853, USA

³ National Radio Astronomy Observatory, 1003 Lopezville Road, Socorro, NM 87801, USA

⁴ Department of Astronomy and Astrophysics, Pennsylvania State University, University Park, PA 16802, USA

⁵ Department of Physics and Astronomy, Franklin & Marshall College, P.O. Box 3003, Lancaster, PA 17604, USA

⁶ Department of Physics, Montana State University, Bozeman, MT 59717, USA

⁷ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada

⁸ Department of Physics, Hillsdale College, 33 E. College Street, Hillsdale, MI 49242, USA

⁹ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA; Justin.A.Ellis@jpl.nasa.gov

¹⁰ Department of Physics, McGill University, 3600 University Street, Montreal, QC H3A 2T8, Canada

¹¹ Department of Physics and Astronomy, West Virginia University, P.O. Box 6315, Morgantown, WV 26506, USA

¹² Department of Nuclear Medicine, Vancouver Coastal Health Authority, Vancouver, BC V5Z 1M9, Canada

¹³ Center for Gravitational Wave Astronomy, University of Texas at Brownsville, Brownsville, TX 78520, USA

¹⁴ Department of Physics, Columbia University, NY 10027, USA

¹⁵ National Radio Astronomy Observatory, P.O. Box 2, Green Bank, WV 24944, USA

¹⁶ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA

¹⁷ TAPIR, MC 350-17, California Institute of Technology, Pasadena, CA 91125, USA

¹⁸ Max Planck Institute for Radio Astronomy, Auf dem Hügel 69, D-53121 Bonn, Germany

¹⁹ Department of Physics, Lafayette College, Easton, PA 18042, USA

²⁰ University of Virginia, Department of Astronomy, P.O. Box 400325, Charlottesville, VA 22904, USA

²¹ Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH, Amsterdam, The Netherlands

²² Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester, M13 9PL, UK

²³ School of Physics & Astronomy, University of Birmingham, Birmingham, B15 2TT, UK

²⁴ Center for Gravitation, Cosmology and Astrophysics, Department of Physics, University of Wisconsin-Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA

²⁵ Department of Physics and Astronomy, Oberlin College, Oberlin, OH 44074, USA

²⁶ Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM 87131, USA

²⁷ School of Physics, Huazhong University of Science and Technology, 1037 Luoyu Road, Wuhan, Hubei Province 430074, China

Received 2015 October 5; accepted 2016 February 18; published 2016 April 4

ABSTRACT

We compute upper limits on the nanohertz-frequency isotropic stochastic gravitational wave background (GWB) using the 9 year data set from the North American Nanohertz Observatory for Gravitational Waves (NANOGrav) collaboration. Well-tested Bayesian techniques are used to set upper limits on the dimensionless strain amplitude (at a frequency of 1 yr^{-1}) for a GWB from supermassive black hole binaries of $A_{\text{gw}} < 1.5 \times 10^{-15}$. We also parameterize the GWB spectrum with a broken power-law model by placing priors on the strain amplitude derived from simulations of Sesana and McWilliams et al. Using Bayesian model selection we find that the data favor a broken power law to a pure power law with odds ratios of 2.2 and 22 to one for the Sesana and McWilliams prior models, respectively. Using the broken power-law analysis we construct posterior distributions on environmental factors that drive the binary to the GW-driven regime including the stellar mass density for stellar-scattering, mass accretion rate for circumbinary disk interaction, and orbital eccentricity for eccentric binaries, marking the first time that the *shape* of the GWB spectrum has been used to make astrophysical inferences. Returning to a power-law model, we place stringent limits on the energy density of relic GWs, $\Omega_{\text{gw}}(f)h^2 < 4.2 \times 10^{-10}$. Our limit on the cosmic string GWB, $\Omega_{\text{gw}}(f)h^2 < 2.2 \times 10^{-10}$, translates to a conservative limit on the cosmic string tension with $G\mu < 3.3 \times 10^{-8}$, a factor of four better than the joint Planck and high- l cosmic microwave background data from other experiments.

Key words: gravitational waves – methods: data analysis – pulsars: general

²⁸ Jansky Fellow.

²⁹ Einstein Fellow.

³⁰ Marie Curie Fellow.

1. INTRODUCTION

Pulsars are remarkable astrophysical laboratories for Einstein’s theory of gravitation: the outstanding stability of their periodic emission made it possible to detect subtle general-relativistic effects in the conservative dynamics of close binary systems (see, e.g., Kramer et al. 2006), and to develop conclusive evidence that binaries lose energy to gravitational waves (GWs) in accord with the general-relativistic quadrupole formula (Weisberg et al. 2010). Moreover, the systematic timing of large arrays of millisecond pulsars has become the dominant experimental technique for the direct detection of low-frequency (\sim nHz) GWs (see, e.g., Burke-Spolaor 2015).

In this paper, we analyze the NANOGrav 9 year pulsar-timing data set (Arzoumanian et al. 2015, hereafter NG9) to set new upper limits on isotropic stochastic GW backgrounds (GWBs) with power-law spectra, such as those generated by inflation-amplified quantum fluctuations in the early universe, by cosmic-string networks, or (most notably) by a population of supermassive black hole binaries (SMBHBs) on circular orbits and evolving due to GW radiation alone. We then interpret our results in terms of cosmological variables and of the properties of source SMBHB populations.

Moving beyond the power-law approximation, we examine what inferences can be drawn from the data about stochastic GWBs with broken-power-law spectra that become less steep at lower frequencies. Such spectra would arise if SMBHBs enter the nanohertz GW band with significant eccentricity, or if they are strongly affected by stellar scattering or mass accretion (Sesana 2013b). Our data are not sufficiently informative to constrain the amplitude and shape of a GWB simultaneously in the absence of other information, but we can use them to investigate the following question. Under astrophysical assumptions that lead to pure power-law spectra, current theoretical models of SMBHB populations predict GWB levels that are marginally (or significantly) too high compared to newer upper limits (Shannon et al. 2015; this article). If we take these models as priors, do our data then favor broken-power-law spectra? Furthermore, how strong and how plausible are the eccentricity and environmental conditions needed to generate such spectra?

In this work we strive to always be transparent in our use of different models and their various assumptions. Furthermore, we refrain from making unqualified statements about the nature of the stochastic GWB and our ability to detect it. In Shannon et al. (2015), the authors suggest that GWs are “missing,” implying that a detection of the stochastic GWB was expected based on current models. However, those models, *and the corresponding data analysis*, assume that the SMBHB evolves without stalling down to the pulsar timing array (PTA) frequency band, and that interactions with stars and gas or binary eccentricity do not reduce the GWB signal as the lowest, most sensitive frequencies. The possibility of such interactions and the uncertainty in the underlying black hole–host relations make a non-detection an unsurprising result.

Performing these analyses required novel theoretical developments in the treatment of pulsar-timing noise and (especially) SMBHB astrophysics. We describe these treatments in detail, since we expect them to be useful for future studies. As a result, this is a long paper. To help readers absorb the overall picture, and if desired concentrate on the sections that most interest them, we collect our main results in this introduction, and we lay out the plan for our exposition.

1.1. Gravitational Wave Detection with PTAs

Sazhin (1978) and Detweiler (1979) realized that GWs could manifest as otherwise unexplained residuals in the times of arrival (TOAs) of pulsar signals after subtracting a deterministic *timing model*. This timing model accounts for the intrinsic evolution of pulsar spin, radio frequency-dependent delays due to interstellar propagation effects, the astrometric time delays and advances due to the relative motion of the pulsar system with respect to the Earth (indeed, to the observatory), as well as orbital-kinematic and light-propagation effects for pulsars that orbit a binary companion (see, e.g., Lorimer & Kramer 2005). Foster & Backer (1990) pointed out that the timing residuals from an array of pulsars (a *pulsar timing array*, or PTA) could be analyzed coherently to separate GW-induced residuals, which have distinctive correlations among different pulsars (Hellings & Downs 1983), from other systematic effects, such as clock errors or delays due to light propagation through the interstellar medium.

Today, three international consortia (NANOGrav (McLaughlin 2013), the EPTA (Kramer & Champion 2013), and the PPTA (Hobbs 2013)) are conducting extensive campaigns to search for GWs by timing dozens of individual millisecond pulsars (MSPs), in which the best-timed have root-mean-squared (rms) residuals less than 100 ns (corresponding to GW strain sensitivities $\sim 10^{-15}$). The three PTAs collaborate and share data under the aegis of the International Pulsar Timing Array (Hobbs et al. 2010).

In order to robustly detect GWs, one must have a thorough understanding of the underlying noise in the pulsar timing data (see, e.g., Cordes 2013; Stinebring 2013, for a detailed review). Template matching errors due to radiometer noise are uncorrelated in both time and frequency, but pulse-jitter noise (Cordes & Shannon 2010) appears to affect all TOAs obtained simultaneously in different frequency channels. Timing noise with a red power spectrum occurs to varying degree in different pulsars. This includes spin noise (Shannon & Cordes 2010), which is achromatic and is much smaller in MSPs compared to objects with stronger magnetic fields and longer spin periods. Chromatic red noise due to propagation through intervening plasmas (ISM, interplanetary medium, and ionosphere) may also be present if dispersive delays are not removed perfectly or if scattering and refraction effects contribute significantly.

1.2. The Stochastic GWB from SMBHBs

GWs enter timing residuals through a time integral of gravitational strain. White TOA measurement uncertainties dominate the noise budget at the highest accessible frequencies in the PTA band, leading to an $\propto f^{3/2}$ scaling of the strain sensitivity. At the opposite end of the band, the strain sensitivity is inhibited at frequencies much less than the inverse observation time by the “quadratic wall,” where the fitting of (and marginalization over) quadratic-in-time functions describing the pulsar spindown leads to a strain sensitivity scaling of $\propto f^{-2}$ (Moore et al. 2015). Thus, a PTA’s GW-strain sensitivity is greatest at frequencies on the order of the inverse timespan of timing observations ($\sim \mathcal{O}$ (nHz)). This behavior has been confirmed both theoretically (Thrane & Romano 2013; Moore et al. 2015) and empirically (e.g., Yardley et al. 2010; Arzoumanian et al. 2014; Lentati et al. 2015; Shannon et al. 2015; Babak et al. 2016), where in the latter this behavior remains true even in the presence of red timing noise. The

strongest expected sources in this low-frequency band are SMBHBs with masses of 10^8 – $10^{10} M_\odot$, out to $z \simeq 1$ (Rajagopal & Romani 1995; Jaffe & Backer 2003; Wyithe & Loeb 2003). The binaries form after the hierarchical mergers (Sesana et al. 2004, 2008) of galaxies hosting individual SMBHBs (as most galaxies are thought to do, cf. Kormendy & Ho 2013). The most massive and nearest binaries may be detected individually by PTAs through their continuous GW emission. Moreover, the cosmic population of SMBHBs may be observed collectively as a stochastic GWB composed of the incoherent superposition of signals from the binaries. Rosado et al. (2015) discuss which detection is likely to come first.

The main focus of this paper is the measurement of an isotropic stochastic GWB from SMBHBs. Anisotropic-background searches based on the formalism and techniques developed by Gair et al. (2014), Taylor & Gair (2013), Mingarelli et al. (2013) are currently underway and will be the subject of a follow-up paper. The simplest characterization of the stochastic GWB—a power-law Gaussian process with isotropic inter-pulsar correlations—applies if:

1. all binaries are assumed to have circular orbits (so each component signal is instantaneously monochromatic);
2. all binaries evolve through the PTA band due purely to GW emission, as opposed to environmental effects such as interactions with nearby gas or with stars in the galactic nucleus;
3. all binaries are distributed isotropically across the sky in sufficient numbers to fulfill the central limit theorem at all frequencies.

Under these conditions, the observed timing residuals due to the GWB are described fully by the (cross-) power spectral density

$$S_{ab}(f) = \Gamma_{ab} \frac{A_{\text{gw}}^2}{12\pi^2} \left(\frac{f}{\text{yr}^{-1}} \right)^{-\gamma} \text{yr}^3, \quad (1)$$

where A_{gw} is the dimensionless strain amplitude, a and b range over the pulsars in the array, $\gamma = 13/3$ for a background composed of SMBHBs, and Γ_{ab} is the Hellings–Downs Hellings & Downs (1983) isotropic correlation coefficient for pulsars a and b (a function of the separation angle between their lines of sight, which is normalized so that $\Gamma_{aa} = 1$; see also Mingarelli & Sidery 2014). Power-law GWBs are also described (independently of observations) in terms of their *characteristic strain*

$$h_c(f) = A_{\text{gw}} \left(\frac{f}{\text{yr}^{-1}} \right)^\alpha, \quad (2)$$

which is related to Equation (1) by $S_{ab}(f) = \Gamma_{ab} h_c(f)^2 / (12\pi^2 f^3)$ and $\gamma = 3 - 2\alpha$ ($\alpha = -2/3$ for SMBHBs).

Current predictions for the value of A_{gw} , based on models of SMBH–galaxy coevolution and on observational constraints of galaxy assembly and SMBH mass functions, range between 10^{-15} and 10^{-14} (Sesana 2013a; McWilliams et al. 2014; Ravi et al. 2014—hereafter S13, MOP14, and RWS14). In this paper, we rely on the population models developed in these works to provide context (and Bayesian priors) for our analysis of the NANOGrav 9 year data set.

The MOP14 model assumes that all SMBHBs are in circular orbits evolving under GW emission alone, and adopts a single black hole–host correlation from McConnell & Ma (2013), yielding an estimate of the stochastic GWB that is roughly four times as large as S13 and RWS14. The S13 model uses a wide variety of galaxy merger rates and empirical black hole–host relations to yield a collection of phenomenological SMBHB merger rates, which are used to compute a distribution of possible GW signals. In this paper, we consider the specific S13 distribution plotted as the blue dashed curve in Figure 3 of that paper; this distribution includes the black hole–host galaxy relations of McConnell & Ma (2013) and Graham & Scott (2013). The RWS14 model also assumes the black hole–host relation from McConnell & Ma (2013) but accounts for SMBHB evolution in dense stellar environments, as well as significant binary eccentricities. When these models predict spectral densities that deviate from pure power laws at low frequencies, we use their high-frequency predictions (specifically, their value $f = 1 \text{ yr}^{-1}$) to calibrate our fiducial A_{gw} . Finally, we note that recent observations (Kormendy & Ho 2013) find higher black hole masses for a given host than previous work, indicating that even larger A_{gw} may be realized in nature.

We note however that there is a history of BH masses being overestimated, particularly in the most massive galaxies (Merritt & Ferrarese 2001), and recent work suggests that systematic errors greater than 100% may still be common, even in some of the most nearby and well-studied galaxies. Examples include M87, for which the stellar- and gas-dynamical BH masses differ by a factor of two (Gebhardt et al. 2011; Walsh et al. 2013), and NGC 1277, which was originally claimed to have a $\sim 2 \times 10^{10} M_\odot$ BH (van den Bosch et al. 2012), but more recent analyses have decreased that estimate by factors of 3–5 (Emsellem 2013; Walsh et al. 2015).

1.3. New Results in this Paper

Over the last few years, the three PTAs have reported steadily improving upper limits on SMBHB GWBs of the form (1), which are summarized in the uppermost panel of Table 1. The most stringent limit so far is $A_{\text{gw}} < 10^{-15}$ (at 95% confidence and at $f = 1 \text{ yr}^{-1}$; Shannon et al. 2015). It is clear that there is significant tension between these observational limits and astrophysical expectations for A_{gw} . It is important to note that a limit on A_{gw} does not translate *directly* to a limit on the SMBHB population because of the finite number of pulsars that contribute to the limit and the stochasticity of the GW signal itself.

In Section 4.2 of this paper (and in Table 1) we report a new 95% upper limit $A_{\text{gw}} < 1.5 \times 10^{-15}$, obtained from the Bayesian analysis of NANOGrav’s 9 year, 37-pulsar data set released in 2015 (NG9). This limit is five times more constraining than the same analysis applied to NANOGrav’s 5 year data set (Demorest et al. 2013, henceforth NG5); see the top of Figure 4 for a comparison of the two posterior probability distributions. Now, following Shannon et al. (2013), we can assess the consistency of our result with astrophysical GW-background models. We find a 0.8% probability that the observed A_{gw} , as characterized probabilistically by its Bayesian posterior, is drawn from the amplitude distribution developed in MOP14, and a 20% probability that it is drawn from the (very similar) RWS14 and S13 distributions.

Table 1
Summary of Power-law Upper Limits

Source	PTA	Limit	Reference
SMBHB ($\gamma = 13/3$)	NANOGrav	$A_{\text{gw}} < 1.5 \times 10^{-15}$	This paper
	PPTA	$A_{\text{gw}} < 1.0 \times 10^{-15}$	Shannon et al. (2015)
	EPTA	$A_{\text{gw}} < 3.0 \times 10^{-15}$	Lentati et al. (2015)
	NANOGrav	$A_{\text{gw}} < 7 \times 10^{-15}$	Demorest et al. (2013)
Relic ($\gamma = 5$)	NANOGrav	$\Omega_{\text{gw}}(f)h^2 < 4.2 \times 10^{-10}$	This paper
	PPTA	$\Omega_{\text{gw}}(f)h^2 < 1.0 \times 10^{-10}$	Lasky et al. (2015)
	EPTA	$\Omega_{\text{gw}}(f)h^2 < 1.2 \times 10^{-9}$	Lentati et al. (2015)
Cosmic string	NANOGrav	$G\mu < 3.3 \times 10^{-8}$	This paper
	EPTA	$G\mu < 1.3 \times 10^{-7}$	Lentati et al. (2015)

Note. All numbers are quoted at 95% confidence and a reference frequency of 1 yr^{-1} , although differences in the statistical analyses and in the availability and selection of pulsar data sets mean that these numbers are not entirely comparable. Lasky et al. (2015) assume $h = 0.67$; here we multiply their result by h^2 for easy comparison with other PTA limits.

Correspondingly, the two bottom panels of Figure 4 shows that 9 year observations update significantly the MOP14 and RWS14/S13 amplitude priors, much more so than our 5 year data set.

In Section 4.1 we report also a frequentist, *optimal-statistic* 95% upper limit $A_{\text{gw}} < 1.3 \times 10^{-15}$, a fivefold improvement on the analogous result of NG5; however, the optimal statistic is problematic in the presence of marginally detectable GW signals, so we offer it only as a proxy for the improving sensitivity of NANOGrav’s observations.

Stochastic GWBs in the PTA band may also originate from quantum fluctuations amplified during inflation (Grishchuk 2005) and from topological broken-symmetry remnants such as (Damour & Vilenkin 2001; Ölmez et al. 2010). Equation (1) can be used to describe the power spectral density of either of these. For inflation-era quantum fluctuations, $\gamma = 5$ (depending on the equation of state w); for cosmic strings, $\gamma = 16/3$.

In Section 6.1 we obtain 95% upper limits $A_{\text{gw}} < 8.1 \times 10^{-16}$ for relic GWs, corresponding to energy-density limits $\Omega_{\text{gw}}(f = \text{yr}^{-1})h^2 < 4.2 \times 10^{-10}$, where h parametrizes the Hubble constant $H_0 \equiv h \times 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. We then obtain limits on the Hubble parameter during inflation, $H_* = 1.6 \times 10^{-2} m_{\text{pl}}$, where $m_{\text{pl}} \equiv 1/\sqrt{G}$ is the Planck mass, using the method developed by Zhao (2011). Note that we write the Planck mass using the convention $c = \hbar = 1$ such that $m_{\text{pl}} = \sqrt{\hbar c/G} \equiv 1/\sqrt{G}$, and the dimensionless cosmic string tension is written as $G\mu/c^2 \equiv G\mu$.

In Section 6.2, we find $A_{\text{gw}} < 6 \times 10^{-16}$ and $\Omega_{\text{gw}}(f = \text{yr}^{-1})h^2 < 2.2 \times 10^{-10}$, corresponding to a conservative limit on the string tension of $G\mu < 3.3 \times 10^{-8}$, a factor of four better than the joint Planck and high- l cosmic microwave background (CMB) data from other experiments. If we then restrict ourselves to a GWB produced by the

production of large cosmic string loops, as described by Blanco-Pillado et al. (2014), then our string tension limit is much more restrictive: $G\mu < 1.3 \times 10^{-10}$. These new limits are also summarized in the middle and bottom panels of Table 1, where they can be compared with earlier results.

1.4. Astrophysical Inferences

The mismatch between PTA observations and theoretical expectations for the SMBHB background can be explained in different ways. First, it is possible that the astrophysical models and the three assumptions listed above *are* correct, but the background amplitude realized in nature lies in the tails of the predicted distributions. This hypothesis obviously wanes as upper limits become more stringent.

Second, it is possible that some of the *inputs* of the astrophysical models are not estimated correctly; we can then use GW observations to constrain these inputs. As an example, in Section 5.1 we assume measurements of the galaxy mass function and merger rate, and we constrain the scaling between the galaxy bulge mass and the central SMBH mass, which affects the observed A_{gw} most significantly, through the distribution of binary chirp masses. Assuming a purely GW-driven spectrum, we find the scaling relations reported in Kormendy & Ho (2013) and McConnell & Ma (2013) to be inconsistent with our data at the 95% and 90% level, respectively. The McConnell & Ma (2013) black hole–stellar velocity dispersion relation underlies the MOP14 predictions, while S13 and RWS14 take into account a variety of alternative black hole–host estimates.

Third, the simple GW-background characterization that yields Equation (1) may not be realistic, because SMBHBs may form with significant eccentricity and retain it into the PTA band, distributing GW emission over a range of frequency harmonics. Furthermore, environmental effects (interactions with stars on centrophilic orbits in galactic nuclei, or with circumbinary gas disks) can accelerate the transit of individual binary systems through the PTA band (see Sections 5.3 and 5.2 for details and references). These environmental effects deplete the GWB at low frequencies where PTA measurements are most sensitive (i.e., frequencies \sim the inverse observation timespan), so PTA upper limits may yet be compatible with the MOP14/S13/RWS14 predictions at the higher frequencies where the $\gamma = 13/3$ power law is realized.

To investigate this point, in Section 4.2.2 we reanalyze the NANOGrav 9 year data set using a phenomenological $S_{ij}(f)$ in the form of an inflected power law, parametrized by a turnover frequency f_{bend} and by a shape parameter κ , as proposed by Sampson et al. (2015) (see Equation (24)). By combining this enhanced GW-background model with MOP14 and S13/RWS14 amplitude priors, we conclude that the data prefer an inflected spectrum to a moderate degree for MOP14, and weakly for S13/RWS14 models. Quantitatively, the Bayes factors between enhanced and pure-power-law spectral models for each of the two priors are 22 and 2.2, respectively; graphically, the shading in Figure 5 represents the frequency-by-frequency posterior probability density for the GW spectrum, which appears significantly inflected for MOP14, and only slightly so for S13/RWS14. The data are not sufficiently informative to constrain the amplitude *and* shape of the spectrum jointly in the absence of a compact prior, so we cannot produce a unique metric of consistency, as we did above in the case of the simple power-law spectrum.

Beyond this phenomenological characterization, the joint $(A_{\text{gw}}, f_{\text{bend}}, \kappa)$ posteriors can be mapped into constraints for the SMBHB eccentricities (which we do in Section 5.3) and for the astrophysical variables that govern environmental interactions (Section 5.2). We analyze each effect separately: doing so is neither accurate nor robust, since in reality multiple effects may occur together, and may even interact with each other (for instance gas interactions would affect orbital eccentricity). Nevertheless, we believe that such an analysis has interest—it will not provide fully reliable constraints on astrophysical parameters, but it will establish whether the corresponding mechanisms could be responsible for GWB attenuation at low frequencies, under the hypothesis that either the MOP14 and S13/RWS14 population priors are correct.

For eccentricity, we assume for simplicity that all binaries had the same e_0 when the semimajor axis of their orbits was 0.01 pc and that they evolved purely by GW emission since; we follow Huerta et al. (2015) to construct eccentric binary populations and GW strain spectra. The resulting posteriors on e_0 indicate that $e_0 \gtrsim 0.7$ is preferred for the MOP14 prior, and $e_0 \gtrsim 0.5$ is preferred for S13/RWS14 (though still consistent with smaller values). These limits suggest that either SMBHBs form with rather high e_0 , or that binary eccentricity is not a good explanation for the mismatch between A_{gw} observations and predictions.

To characterize environmental interactions (see Section 5.2), we compute the evolution of orbital frequency due to stellar scattering events and to circumbinary gas disk interactions as $df/dt \propto f^{1/3}$ and $f^{4/3}$, respectively, corresponding to $\kappa = 10/3$ and $7/3$ (since for GW-driven evolution $df/dt \propto f^{11/3}$); the frequency f_{bend} then marks the transition between environmentally and GW-driven evolution.

For the case of stellar scattering, f_{bend} depends most significantly on the mass density ρ of galactic-core stars. Astrophysical estimates for ρ are quite uncertain with typical values between 10 and $10^4 M_{\odot} \text{pc}^{-3}$ assuming that a majority of our GW sources come from merging elliptical galaxies (Dotti et al. 2007). Under several simplifying assumptions (e.g., that all binaries have circular orbits, that all galaxies have comparable densities, and that only a single environmental effect is active), we find that $\rho \gtrsim 10^4 M_{\odot} \text{pc}^{-3}$ is strongly preferred for the MOP14 amplitude prior, and the data are unconstraining for the S13/RWS14 prior.

For the circumbinary-disk case, f_{bend} depends on the accretion rate onto the primary (most massive) BH, \dot{M}_1 . The accretion rate of the primary BH is $\dot{M}_1 \propto M_1 \epsilon^{-1} \kappa_{\text{opp}}^{-1}$, where M_1 is the mass of the primary BH; ϵ is the radiative efficiency parameter, with canonical value $\epsilon = 0.1$; and κ_{opp} is the disk opacity. Hence \dot{M}_1 takes on a range of values, typically $10^{-3} M_{\odot} \text{yr}^{-1}$ to $1 M_{\odot} \text{yr}^{-1}$, see, e.g., Di Matteo et al. (2001), Armitage & Natarajan (2002), Goicovic et al. (2015). In our analysis, we find that the accretion rate $\dot{M}_1 \gtrsim 10^{-1} M_{\odot} \text{yr}^{-1}$ is strongly preferred for the MOP14 amplitude prior, and again, the data are unconstraining for the S13/RWS14 prior.

1.5. Plan of the Paper

This paper presents the first analysis that characterizes the spectral amplitude and shape of the GWB; in some cases we find tension between observations and predictions for the GWB from SMBHBs.

The rest of this paper is organized as follows:

1. in Section 2 we describe the NANOGrav 9 year pulsar-timing and data set;
2. in Section 3 we discuss our signal and noise models, as well as the statistical framework of our analysis;
3. in Section 4 we document important implementation details, and report our results in detail;
4. in Section 5 we discuss the astrophysical interpretation of our analysis: specifically, Section 5.1 focuses on the interpretation of power-law limits in terms of black hole–host scaling relations, while Sections 5.2 and 5.3 are concerned with the constraints that can be put on SMBHB environments and eccentricities using a broken-power-law model;
5. in Section 6 we describe the consequences of power-law limits on cosmological quantities;
6. in Section 7 we offer our conclusions.

2. OBSERVATIONS

This paper uses the observations from the NANOGrav 9 year data set, recently presented in NG9, which contains observations made from 2004 to 2013. Initially the array consisted of 15 pulsars, and it grew to 37 pulsars over the course of the project. The first five years of data on 17 pulsars constituted the NANOGrav 5 year data set, previously reported by NG5. For the present work, NG5 data were reprocessed using NG9 procedures.

2.1. Observatories

The Arecibo Observatory (AO) was used to observe pulsars with declinations in the range $0^\circ < \delta < 39^\circ$; pulsars outside of this range were observed with the Green Bank Telescope (GBT) of the National Radio Astronomy Observatory (NRAO). The pulsars PSRs J1713+0747 and B1937+21 were observed at both AO and the GBT. The typical observation cadence was about once every month.

At AO, four receivers were used: 327, 430, 1400, and 2100 MHz. Of those, typically two (or more) were used in immediate succession within ~ 1 hr per observation session. At the GBT, observations were made with two receivers at 800 and 1400 MHz typically separated by a few days. Observations were only included in the data set if observations were made within a time span of 14 days or if wideband data were available; otherwise the observations were discarded for the lack of information about variations in the interstellar medium dispersion.

During an observation, digitized telescope receiver voltages were coherently dedispersed, detected, and folded the data using pre-computed pulsar ephemerides. At each observatory, a narrowband instrument (up to 64 MHz, depending on receiver bandwidth) was used in early years and a wideband instrument (up to 800 MHz) instrument was used in later years. Time offsets between the instruments at each observatory were measured and removed to produce seamless data sets for each pulsar. Data were integrated for short subintervals over the course of a typical 20–30 minute observation. After RFI excision, polarization calibration, and flux calibration, the end product of each observation was a set of total-intensity pulse profiles for a series of subbands.

2.2. Time of Arrival Data

TOAs were calculated using standard techniques for each series of subbands recorded simultaneously in a given observation. The data are available online.³¹ The effect of time-varying dispersion was taken into account by including “DMX” parameters in the timing model (NG5), which essentially allows for an extra delay proportional to $1/\nu^2$ to be fit for, where ν is radio frequency. In other words, a separate dispersion delay parameter is included in the timing model for each data epoch (defined as data collected within a 14 day span). Pulse shape variation is accounted for by an empirical model that parameterizes perturbations of pulse arrival times as a function of frequency; details are given in NG9.

2.3. Timing Models

In NG9, standard timing models were fit to each each pulsar using the TEMPO timing package.³² The timing models were also found to work equally well with TEMPO2.³³ The model for each pulsar included spin-down (pulse period and first derivative), position, proper motion, parallax, and time-variable dispersion measure (DMX). The timing models for binary pulsars included a Keplerian orbital model parametrized in a way appropriate for that orbit. Additional timing parameters, most notably proper motion and post-Keplerian orbital elements, were included in the timing models if they satisfied an F -test with significance value of 0.0027 (essentially detection with 3σ significance). The removal of such a timing model puts each pulsar in a regime in which the response to a perturbation in any parameter is linear, opening the way for the analysis described in Section 3 below. We emphasize that all parameters that we included in the timing model were free parameters, not fixed in the subsequent analysis, in order to prevent systematic biases that would otherwise arise.

3. DATA ANALYSIS METHODS

All the data analysis methods we use in this manuscript (both the Bayesian and frequentist methods) are effectively carried out in the time-domain. We start in Section 3.1 by defining the likelihood function and introducing our notation, which follows that of (NG9) to the letter. For figures, noise realizations, and other quantities we refer the reader to (NG9). We continue our discussion with the Bayesian approach in Section 3.3, and the frequentist approach in the form of the optimal cross-correlation statistic in Section 3.2.

3.1. Likelihood

We start our discussion by decomposing our N_{TOA} timing residual data for a single pulsar $\delta\mathbf{t}$ in its individual constituents as follows:

$$\delta\mathbf{t} = M\boldsymbol{\epsilon} + F\mathbf{a} + U\mathbf{j} + \mathbf{n}. \quad (3)$$

The term $M\boldsymbol{\epsilon}$ describes inaccuracies in the subtraction of the timing model, where M is the timing model design matrix, and $\boldsymbol{\epsilon}$ is a vector of small timing model parameter offsets. The term $F\mathbf{a}$ describes all low-frequency signals, including

low-frequency (“red”) noise, with a limited number of Fourier coefficients \mathbf{a} . Our harmonics are chosen as integer multiples of the harmonic base frequency $1/T$, with T the length of our data set (either of a single pulsar, or the entire array of pulsars). The number of harmonics is chosen high enough to the point that adding in extra harmonics do not change results (set to 30 for all pulsars). The matrix F then has alternating sine and cosine functions. We note that this is just a particular choice of rank-reduced basis, and we could have chosen many others without influencing our results. The term $U\mathbf{j}$ describes noise that is fully correlated (correlation coefficient of 1) for simultaneous observations at different observing frequencies, but fully uncorrelated in time. Pulse jitter is an example of an effect that would look like that, with the note that pulse jitter should in principle scale inversely with integration time. Our integration times are (nearly) constant over the data set, and we do not take integration time into account at this level of the analysis. The matrix U is an $N_{\text{TOA}} \times N_s$ matrix that maps the N_s subchannels (i.e., TOAs measured simultaneously) to the N_{TOA} TOAs. The vector \mathbf{j} describes the white noise per observation session that is fully correlated across all observing frequencies. The last term, \mathbf{n} , describes Gaussian white noise. The white noise is assumed to be an uncorrelated, Gaussian noise process, with variance-covariance:

$$N_{ij,k} \equiv \langle \mathbf{n}_i \mathbf{n}_j \rangle_k = E_k^2 (W_{ij} + Q_k^2 \delta_{ij}), \quad (4)$$

where E_k and Q_k are the TEMPO and TEMPO2 “EFAC” and “EQUAD” parameters for each observing system (Receiver + backend) k , δ_{ij} is the Kronecker delta function, and $W = \text{diag}\{\sigma_i^2\}$ are the TOA uncertainties. The matrix elements (i, j) apply to the TOAs corresponding to the observing system labeled by k . In practice we cannot fully separate various contributions to our TOAs, so we have to take into account that our corrections for the various terms of Equation (3) are imprecise. We do this by constructing our likelihood from the noise-mitigated timing residuals:

$$\mathbf{r} = \delta\mathbf{t} - M\boldsymbol{\epsilon} - F\mathbf{a} - U\mathbf{j}, \quad (5)$$

where \mathbf{r} is our best approximation of \mathbf{n} , given our knowledge of all the noise and signal parameters. The likelihood can now be written down as:

$$p(\delta\mathbf{t}|\boldsymbol{\epsilon}, \mathbf{a}, \mathbf{j}, \boldsymbol{\eta}) = \frac{\exp\left(-\frac{1}{2}\mathbf{r}^T N^{-1}\mathbf{r}\right)}{\sqrt{\det(2\pi N)}}, \quad (6)$$

where $N = \sum_k N_k$ represents the total effect of all white noise. We have collectively denoted all parameters not directly represented by $\boldsymbol{\epsilon}$, \mathbf{a} , and \mathbf{j} as $\boldsymbol{\eta}$. Henceforth, we shall refer to $\boldsymbol{\eta}$ as the hyperparameters. We group the reduced-rank signals as follows:

$$T = [M \ F \ U], \quad \mathbf{b} = \begin{bmatrix} \boldsymbol{\epsilon} \\ \mathbf{a} \\ \mathbf{j} \end{bmatrix}, \quad (7)$$

³¹ <http://data.nanograv.org>

³² <http://tempo.sourceforge.net>

³³ <https://bitbucket.org/mkeith/tempo2>

which allows us to place a Gaussian prior on the coefficients of these random processes. The prior covariance is:

$$B = \begin{bmatrix} \infty & 0 & 0 \\ 0 & \varphi & 0 \\ 0 & 0 & \mathcal{J} \end{bmatrix}, \quad (8)$$

resulting in a prior:

$$p(\mathbf{b}|\boldsymbol{\eta}) = \frac{\exp\left(-\frac{1}{2}\mathbf{b}^T B^{-1}\mathbf{b}\right)}{\sqrt{\det(2\pi B)}}, \quad (9)$$

where ∞ is a diagonal matrix of infinities, which effectively means we have a uniform unconstrained prior on the timing model parameters ϵ . As described in (NG9), this representation allows us to analytically marginalize Equation (6) times Equation (9) over the waveform coefficients \mathbf{b} of the noise, resulting in a drastic reduction of the dimensionality of our posterior (Lentati et al. 2013; van Haasteren & Vallisneri 2014, 2015):

$$p(\delta\mathbf{t}|\boldsymbol{\eta}) = \frac{\exp\left(-\frac{1}{2}\delta\mathbf{t}^T C^{-1}\delta\mathbf{t}\right)}{\sqrt{\det(2\pi C)}}, \quad (10)$$

with $C = N + \text{TBT}^T$. The Woodbury matrix identity (Woodbury 1950) can be used to evaluate Equation (10) efficiently.

The parameters that describe B are the hyperparameters $\boldsymbol{\eta}$. The hyperparameters of the diagonal matrix \mathcal{J} are the per-system TEMPO2 ECORR parameters, which represent the variance of the per-epoch correlated effects given by Uj . The matrix φ represents the spectrum of the low-frequency noise and the stochastic gravitational waves, and it therefore contains terms correlated between pulsars. Denoting pulsar number with (a, b) , and frequency bin with (i, j) , we can write:

$$[\varphi]_{(ai),(bj)} = \Gamma_{ab}\rho_i\delta_{ij} + \eta_{ai}\delta_{ab}\delta_{ij}, \quad (11)$$

where η_a is the spectrum of the low-frequency noise of pulsar a , ρ is the spectrum of the GWB, and Γ_{ab} is the signal correlation matrix. The elements of the signal correlation matrix consist of the overlap reduction function for a GWB signal, which is a dimensionless function that quantifies the correlated response of the pulsars to a stochastic GWB (Mingarelli & Sidery 2014). The quantities ρ_i and η_{ai} can be either modeled as independent model parameters (i.e., per frequency), or as modeled spectral density with a specific shape (e.g., a power-law model). We note that φ can be represented by a block-diagonal matrix, where each block corresponds to a specific frequency bin; all frequencies are nominally independent degrees of freedom.

3.2. Optimal Cross-correlation Statistic

The *optimal statistic* (Anholm et al. 2009; Demorest et al. 2013; Chamberlin et al. 2015) is a frequentist estimator of the isotropic GWB strain-spectrum amplitude in the weak-signal regime. This estimator is derived by maximizing the signal-to-noise ratio (S/N) of a cross-correlation statistic and can be written as:

$$\hat{A}_{\text{gw}}^2 = \frac{\sum_{ab} \delta\mathbf{t}_a^T \mathbf{P}_a^{-1} \tilde{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \delta\mathbf{t}_b}{\sum_{ab} \text{tr}[\mathbf{P}_a^{-1} \tilde{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \tilde{\mathbf{S}}_{ba}]}, \quad (12)$$

where $\mathbf{P}_a = \langle \delta\mathbf{t}_a \delta\mathbf{t}_a^T \rangle$ is the auto-covariance matrix of the residuals of a single pulsar. This is C of Equation (10) for a noise-only model. The term $A_{\text{gw}}^2 \tilde{\mathbf{S}}_{ab} = \mathbf{S}_{ab} = \langle \delta\mathbf{t}_a \delta\mathbf{t}_b^T \rangle$ represents the signal covariance between pulsar a and pulsar b . As described in the previous section, our model contains no other signals with a non-zero correlation coefficient between different pulsars. The normalization of the optimal statistic is such that $\langle \hat{A}_{\text{gw}}^2 \rangle = A_{\text{gw}}^2$.

Following Chamberlin et al. (2015), we also quote expressions for the variance, and the S/N (in power) of the optimal statistic. The variance of the optimal statistic in the absence of a GWB signal is given by:

$$\sigma_0^2 = \left(\sum_{ab} \text{tr}[\mathbf{P}_a^{-1} \tilde{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \tilde{\mathbf{S}}_{ba}] \right)^{-1}. \quad (13)$$

Although this expression is not valid in general, in the weak-signal regime, in which the cross correlated power is ignored from the variance, this can be used as an approximation of the variance in \hat{A}_{gw}^2 . The S/N for a given signal and noise realization is:

$$\rho = \frac{\hat{A}_{\text{gw}}^2}{\sigma_0} = \frac{\sum_{ab} \delta\mathbf{t}_a^T \mathbf{P}_a^{-1} \tilde{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \delta\mathbf{t}_b}{\left(\sum_{ab} \text{tr}[\mathbf{P}_a^{-1} \tilde{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \tilde{\mathbf{S}}_{ba}] \right)^{1/2}}, \quad (14)$$

which has an expectation value of

$$\langle \rho \rangle = A_{\text{gw}}^2 \left(\sum_{ab} \text{tr}[\mathbf{P}_a^{-1} \tilde{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \tilde{\mathbf{S}}_{ba}] \right)^{1/2}. \quad (15)$$

These expressions generalize the detection significance estimator provided in Jenet et al. (2005), properly taking into account the spectra of the signal and the noise, as well as details such as the irregular sampling. The S/N here is essentially a ratio of the significance of the cross-correlated power to the power of the same signal without cross-correlations. If interpreted as a zero-mean unit-variance normal distribution³⁴, the optimal statistic can be used to place upper-limits on the GWB amplitude. Clearly, this interpretation is only appropriate in the weak-signal regime, but it serves as an independent sanity check for our other methods. Setting our statistical significance threshold $\alpha = 0.95$, we can place a one-sided upper-limit as:

$$A_{\text{ul}}^2 = \hat{A}_{\text{gw}}^2 + \sqrt{2} \sigma_0 \text{erfc}^{-1}[2(1 - \alpha)]. \quad (16)$$

We note that all the usual caveats of frequentist-type upper-limits apply to this methodology as well, as no prior information is used. For instance, it is possible to set an upper-limit of $A_{\text{ul}} = 0$ in a data set without a (detectable) signal, which theoretically happens with probability α .

It was shown in Chamberlin et al. (2015) that the optimal statistic is identical to the cross-correlation statistic used by NG5. This alternative interpretation of the optimal statistic allows us to obtain a measure of the cross power between pulsars. The cross power is the amount of correlated power between the timing residuals of different pulsars, and one

³⁴ In general this will be a chi-squared distribution; however, since we are summing over the number of pulsar pairs, the central limit theorem allows us to approximate the optimal statistic distribution as Gaussian.

would expect this cross power to follow the Hellings and Downs cross-correlation signature for a detectable GWB signal. The cross power and the uncertainty estimates are given by

$$\xi_{ab} = \frac{\delta \mathbf{t}_a^T \mathbf{P}_a^{-1} \hat{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \delta \mathbf{t}_b}{\text{tr}[\mathbf{P}_a^{-1} \hat{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \hat{\mathbf{S}}_{ba}]}, \quad \sigma_{0,ab} = (\text{tr}[\mathbf{P}_a^{-1} \hat{\mathbf{S}}_{ab} \mathbf{P}_b^{-1} \hat{\mathbf{S}}_{ba}])^{-1/2}, \quad (17)$$

which is independent of any specific overlap reduction function and $\hat{A}_{\text{gw}}^2 \Gamma_{ab} \hat{\mathbf{S}}_{ab} = \mathbf{S}_{ab}$. One can then fit these correlation coefficients assuming a particular overlap reduction function by minimizing

$$\chi^2 = \sum_{ab} \left| \frac{\xi_{ab} - A_{\text{gw}}^2 \Gamma_{ab}}{\sigma_{0,ab}} \right|^2, \quad (18)$$

where Γ_{ab} are the cross-correlation coefficients given by the overlap reduction function for an isotropic GWB. It can easily be shown that the best fit value of A_{gw} is then the optimal statistic value of Equation (12).

3.3. Bayesian Analysis

As an alternative to the frequentist approach to data analysis, Bayesian inference is a method of statistical inference in which Bayes' rule of conditional probabilities is used to update one's knowledge as observations are acquired. Given a model \mathcal{H} , model parameters Θ , and observations \mathcal{D} , we write Bayes rule as:

$$\Pr(\Theta|\mathcal{D}, \mathcal{H})\Pr(\mathcal{D}|\mathcal{H}) = \Pr(\mathcal{D}|\Theta, \mathcal{H})\Pr(\Theta|\mathcal{H}), \quad (19)$$

where $\Pr(\Theta|\mathcal{D}, \mathcal{H}) \equiv \Pr(\Theta)$ is the posterior (probability distribution) of the parameters, $\Pr(\mathcal{D}|\Theta, \mathcal{H}) \equiv L(\Theta)$ is the likelihood, $\Pr(\Theta|\mathcal{H}) \equiv \pi(\Theta)$ is the prior (probability distribution), and $\Pr(\mathcal{D}|\mathcal{H}) \equiv \mathcal{Z}$ is the marginal likelihood or evidence.

The left-hand side of Equation (19) can be regarded as the ‘‘output’’ of the Bayesian analysis, and the right-hand side is the ‘‘input.’’ Indeed, provided we have a generative model of our observations (meaning we can simulate data, given the model parameters), we know the likelihood and prior. However, for parameter estimation we would like to know the posterior, and for model selection we need the evidence.

For parameter estimation, the evidence \mathcal{Z} is usually ignored, and one can use $L(\Theta)\pi(\Theta)$ directly to estimate confidence intervals. Typically one provides confidence intervals for single components and pairs of elements of Θ . This involves an integral over $\Pr(\Theta)$ over all but one or two parameters, a process called marginalization. When Θ is higher-dimensional, Monte-Carlo sampling methods are typically used to perform this multi-dimensional integral. We use Markov chain Monte Carlo methods in this work to sample the posterior distribution.

Model selection between two models \mathcal{H}_0 and \mathcal{H}_1 can be carried by calculating the ‘‘Bayes factor’’: the ratio between the evidence for the two models. Assuming we have a prior degree of belief of how likely the two models are ($\Pr(\mathcal{H}_0)$ and $\Pr(\mathcal{H}_1)$), we can write:

$$O = \frac{\Pr(\mathcal{H}_1|\mathcal{D})}{\Pr(\mathcal{H}_0|\mathcal{D})} = \mathcal{B}_{10}(\mathcal{D}) \frac{\Pr(\mathcal{H}_1)}{\Pr(\mathcal{H}_0)}, \quad (20)$$

where $\mathcal{B}_{10}(\mathcal{D}) \equiv \mathcal{Z}_1/\mathcal{Z}_0$ is the Bayes factor, and O is the odds ratio. The odds ratio can be obtained by calculating the

evidence \mathcal{Z} for each model separately (e.g., with nested sampling or thermodynamic integration), or by calculating the Bayes factor \mathcal{B} between two models directly (e.g., with transdimensional Markov chain Monte Carlo methods).

4. RESULTS

Below we first analyze the data using the optimal statistic of Section 3.2 primarily to facilitate comparison with previous work, and then we analyze it using the Bayesian methods of Section 3.3.

4.1. Optimal Statistic Analysis

The optimal cross correlation statistic was applied to the full 37 pulsar 9 year NANOGrav data set. The maximum-likelihood single-pulsar noise values were obtained by independent noise analyses on each pulsar. The maximum-likelihood amplitude and S/N of this search are $\hat{A}_{\text{gw}} = 8.9 \times 10^{-16}$ and $\rho = 1.5$, respectively, indicating little evidence for the expected GWB cross correlations. The resulting upper limit using this method is $A_{\text{gw}} < 1.3 \times 10^{-15}$ at a reference frequency of yr^{-1} , which is 5.4 and 1.5 times more stringent than the limits using this method presented in NG5 and (Lentati et al. 2015, henceforth LTM15), respectively. It should be noted that the limiting technique presented in Section 3.2 does not strictly have proper frequentist coverage (i.e., the $x\%$ upper limit amplitude is not greater than the injected value in $x\%$ of realizations) in the presence of any measurable GWB signal;³⁵ therefore, this limit will serve as a proxy to our improved sensitivity and not a strict upper limit.

As a by-product of the optimal statistic analysis, we can obtain the maximum-likelihood cross-correlation values for each pulsar pair in the analysis. In Figure 1 we plot the cross-correlated power versus angular separation in the top panel. We have binned the values into 10° bins using a weighted averaging technique. The red curve shows the maximum-likelihood correlated power fit. It is clear that the cross-correlated power is still consistent with zero signal and that we have much more sensitivity at values of $\zeta < 100^\circ$, which is a by-product of the fact that our best timed pulsars all lie at a similar position on the sky. This is illustrated in the bottom panel of Figure 1, where we plot the histogram of the number of pulsars in each ten-degree bin (red) as well as the most significant bins (blue).

4.2. Bayesian Analysis

We have used the Bayesian data analysis techniques described in Section 3 to place upper limits on the GWB parameterized using the standard power law, spectral components and a broken-power-law. Bayesian data analysis of modern PTA data sets is quite difficult to perform in general due to the large parameter spaces and likelihood evaluation time. For the current set of 37 pulsars a full GWB analysis would have to search over at least 8 parameters per pulsar (EFAC, EQUAD, and ECORR per observing system, assuming two observing systems per pulsar, and a red noise amplitude

³⁵ In the limit of zero GWB signal, this method actually over covers in the frequentist sense. We have found by comparison with other frequentist bounding techniques and Bayesian upper limits that the 5 year limit published in NG5 is indeed robust and does not suffer from under coverage.




Figure 1. Top panel: cross-correlated power vs. angular separation obtained through an optimal statistic analysis. The dashed red curve shows the maximum-likelihood amplitude mapped onto the Hellings and Downs coefficients. The blue points show the measured cross-correlation values, ξ_{ab} , and their 1-sigma uncertainties. Bottom panel: histogram of the number of pulsars in each bin (red right axis) and a weighted histogram (blue left axis) using the uncertainties in the cross correlation as weights.

and spectral index) and up to 14 for pulsars with observations at both AO and GBT. Overall this means that we would have to search over 524 parameters in the analysis. Even with the efficient methods described in Section 3 we are still prohibited from carrying out a full Bayesian search where we vary all noise parameters and GW parameters for all pulsars simultaneously.

In this work we ameliorate these two problems by using a two-tiered approach to noise modeling and GWB analysis. To avoid including all noise parameters in our GWB analysis, we first carry out single-pulsar noise analyses in which we include EFAC, EQUAD, and ECORR for each backend/receiver system and red noise parameterized as a power law. We then perform the GWB analysis while holding all EFAC, EQUAD, and ECORR values fixed to their mean value from the single-pulsar analysis but allowing the red noise and GWB parameters to vary. By holding these white noise parameters fixed, we reduce the computational burden since large matrix products can be pre-computed. Furthermore, by holding all white noise values fixed we reduce the number of free parameters drastically from ~ 524 to $2N_{\text{psr}} + N_{\text{gw}}$, where N_{psr} and N_{gw} are the number of pulsars in the array and number of parameters describing the GWB, respectively.

We further reduce the computational burden in two ways: by ignoring cross correlations and by only using a subset of pulsars. We choose to ignore cross correlations in this work as they do not contribute to the upper limit in the absence of a signal and only serve to greatly increase our computational burden by requiring that we invert a large dense matrix for every iteration in the analysis. That the cross correlations have no bearing on the upper limit has been argued in Shannon et al. (2013) and further shown in LTM15. A real GWB signal will reveal itself as a strong *common* red noise term well before the

Table 2
Pulsars Used in GWB Analysis

PSR Name	$A_{\text{gw}}^{95\%}$ (10^{-15})	rms ^a (μs)	Timing Baseline (year)
J1713+0747	1.96	0.116	8.76
J1909-3744	4.50	0.081	9.04
J1640+2224	11.8	0.158	8.90
J1600-3053	12.3	0.197	5.97
J2317+1439	13.6	0.267	8.87
J1918-0642	16.0	0.340	9.01
J1744-1134	16.1	0.334	9.21
J0030+0451	19.8	0.723	8.76
J0613-0200	23.4	0.592	8.58
J1614-2230	23.9	0.189	5.09
B1855+09	26.6	1.338	8.86
J1853+1303	31.0	0.235	5.60
J2145-0750	33.0	0.370	9.07
J1455-3330	37.9	0.694	9.21
J1012+5307	43.3	1.197	9.21
J1741+1351	56.8	0.103	4.24
J2010-1323	83.3	0.312	4.08
J1024-0719	92.9	0.280	4.01

Note.

^a Weighted root-mean-square of epoch-averaged post-fit timing residuals. See Arzoumanian et al. (2015) for more details.

cross correlations become detectable, and since we see no evidence for a common red noise term, we are justified in dropping these terms. Lastly, we choose to only include a subset of pulsars in our analysis as not all pulsars contribute to our upper limit either due to short timing baselines or large measurement errors. To choose this subset of pulsars, we first carry out our single pulsar analyses mentioned above but now include an extra red noise process with power spectral index fixed to 13/3 (i.e., SMBHBs). The pulsars are then sorted based on their single pulsar upper limits on the GWB. We then carry out the GWB analysis mentioned above to compute an upper limit using an increasing number of pulsars in the sorted list. This process is continued until the upper limit saturates. In other words, including more pulsars beyond this point does not change the upper limit. Using this method we choose to use the 18 pulsars shown in Table 2. Note that it may seem counterintuitive that PSR J1909-3744 gives a significantly worse single-pulsar upper limit than PSR J1713+0747 when it has a similar timespan and rms. However, as mentioned in NG9, PSR J1909-3744 has evidence for weak red noise that acts to give a less constraining single-pulsar limit. In Shannon et al. (2015), no evidence of red noise is seen in J1909-3744; however, only 10 cm data was used in that work thereby making a direct comparison difficult.

To compute the posterior probability of Equation (10) and to map out the multi-dimensional parameter space we make use of the pulsar timing data analysis suite PAL2³⁶ in conjunction with a parallel-tempering Markov Chain Monte-Carlo (PTMCMC) code.³⁷ The details of the PTMCMC sampler can be found in Appendix C of Arzoumanian et al. (2014). The parameters and prior ranges used in the analysis are shown in Table 3. As shown in the table, we have chosen to use uniform priors on both the red noise and GWB amplitude parameters. We chose a uniform GWB amplitude prior so that it is proper at

³⁶ <https://github.com/jellis18/PAL2>

³⁷ <https://github.com/jellis18/PTMCMCSampler>

Table 3
Summary Model Parameters and Prior Ranges

Parameter	Description	Prior	Comments
White Noise			
E_k	EFAC per backend/receiver system	uniform in $[0, 10]$	Only used in single pulsar analysis
Q_k (s)	EQUAD per backend/receiver system	uniform in logarithm $[-8.5, -5]$	Only used in single pulsar analysis
J_k (s)	ECORR per backend/receiver system	uniform in logarithm $[-8.5, -5]$	Only used in single pulsar analysis
Red Noise			
A_{red}	Red noise power-law amplitude	uniform in $[10^{-20}, 10^{-11}]$	1 parameter per pulsar
γ_{red}	Red noise power-law spectral index	uniform in $[0, 7]$	1 parameter per pulsar
GWB			
A_{gw}	GWB power-law amplitude	uniform in $[10^{-18}, 10^{-11}]$	1 parameter for PTA for power-law models
γ_{gw}	GWB power-law spectral index	delta function	Fixed to different values depending on analysis
ρ_i (s^2)	GWB power spectrum coefficients at frequency i/T	uniform in $\rho_i^{1/2}$ $[10^{-18}, 10^{-8}]^{\text{a}}$	1 parameter per frequency
A	GWB broken power-law amplitude	log-normal ^b for models A(B) $\mathcal{N}(-14.4(-15), 0.26(0.22))$	1 parameter for PTA for broken power-law models
κ	GWB broken power-law low-frequency spectral index	uniform in $[0, 7]$	1 parameter for PTA for broken power-law models
f_{bend} (Hz)	GWB broken power-law bend frequency	uniform in logarithm $[-9, -7]^{\text{c}}$	1 parameter for PTA for broken power-law models

Notes.

^a The prior uniform in $\rho_i^{1/2}$ is chosen to be consistent with a uniform prior in A_{gw} for the power-law model since $\varphi_{ii} \propto A_{\text{gw}}^2$.

^b These values are quoted in log base 10 and are obtained from MOP14 to S13.

^c We choose different prior values on f_{bend} when mapping to astrophysical model parameters as described in Section 5.

$A_{\text{gw}} = 0$, that is to say, it must converge to a finite value in the limit of zero amplitude, otherwise the upper limit would depend on the lower bound of the prior which is undesirable. However, for the red noise amplitude we have no such restriction. A log-uniform prior is non-informative and would result in the most unbiased parameter estimation and a uniform prior would bias the parameter estimation toward higher red noise amplitude and lower red noise spectral indices (since there is a strong covariance between these two parameters); however, the log uniform prior along with a uniform prior on the GWB amplitude could cause some intrinsic red noise to be modeled by the GWB amplitude if the intrinsic red noise spectral index is consistent with that of the modeled GWB. Following the precedent of LTM15 and taking in to account the considerations above, we choose uniform priors on both as it gives equal weight to both red noise and the GWB.

4.2.1. Power-law and Spectral Limits

When computing upper limits on the dimensionless strain amplitude we fix the spectral index (13/3 in the case of SMBMBs) and adopt a uniform prior on A_{gw} . When performing a spectral analysis we again use priors that correspond to a uniform prior on A_{gw} ; this results in a prior that is uniform in the square root of the power spectrum coefficients.

A summary of all power-law limits is shown in Table 1 in the introduction. Figure 2 shows the results of the power law and spectral analysis along with relevant astrophysical model predictions. The solid black and long dashed black lines are the 95% upper limits from the spectral and power-law analyses, respectively. The blue, gray, and red shaded regions are the one-sigma prediction on the strain spectra from MOP14, RWS14, and S13, respectively. We find an upper limit on the dimensionless strain amplitude of $A_{\text{gw}} < 1.5 \times 10^{-15}$, slightly less constraining than the most stringent published upper limit to date (Shannon et al. 2015), a

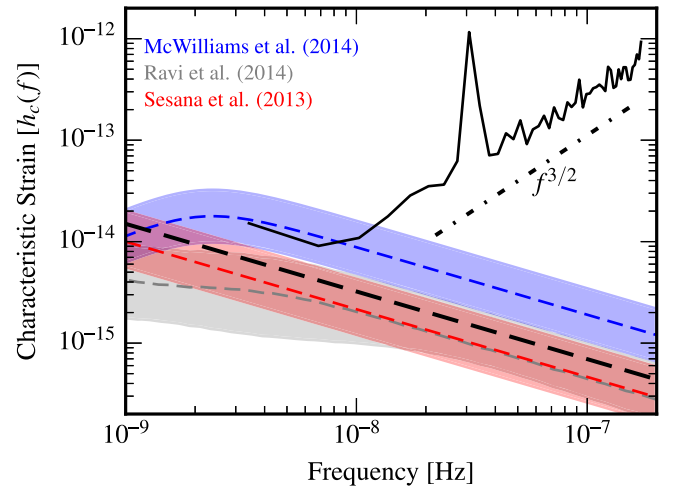


Figure 2. Strain amplitude vs. GW frequency. The solid black and long dashed black lines are the 95% upper limits from our spectral and power-law analyses. The red, gray, and blue shaded regions are the one-sigma predictions from the models of S13, RWS14, and MOP14. The blue shaded region uses the simulation results from MOP14, but replaces the fit to the GWB predictions used in that paper with the functional form given by Equation (24). The dashed-dotted line shows the expected slope of the strain spectrum for white noise.

factor of two more constraining than the recent EPTA upper limit of LTM15, and a factor of five more constraining than the 5 year data set upper limit when applying the same Bayesian analysis. Furthermore, we find a slightly less constraining upper limit when using the free spectrum model (i.e., the power at each frequency is a free parameter and not parameterized by a power law). This is to be expected since the free spectrum model has many more degrees of freedom (we use 50 free amplitudes for each of the 50 frequencies in this case) over the power law parameterization (1 degree of freedom). We also find that the upper limit on the strain spectrum from the spectrum analysis is consistent with white noise (i.e.,

$h_c^{\text{white}}(f) \propto f^{3/2}$) at frequencies $\gtrsim 2/T$, where T is the length of the longest set of residuals in the data set, which indicates that our GWB upper limits are coming from the three lowest frequency bins. This behavior is to be expected since we have several well timed pulsars that do not span the full 9 year baseline (see Table 2) and thus will have peak sensitivity at frequencies greater than $1/T$.

From inspection of Figure 2 we see that our 95% upper limit is within at least the 2-sigma confidence region of all three astrophysical models and is sensitive to a potential turnover in the spectrum due to environmental coupling factors. We wish to determine the level of consistency between our data and the power-law models displayed in Figure 2. To accomplish this we follow the method applied in Shannon et al. (2013). Given that we have a model M for the value of the GW amplitude A_{gw} whose probability distribution function is denoted $p(A_{\text{gw}}|M)$ and that we have a probability distribution function for A_{gw} given the data, denoted $p(A_{\text{gw}}|d)$, where d represents the data, the probability that we measure a value of \hat{A}_{gw} greater than that predicted by the model, A_{gw}^M , is given by the law of total probability

$$P(\hat{A}_{\text{gw}} > A_{\text{gw}}^M) = \int_{-\infty}^{\infty} p(A_{\text{gw}}|M) dA_{\text{gw}} \int_{A_{\text{gw}}}^{\infty} p(A'_{\text{gw}}|d) dA'_{\text{gw}}. \quad (21)$$

Therefore, low values of $P(\hat{A}_{\text{gw}} > A_{\text{gw}}^M)$ indicate that the range of A_{gw} that is consistent with our data is *inconsistent* with the model M , and vice versa. To carry out this procedure the distribution $p(A_{\text{gw}}|d)$ is simply the marginalized posterior distribution when using the uniform prior on A_{gw} . We use log-normal distributions to model the MOP14, S13, and RWS14, models. Since the models of RWS14, and S13 predict nearly the same GWB amplitude distribution (assuming a power law only) we make no distinction between these two models. Furthermore, the model distributions on A_{gw} , given by log-normal distributions have mean and standard deviations of $(-14.4, -15)$ and $(0.26, 0.22)$ for the MOP14 (hereafter model A) and S13/RWS14 (hereafter model B) models, respectively. Using the aforementioned distributions and Equation (21) we find that our data are 0.8% and 20% consistent with models A and B, respectively, under the assumption of a power law. This indicates that either the assumptions that go into these models are incorrect, our universe is a realization of the GWB that has an amplitude in the tail of the probability distributions mentioned above, or that environmental effects are depleting SMBHB sources at low frequencies making the power-law assumption faulty. The implications of this last point are discussed in Section 5.

In addition to our power-law limits on the stochastic background (i.e., strain spectral index $-2/3$), we have also computed the upper limit on the GWB for a range of different spectral indices. In Figure 3 we plot the upper limits obtained at varying spectral indices (red points) versus power spectral index. We also provide the best fit model for the upper limit as a function of power spectral index where we find $A_{\text{gw}} = 7.9 \times 10^{-14} \times 10^{-0.4\gamma} \propto T^{0.83\alpha}$. This differs from NG5 where they find $A_{\text{gw}} \propto T^\alpha$, arguing that this is due to the fact that the sensitivity is dominated by the lowest frequency of $1/T$. Our fit, giving a slightly weaker dependence on α is consistent with what we have seen above, namely that




Figure 3. Upper limit on the GWB as a function of power spectral index.




Figure 4. Marginalized posterior probability density of $\log_{10} A_{\text{gw}}$ computed using the nine (blue) and five (red) year NANOGrav data sets for uniform, MOP14 model Gaussian, and S13/RWS14 model Gaussian prior distributions. The Gaussian priors are shown in green.

our limits are not completely dominated by the lowest frequency.

In a Bayesian analysis, the posterior distribution is the prior distribution updated by the data. Here we illustrate this by comparing our power-law upper limits, using identical methods, on the 5 year NG5 and 9 year NG9 NANOGrav data sets. The results of this comparison are shown in Figure 4 where we plot the marginalized posterior distributions of $\log_{10} A_{\text{gw}}$ for the 9 and 5 year data sets in blue and red, respectively. The Gaussian prior distributions described above are shown in green for model A and model B. For the uniform prior case we see quite a dramatic improvement (i.e., the factor of five mentioned above) in the upper limits. For model A; the 5 year data set does marginally inform the prior, whereas the 9 year data set results in a posterior that is largely inconsistent with the prior distribution. For model B the 5 year data set does not inform the prior at all, whereas the 9 year data set does indeed update the prior, therefore showing that we have actually *learned* new information from the 9 year data set. Note

that this analysis only applies to *power-law* models and that the results would significantly differ if more flexible models were used.

4.2.2. Broken Power-law Limits

We place constraints on the strength of environmental coupling effects that will likely affect our GWB signal at low frequencies (i.e., large orbital separations) via a simple parameterization of the GWB spectrum that allows for a “bend” frequency at which there is a transition from environmentally driven evolution to GW-driven evolution. The following discussion and analysis techniques are based on Sampson et al. (2015). Here we give a brief overview of this more generalized GWB spectrum.

The characteristic amplitude of a stochastic background from an ensemble of SMBHBs in circular orbits is (Phinney 2001; Sesana et al. 2008; McWilliams et al. 2014)

$$h_c(f)^2 = \int_0^\infty dz \int_0^\infty d\mathcal{M} \frac{d^3N}{dz d\mathcal{M} dt} \frac{dt}{d \ln f} h^2(f), \quad (22)$$

where $d^3N/(dz d\mathcal{M} dt)$ is the differential number of inspiraling binaries per unit z , \mathcal{M} and t , where z is the redshift, $\mathcal{M} = (m_1 m_2)^{3/5}/(m_1 + m_2)^{1/5}$ is the chirp mass of the binary, and t is the time measured in the binary rest frame. The $dt/d \ln f$ term describes the frequency evolution of the binary system, and $h(f)$ is the strain spectrum emitted by a single circular binary with orbital frequency $f/2$. Typically, it has been assumed that the binary is purely GW-driven, which results in our usual expression for $h_c(f)$ given in Equation (2); however, physical mechanisms other than GW radiation that are necessary to drive the binary to coalescence (Milosavljević & Merritt 2003) will change the frequency dependence (i.e., the $dt/d \ln f$ term) of this equation (see Colpi 2014, for a review of SMBHB coalescence). Following Sampson et al. (2015) we can generalize the frequency dependence of the strain spectrum to

$$\frac{dt}{d \ln f} = f \left(\frac{df}{dt} \right)^{-1} = f \left(\sum_i \left(\frac{df}{dt} \right)_i \right)^{-1}, \quad (23)$$

where i ranges over many physical processes that are driving the binary to coalescence. If we restrict this sum to GW-driven evolution and an unspecified physical process then the strain spectrum can be written

$$h_c(f) = A \frac{(f/f_{\text{yr}})^\alpha}{(1 + (f_{\text{bend}}/f)^\kappa)^{1/2}}, \quad (24)$$

where f_{bend} and κ are the parameters that encode information about the physical processes (other than GW radiation) driving the binary evolution. Note here that the amplitude, A , is not strictly the same parameter as A_{gw} unless $f_{\text{bend}} \gg 1/T$. As mentioned above, there could be many physical processes contributing to the frequency evolution of the SMBHB system; however, at current sensitivity it is very unlikely that our data can distinguish them. Thus we follow Sampson et al. (2015) and adopt this simplified spectrum to place constraints on possible environmental coupling mechanisms.

The above discussion has focused on the frequency evolution of SMBHB. The other piece of the equation is the

merger rate of SMBHBs. Assuming that the final-parsec problem (Milosavljević & Merritt 2003) is solved and SMBHBs do not stall, the SMBHB merger rate will set the overall amplitude scale for the GWB somewhat independently of the physical mechanisms that drive the system to merger. Therefore, by specifying a prior on the GWB amplitude in this case, we are assuming a model for the behavior of the SMBHBs at small orbital separations (i.e., high GW frequencies) but we are *not* assuming that the dynamics are entirely GW driven at large orbital separations (i.e., low GW frequencies). Here we use the same log-normal distributions introduced in Section 4.2.1 for models A and B as prior probability distributions for the GWB amplitude A in Equation (24). In the following analysis we also fix $\alpha = -2/3$ and use uniform priors on $\log_{10} f_{\text{bend}} \in [-9, -7]$ and $\kappa \in [0, 7]$ unless stated otherwise. It is worth noting that by specifying a prior for the GWB amplitude we are in no way claiming that these models include all of the relevant physics and are *correct*. We have chosen these models based on the fact that they essentially cover the plausible range of GWB amplitudes under the assumption of negligible stalling. Therefore, it is very useful to keep in mind that the results presented below only apply to these specific models but are nonetheless very useful as it allows us to assess the validity of these predictions under more flexible spectral shapes.

Here, we have run an identical analysis to that of Section 4.2.1 except that the GWB spectrum model is now that of Equation (24) and we adopt the aforementioned priors on A , f_{bend} , and κ . Figures 5 and 6 show the results of this analysis. Figure 5 shows the posterior probability density of the GWB spectrum defined in Equation (24) with model A on the left and model B on the right. This probability density is constructed by drawing values of A , f_{bend} , and κ from the joint probability distribution shown in Figure 6, constructing the spectrum at each frequency via Equation (24), and then producing a histogram of the spectral power at each frequency. The solid black lines in Figure 5 represent the 95% credible region and the median of the GWB spectrum. The dashed line is the upper limit on A_{gw} using the purely GW-driven spectrum (i.e., no transition frequency) and the Gaussian amplitude priors from models A and B, respectively. Lastly the thin solid black line is the 95% upper limit on the GWB spectrum from the spectral analysis presented in Section 4.2.1. By inspecting the inferred GWB spectrum one can determine that the data prefer a GWB spectrum that has a definitive transition from GW-driven to environmentally driven within the pulsar timing frequency band for model A, whereas the data does not significantly constrain the shape of the spectrum for model B.

This can be seen further in the joint posterior distributions in Figure 6 in which the probability distributions (blue) for the bend frequency parameter, f_{bend} , and spectral index parameter, κ , are significantly different from the prior distribution (green) for model A and not significantly different for model B. When this same analysis is carried out on the 5 year data set we find that the data can only slightly inform the prior on f_{bend} for model A and gives no information on the other parameters in either model A or B. This, again, indicates that the 9 year data set provides us with much more information about the shape of the GWB strain spectrum.

Finally, we can be more quantitative and apply Bayesian model selection to this problem by computing the Bayes factor between the broken-power-law (Equation (24)) and pure-




Figure 5. Probability density plots of the recovered GWB spectra for models A and B using the broken-power-law model parameterized by $(A_{\text{gw}}, f_{\text{bend}}, \text{and } \kappa)$ as discussed in the text. The thick black lines indicate the 95% credible region and median of the GWB spectrum. The dashed line shows the 95% upper limit on the amplitude of purely GW-driven spectrum using the Gaussian priors on the amplitude from models A and B, respectively. The thin black curve shows the 95% upper limit on the GWB spectrum from the spectral analysis.

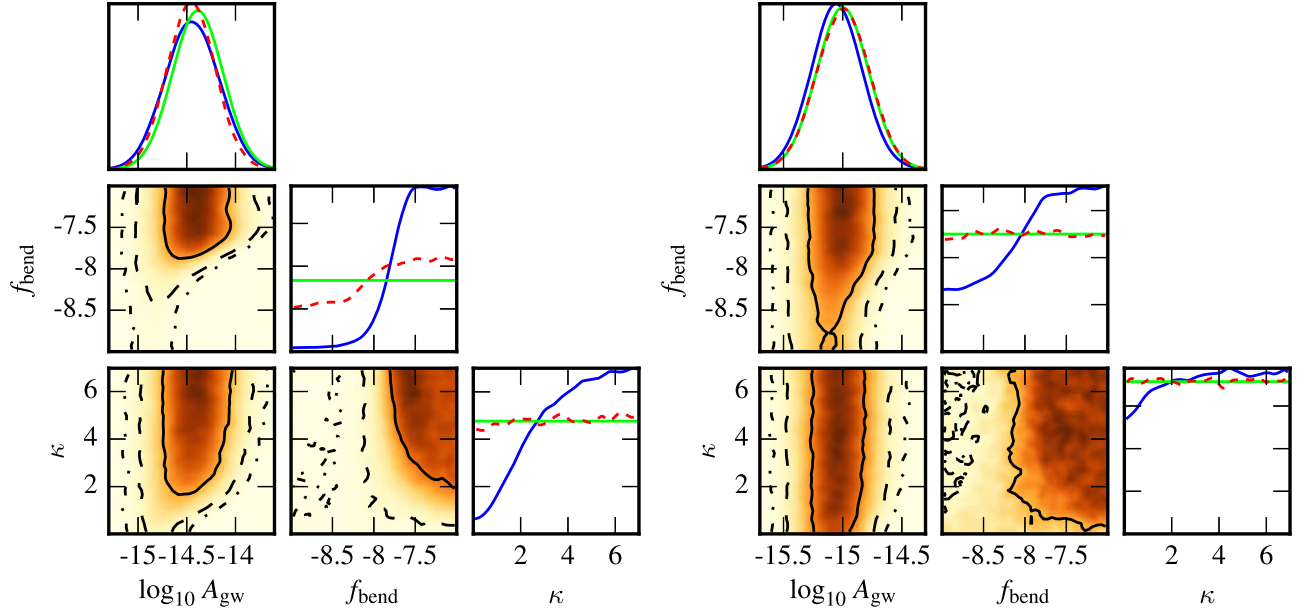


Figure 6. One- and two-dimensional posterior probability density plots of the spectrum model parameters A_{gw} , f_{bend} , and κ . In the one-dimensional plots, we show the posterior probability from the 9 year data set (blue), the 5 year data set (dashed red) and the prior distribution used in both analyses (green). In the two-dimensional plots we show a heat map along with the one (solid), two (dashed), and three (dashed-dotted) sigma credible regions. Model A is on the left and model B is on the right.

power-law (Equation (2)) parameterizations for both models A and B. When this analysis is carried out we arrive at Bayes factors of 22.2 ± 1.1 and 2.23 ± 0.15 for models A and B, respectively. These Bayes factors were computed using parallel tempering and a custom thermodynamic integration implementation (see Section 6.1.2 of Cornish & Littenberg 2015).

This analysis shows, for the first time, that PTAs are entering a regime where even in the case of a non-detection meaningful constraints can be placed on the dynamical history of the SMBHB population. Furthermore, this analysis shows that when placing upper limits to make statements about the full range of astrophysical merger scenarios it is no longer valid to

consider only the classic strain amplitude, but one must instead frame the question in terms of constraining the amplitude *and shape* of the GWB spectrum. In other words, it is only by relaxing the assumptions of a pure power-law spectrum that we can truly assess the validity of merger scenarios such as models A and B. As we have seen in the above analysis and as can be seen clearly in Figure 6, given a model for the SMBHB merger physics (i.e., a prior on A) and discarding the assumption of a purely GW-driven signal (i.e., a broken-power-law model), it is very difficult to rule out any of the GWB amplitude parameter space with any certainty unless one has strong a-priori knowledge on the shape of the spectrum. However, we can begin to place constraints on the environmental coupling

effects that drive the system to the GW-dominated regime for specific SMBHB merger models.

5. LIMITS ON ASTROPHYSICAL MODELS

While the parameter estimation of the previous section is interesting in and of itself, some of the most interesting science available from the NANOGrav data is accessible only by relating these parameters to properties of the source populations. Here we attempt to interpret the phenomenological posteriors on the GW spectrum in terms of black hole–host galaxy relations, environmental effects, or binary-eccentricity by carrying out our analyses in sequence, investigating different effects *separately*. In Section 5.1 we use the results of our power-law analyses of Section 4.2.1 to provide constraints on the parameters of scaling relations between host galaxies and black holes. We then go beyond the assumption of a power-law spectrum in Section 5.2 to investigate how our constraints on the shape of the characteristic strain spectrum from Section 4.2.2 map to constraints on the environment of SMBH binaries. Finally in Section 5.3, we probe the eccentricity of binaries before they entered the PTA band. A robust analysis should take into account all these ingredients simultaneously. By exploring them separately, we can place astrophysical constraints only within the assumptions of the adopted models. As such, the results presented below should be taken as a proof of principle of the astrophysical potential of PTAs, rather than absolute constraints on the individual physical processes.

5.1. Constraints on Host Galaxy–Black Hole Scaling Relations

If the gravitational wave spectrum is assumed to be created by an ensemble of binary SMBHBs that are formed following galaxy mergers (spectral index of $-2/3$), and whose evolution is assumed to be dominated by GW emission throughout the PTA-sensitive waveband, then we can trace the expected binary SMBHB population using observations of galaxy merger rates, the galaxy stellar mass function, and the black hole–host galaxy relation. This is the approach taken in S13, and Ravi et al. (2015). Assuming equal fractional uncertainties in these parameters, the black hole–host galaxy relation will have the largest impact on the predicted level of the GWB. This is due to the much stronger dependence of the GWB on the chirp mass of each source than on the number of sources.

Middleton et al. (2015) shows that it is difficult to extract information from PTA limits without making significant assumptions about the functional dependence on chirp mass, mass ratio, and redshift of the black hole merger rate density. If instead a galaxy merger rate density calculated from observed galaxy parameters is used as a proxy for the black hole merger rate density, then a limit on the GWB can be directly translated into a limit on the input galaxy parameter spaces (J. Simon & S. Burke-Spolaor 2015, in preparation). For this paper, we focus specifically on the scaling relation between host galaxy properties and black hole mass (e.g., Häring & Rix 2004; McConnell & Ma 2013) as it is the observed parameter that is most easily constrained by NANOGrav data. Specifically, we constrain the M_{\bullet} – M_{bulge} relation:

$$\log_{10} M_{\bullet} = \alpha + \beta \log_{10} (M_{\text{bulge}}/10^{11} M_{\odot}), \quad (25)$$

where M_{bulge} is the stellar mass of the galaxy’s bulge. In addition to α and β , observational measurements of this

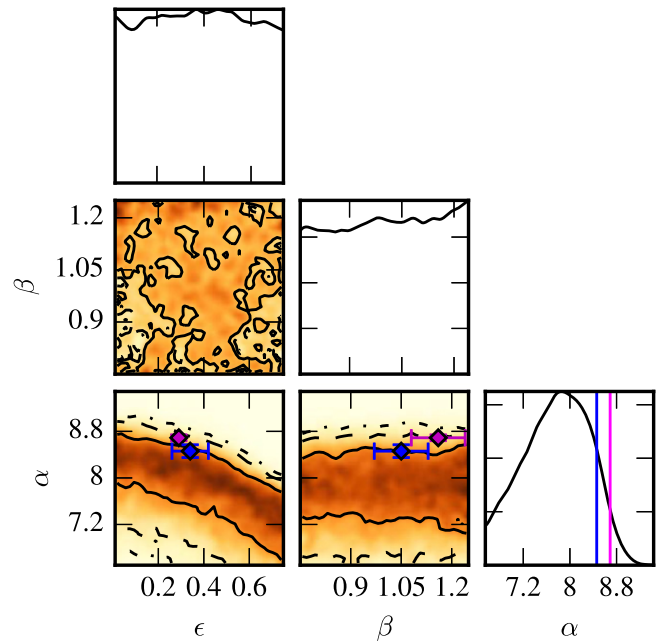


Figure 7. The above plot shows the translation of the marginalized posterior distribution of A_{gw} , Figure 4, into the black hole–host galaxy parameter space, which is characterized by an intercept α , a slope β , and an intrinsic scatter ϵ . Flat priors are used for α , β , and ϵ . β is not informed by the distribution of A_{gw} , while both α and ϵ are, with a limit on α being more strongly set. The curves show the 1, 2, and 3 σ contours. Relevant observational measurements are also shown, with McConnell & Ma (2013) in blue and Kormendy & Ho (2013) in magenta. Since β is not strongly informed by the upper limit, we can set an upper limit in α – ϵ space by marginalizing over β . That upper limit is shown in Figure 8.

relation also fit for ϵ , the intrinsic scatter of individual galaxy measurements around the common α , β trend line. In practice, α and ϵ have the greatest impact on predictions of A_{gw} , and all observational measurements agree with $\beta \approx 1$.

PTAs are most sensitive to binary SMBHBs where both black holes are $\gtrsim 10^8 M_{\odot}$ (e.g., Sesana et al. 2008). Therefore M_{\bullet} – M_{bulge} relations that are derived including the most massive systems are the most relevant to understanding the population in the PTA band. Several recent measurements of the M_{\bullet} – M_{bulge} relation specifically include high-galaxy-mass measurements, e.g., those from Brightest Cluster Galaxies (BCGs). As these fits include the high-mass black holes that we expect to dominate the PTA signals, we take these as the “gold standard” for comparison with PTA limits (Kormendy & Ho 2013; McConnell & Ma 2013, e.g.).

The translation of an upper limit on A_{gw} to the black hole–host galaxy parameter space is calculated as follows:

$$p(\alpha, \beta, \epsilon | \text{PTA}) \propto \int d\theta p(\theta) p(A_{\text{gw}}(\alpha, \beta, \epsilon, \theta) | \text{PTA}), \quad (26)$$

where the posterior of A_{gw} , $p(A_{\text{gw}}(\alpha, \beta, \epsilon, \theta) | \text{PTA})$, is the marginalized posterior distribution of A_{gw} , which is shown in Figure 4; $A_{\text{gw}}(\alpha, \beta, \epsilon, \theta)$ is the prediction of A_{gw} calculated from models similar to S13; θ represents the galaxy stellar mass function and the galaxy merger rate; and $p(\alpha, \beta, \epsilon | \text{PTA})$ is the marginalized posterior distribution of the black hole–host galaxy relation, which is shown in Figure 7. For this analysis, we use two leading measurements of the galaxy stellar mass




Figure 8. The above plot shows the translation of the 95% upper limit on A_{gw} , Figure 4, into the parameter space α - ϵ , which characterizes the black hole–host galaxy relation as described in Equation (25). The parameter space above the line is inconsistent with the power-law analysis of the S13 model. Observational measurements of this parameter space are shown with errorbars.

function, Ilbert et al. (2013) and Tomczak et al. (2014), and two measurements of the galaxy merger rate, Robotham et al. (2014) and Keenan et al. (2014), as the basis for simulating a local population of binary SMBHs. A flat prior is used for α , β , and ϵ , and the posterior on A_{gw} using a uniform prior, as seen in Figure 4, is directly translated into this parameter space. The result of which is shown in Figure 7. β is clearly not informed by a PTA posterior, but the combination of α and ϵ are, with the strongest limit being set on α . Figure 8 shows the translation of our posterior on A_{gw} into α - ϵ parameter space with observational measurements of the parameters from Kormendy & Ho (2013) and McConnell & Ma (2013). Assuming a power-law analysis of the S13 model, as described in J. Simon & S. Burke-Spolaor (2015, in preparation), and using a similar method to Equation (21), we find our data to be 5% and the 10% consistent with the measurements from Kormendy & Ho (2013) and McConnell & Ma (2013), respectively. Potential solutions to the inconsistency of our data with these measurements include: a significant number of black hole binaries do not reach the GW-dominant regime in our assumed timescale (they “stall”); the “classical” assumption of a power-law strain spectrum in the PTA band is incorrect and in fact there is a turn-over in the strain spectrum at lower frequencies (see Section 4.2.2); or that the measured astronomical parameters are not correct for the population of binary SMBHs in the PTA band.

As the possibility for a different strain spectrum curve is discussed in Section 4.2.2, let us explore the potential for “stalling” within the model described so far in this section. Using the galaxy merger rate density as a proxy for the black hole merger rate density implies an assumption that the events occur at a similar cosmological time. If there was significant stalling in the binary black hole population, then these events would be offset in cosmological time by some stalling timescale. We may then add a parameter to account for a stalling timescale to see how such a delay would effect the prediction of A_{gw} . As done in J. Simon & S. Burke-Spolaor (2015, in preparation), we introduce a variable, T_{stall} , which is a measure of this offset in time between the assumed galaxy




Figure 9. By introducing the parameter T_{stall} , as described in J. Simon & S. Burke-Spolaor (2015, in preparation), we can start to explore the inconsistency of our upper limit with power-law models for the GWB. In the above plot, we allow T_{stall} to vary while using the M_{\bullet} - M_{bulge} relation Kormendy & Ho (2013). The probability of T_{stall} is a direct translation of the posterior on A_{gw} from Figure 4. The vertical line is the 95% lower limit on T_{stall} , which we find to be 1.1 Gyr. While this is not sufficiently constraining to make meaningful astrophysical statements, this parameter may be useful for future PTA upper limits.

merger rate density and the black hole merger rate density used to model a GWB. Figure 9 shows the translation of the posterior distribution of A_{gw} , Figure 4, into a probability distribution of T_{stall} using the Kormendy & Ho (2013) measurements of the M_{\bullet} - M_{bulge} relation. Using this we set a 95% lower limit on T_{stall} of 1.1 Gyr. This value is not constraining enough to make statements on what evolutionary processes drive the binary through the “last parsec,” where binary inspiral has been suggested to stall (Begelman et al. 1980). However, this parameter may be useful for future PTA upper limits as a probe of the level of stalling in the binary black hole population.

5.2. Constraints on Binary Environmental Influences

The cores of galactic merger remnants can harbor stars with little angular momentum and almost radial trajectories which intersect the central galactic region (centrophilic orbits). These stars can undergo three-body interactions with the resident SMBHB, causing the stars to be ejected, which results in energy and angular momentum being extracted from the black hole system, and leading to binary hardening (Quinlan 1996).³⁸ Additionally, the formation of circumbinary gaseous disks can lead to interactions which extract energy and angular momentum from the binary orbit, driving it toward smaller orbital separations (Ivanov et al. 1999). We expect that, in the type of gas-poor galaxies which dominate the nanohertz GW signal, hardening from stellar scattering will dominate over circumbinary interactions, but we consider both in the following. We begin with a discussion of how these environmental mechanisms drive the evolution of the binary, then provide constraints on the frequency at which the characteristic strain spectrum exhibits a turnover from the analysis in Section 4.2.2. We finish by mapping these

³⁸ We assume that all galactic merger remnants maintain the same mass density of core stars throughout the binary merger. The subtleties of loss-cone replenishing impact the evolution of the binary and of the central density profile within a factor of approximately two (a few at most), as shown by Sesana & Khan (2015) and Vasiliev et al. (2015).

frequencies to constraints on the astrophysical environment of SMBH binaries emitting GWs in the nanohertz band.

5.2.1. Environmentally Driven Orbital Evolution

We use the formalism of Quinlan (1996) to define the evolution of a (circular) binary due to three-body stellar scattering events, where

$$\frac{d}{dt}\left(\frac{1}{a}\right) = \frac{G\rho H}{\sigma}, \quad (27)$$

where a is the binary orbital semimajor axis, ρ is the mass density of galactic core stars, H is a dimensionless hardening rate which takes a value of ~ 15 , and σ is the velocity dispersion of core stars. Using Kepler's third law, we can rearrange this equation to solve for the rate of orbital frequency evolution, which gives $df/dt \propto f^{1/3}$. Since the binary orbital evolution will be due to a combination of environmental influences and GW emission, df/dt is actually a sum over all mechanisms, as in Equation (23). We know that the rate of frequency evolution due to GW emission is $df/dt \propto f^{11/3}$ (Peters & Mathews 1963), hence, in the language of the parametrized spectrum model in this paper given in Equation (24), $\kappa = 10/3$ for binary hardening by three-body stellar scattering.

Likewise, the evolution of a circular binary due to circumbinary disk interaction is modeled within the α -disk formalism (Ivanov et al. 1999; Haiman et al. 2009; Sesana 2013b) as (Sesana 2013b)

$$\frac{da}{dt} = -\frac{2\dot{M}_1}{\mu}(aa_0)^{1/2}, \quad (28)$$

where \dot{M}_1 is the accretion rate onto the primary black hole, μ is the binary reduced mass, and a_0 is a characteristic orbital separation at which the enclosed disk mass equals the mass of the secondary black hole. The latter can be expressed as (Ivanov et al. 1999)

$$a_0 = 3 \times 10^3 \left(\frac{\alpha}{10^{-2}}\right)^{4/7} \left(\frac{M_2}{10^6 M_\odot}\right)^{5/7} \times \left(\frac{M_1}{10^8 M_\odot}\right)^{-11/7} \left(100 \frac{\dot{M}_1}{\dot{M}_E}\right)^{-3/7} r_g, \quad (29)$$

where α is a disk viscosity parameter, $M_{1,2}$ are the binary black hole masses; $\dot{M}_E = 4\pi GM_1/c\kappa_T$ is the Eddington accretion rate of the primary (κ_T is the Thompson opacity coefficient); and $r_g = 2GM_1/c^2$ is the Schwarzschild radius of the primary.

As in the stellar scattering case, we can rearrange this equation to determine the orbital frequency evolution. In this model, $df/dt \propto f^{4/3}$, and so $\kappa = 7/3$ for α -disk binary interactions.

5.2.2. Constraints on Spectral Turnover Frequency

We define the spectral turnover frequency in the obvious way to mean the frequency at which the characteristic strain spectrum exhibits a change in slope. If the low-frequency slope is positive, this will correspond to the point at which the spectrum is maximized. One must note that setting $f = f_{\text{bend}}$ in Equation (24) does not maximize the characteristic strain spectrum. Rather the turnover frequency will be a function of

both f_{bend} and κ :

$$f_{\text{turn}} = f_{\text{bend}} \left(\frac{3\kappa}{4} - 1\right)^{1/\kappa}. \quad (30)$$

We can combine our measurements of f_{bend} and κ from the analysis in Section 4.2.2 to compute the probability distribution of spectral turnover frequencies. We find that placing numerical constraints on f_{bend} is difficult as the posterior is heavily influenced by the prior, namely the upper and lower bounds for the uniform priors used in this analysis. In the following analysis we set the lower bound on f_{bend} by requiring that the power at $f = 1/T_{\text{span}}$ differs by no more than 10% from a pure-power-law for any value of κ . By using this prior, we ensure that we can recover a pure-power-law spectrum (in our frequency range) for any value of κ . The upper bound of f_{bend} is chosen based on the specific environmental coupling mechanism we are considering.

We emphasize that the probability distributions of f_{bend} , κ , and f_{turn} are not distributions of these parameters over the SMBHB population. Rather our posterior distribution is illustrating the spread of our beliefs in the measurement of the single f_{bend} and κ model parameters. Investigations into the intrinsic spread (rather than our measurement spread) of f_{bend} should be considered in future studies as a means to investigate the astrophysical distribution of galactic center environmental properties.

5.2.3. Constraints on Environmental Parameters

We can now extract astrophysical constraints from our constraints on the transition and spectral turnover frequencies. By equating the rate of orbital evolution, da/dt , due to environmental mechanisms and GW emission, we can deduce the characteristic transition frequency, f_{bend} , between these influences. We first consider stellar-scattering, for which the transition frequency is given by

$$f_{\text{bend}} = 3.13 \text{ nHz} \times \sigma_{200}^{-3/10} \rho_3^{3/10} H_{15}^{3/10} M_8^{-2/5} q_r^{-3/10} \quad (31)$$

$$= 3.25 \text{ nHz} \times \rho_3^{3/10} H_{15}^{3/10} M_8^{-23/50} q_r^{-3/10} \quad (32)$$

where M is the total binary mass, $M_8 \equiv M/(10^8 M_\odot)$, $q = M_2/M_1$, $q_r = q/(1+q)^2$, $\sigma_{200} \equiv \sigma/(200 \text{ km s}^{-1})$, $\rho_3 \equiv \rho/(10^3 M_\odot \text{ pc}^{-3})$, and $H_{15} \equiv H/15$. In the second line we have used the M - σ relation (McConnell et al. 2011) to replace velocity dispersion, where $M_8 \approx 1.9 \sigma_{200}^5$. It follows from the McConnell et al. (2011) M - σ relation that the velocity dispersion term on the first line of Equation (31) has the mass scaling $\sigma_{200}^{-3/10} \propto M_8^{-3/50}$, which is a small modification to the exponent of the other mass term, $M_8^{-2/5}$. Hence, f_{bend} has a relatively weak dependence on the mass scaling of σ . Nevertheless we can express the transition frequency in terms of variables of a parametrized M - σ relation, where $\log_{10}(M/M_\odot) = a + b \log_{10} \sigma_{200}$, such that $M_8 = 10^{a-8} \sigma_{200}^b$, and

$$f_{\text{bend}} = 3.13 \text{ nHz} \times 10^{3(a-8)/10b} \rho_3^{3/10} H_{15}^{3/10} M_8^{-(3+4b)/10b} q_r^{-3/10}. \quad (33)$$

Finally, the weak scaling of f_{bend} with H , and the $\lesssim 20\%$ deviations of this parameter away from 15 seen in numerical

scattering experiments (Sesana & Khan 2015) justifies our keeping this parameter fixed at its fiducial value of $H_{15} = 1$. Astrophysical estimates on ρ are quite uncertain with estimated values around $10\text{--}10^4 M_\odot \text{pc}^{-3}$ for typical environments (Dotti et al. 2007). The variation of estimates over several orders of magnitude is why we choose to investigate only ρ in our stellar-scattering constraints.

The equivalent transition frequency for α -disk interaction is

$$f_{\text{bend}} \approx 0.144 \text{ nHz } M^{-17/14} q_r^{-6/7} \dot{M}_1^{3/7} a_0^{3/14}, \quad (34)$$

where a_0 is as previously defined. We adopt the fiducial value of the disk viscosity parameter $\alpha \sim 10^{-2}$ used in Ivanov et al. (1999). The very weak dependence of the bend frequency on this parameter, $f_{\text{bend}} \propto \alpha^{6/49}$ will significantly dampen the influence of any deviations from this fixed value. Hence, in our constraints on the influence of circumbinary disk interactions, we only vary the accretion rate of gas onto the primary black hole, \dot{M}_1 , of which estimates in the literature vary over several orders of magnitude—typically $10^{-3} M_\odot \text{yr}^{-1}$ – $1 M_\odot \text{yr}^{-1}$, see, e.g., Di Matteo et al. (2001), Armitage & Natarajan (2002), Goicovic et al. (2015).

Equations (31)–(34) indicate the GW frequency at which a *single* circular binary will transition between being environmentally driven and being GW driven. Of course, f_{bend} is not the quantity that can be extracted from a spectral analysis— f_{turn} is what we can measure. Our analysis of the stochastic GWB has provided us with constraints on the characteristic transition frequency of the entire population (given a certain dynamical mechanism) which can be used to infer where the spectrum turns over. It is important to separate the interpretation of these two frequencies, since f_{bend} tells us the frequency above which radiation reaction dominates binary orbital evolution, whereas f_{turn} tells us directly about the shape of the background strain spectrum. Hence, our path to providing constraints on environmental parameters requires us to numerically construct characteristic strain spectra for populations of SMBH binaries in contact with their environment. We can then construct numerical mappings between the environmental parameters of interest (core stellar mass density, ρ , and primary black hole accretion rate, \dot{M}_1) and the turnover of the spectrum. We use the formalism of Phinney (2001) via Equation (22), where the differential comoving number density of merging binaries per redshift and component masses is constructed as in McWilliams et al. (2014):

$$\frac{d^3 n}{dz dM_1 dM_2} \propto P(z) \phi(M_1) \phi(M_2), \quad (35)$$

where $P(z)$ encapsulates redshift dependent factors, and $\phi(M_{1,2})$ is the number density of black holes of a certain mass, which is given by a (redshift dependent) Schechter function modified at the high-mass end by a lognormal component to accommodate recent high mass BCG discoveries (Lin et al. 2010).

The combined influence on the binary orbital evolution of GW emission and environmental couplings is modeled with the sum in Equation (23), where either stellar scattering *or* disk interactions are included i.e., we consider one mechanism at a time. By considering all binary environments to have the same ρ or \dot{M}_1 , we can determine the spectral turnovers from the numerically computed strain spectra, iterating over many

values of these environmental parameters to deduce a mapping. We draw binary systems from the ranges $z \in [0, 1]$, $M_1 \in [10^7, 10^{10}] M_\odot$, and $q \in [0.1, 1]$. The results for our fiducial assumptions are shown in the top panel of Figure 10, with the stellar density required to give a certain turnover frequency shown in the left panel, and the primary accretion rate required to give a certain turnover frequency shown in the right panel.

In the lower panels of Figure 10 we plot the posterior distributions of the stellar density, ρ , for stellar hardening, and mass accretion rate, \dot{M}_1 for circumbinary disk interaction. In this analysis we perform the Bayesian parameter estimation for fixed values of κ corresponding to the appropriate values for stellar hardening ($\kappa = 10/3$) and circumbinary disk interaction ($\kappa = 7/3$). These posteriors are constructed by first converting the marginalized distributions on f_{bend} to a distribution for f_{turn} via Equation (30) and then using the empirical mapping to convert f_{turn} to the appropriate astrophysical parameter. Again, we do not place numerical confidence limits on ρ or \dot{M}_1 since the data does not constrain the prior distribution at large values. Nonetheless, from inspection of Figure 10 we see that the MOP14 model heavily prefers $\rho \gtrsim 10^4 M_\odot \text{pc}^{-3}$ and $\dot{M}_1 \gtrsim 10^{-1} M_\odot \text{yr}^{-1}$, while the S13 model is largely unconstraining for both mechanisms. Typical densities of massive elliptical galaxies at the MBH influence radius is $\sim 10^3 M_\odot \text{pc}^{-3}$, making the MOP14 model hard to reconcile with observations, even if we consider that massive ellipticals were likely a factor 2–3 more compact at $z = 1$ (A. Sesana 2016, private communication). Our results approach the upper end (for the MOP14 prior) of the expected range of \dot{M} , $10^{-3} M_\odot \text{yr}^{-1}$ – $1 M_\odot \text{yr}^{-1}$, observed in the local universe and predicted via simulations, see, e.g., Di Matteo et al. (2001), Armitage & Natarajan (2002), Goicovic et al. (2015). Furthermore, Dotti et al. (2015) predict that $\dot{M}_1 \ll 10^{-1} M_\odot \text{yr}^{-1}$ for BH masses of $10^9 M_\odot$ and redshifts $z < 1$; however, these are average accretion rates, and short, episodic accretion triggered by galaxy mergers could occur at a higher rate.

We go beyond the fiducial assumptions for the case of stellar hardening since it is the most likely environmental coupling mechanism for SMBH binaries. When we increase the low-mass cutoff of systems which contribute to the characteristic strain budget this further constrains the stellar mass density. This is seen most easily in Equation (31), where one must raise the stellar mass density to match a corresponding increase in binary mass so that the transition frequency is maintained. Furthermore, modeling the distribution of black hole masses in Equation (35) without the lognormal component or redshift evolution will increase the contribution of lower mass binaries to the GW strain budget, leading to smaller stellar mass density constraints than reported in Figure 10. Varying the normalization, a , and exponent, b , of the M – σ relation such that $a \in [7, 9]$ and $b \in [4, 6]$ has very little impact on the environmental constraints.

5.3. Constraints on SMBHB Population Eccentricity

It is not only the astrophysical environment of SMBH binaries that can induce a bend in the characteristic strain spectrum. Binaries with non-zero eccentricity emit GWs at a spectrum of harmonics of the orbital frequency. The cumulative effect over the entire population can lead to a depletion of the low frequency strain spectrum (Enoki et al. 2007; Sesana 2013b; Ravi et al. 2014; Huerta et al. 2015), and a turnover




Figure 10. Top: empirical mapping from f_{turn} to ρ (left) and \dot{M}_1 (right). Bottom: posterior distributions for the mass density of stars in the galactic core (left) and the accretion rate of the primary black hole from a circumbinary disk (right). These distributions are constructed by first converting the marginalized distribution of f_{bend} to a distribution of f_{turn} via Equation (30), and then using the empirical mapping described in the text to convert from f_{turn} to the astrophysical quantities ρ and \dot{M}_1 , respectively.

whose shape can be captured with the parametrized spectrum model employed in this paper. Hence, we can use our f_{turn} posterior from the marginalization of f_{bend} over all κ to deduce constraints on the eccentricity of binaries at some reference orbital separation. Our approach follows from the previous astrophysics constraints, where we build populations and strain spectra which have varying eccentricities at a fixed semimajor axis of 0.01 pc, then construct a relationship between this eccentricity and the spectral turnover. An important modeling assumption we make here is that binaries are (and always have been) driven entirely by GW emission. This factors into how we model df/dt and how we evolve the binary eccentricity, where we assume binaries could have eccentricities arbitrarily close to 1 in the past.

We construct eccentric populations and corresponding strain spectra using the formalism of Huerta et al. (2015). The resulting relationship between the spectral turnover frequency and the eccentricity of all binaries at a semimajor axis of 0.01 pc is shown in the top panel of Figure 11, along with the corresponding eccentricity posteriors from each amplitude prior in the bottom panel. The high turnover frequency obtained with the MOP14 prior leads to an eccentricity posterior distribution that largely favors $e_0 \gtrsim 0.7$ while the S13 prior leads to an eccentricity posterior that is consistent with smaller eccentricities, more weakly favoring $e_0 \gtrsim 0.5$. We emphasize that, while these eccentricities seem rather large, it is well established that binaries evolving in stellar environments tend to increase their eccentricity (Quinlan 1996). It is therefore likely that most binaries can get to $e \sim 0.5$ – 0.7 along their evolution (see tracks in Sesana 2010). The eccentricity growth rate is generally larger for smaller binary mass ratios, and for larger initial eccentricities. The latter is indeed a crucial parameter; if, following galaxy mergers, the SMBHB already has a significant ($e \gtrsim 0.5$) eccentricity at the moment of formation, the subsequent evolution will almost certainly drive it to $e > 0.9$. Given that the SMBHB eccentricity at formation is hard to determine (Hemsendorf et al. 2002; Aarseth 2003; Amaro-Seoane et al. 2009; Berentzen et al. 2009), it is impossible to draw astrophysical conclusions from the

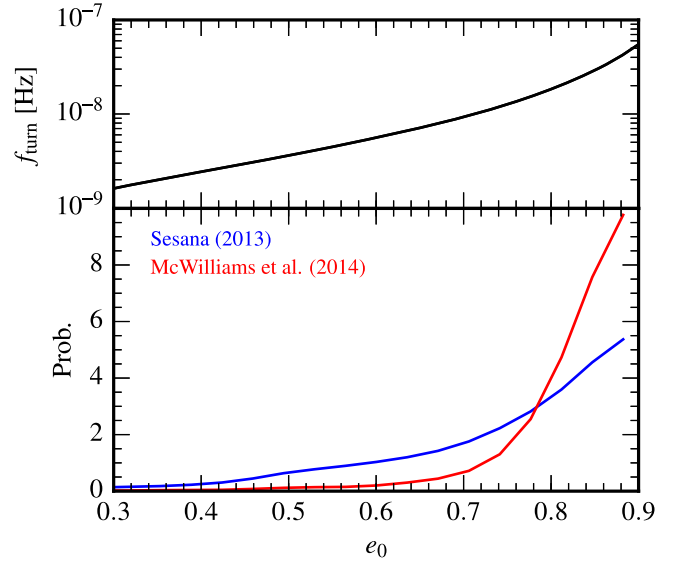


Figure 11. Same as Figure 10, except now we display the empirical mapping (top) and posterior distribution (bottom) for the eccentricity of SMBH binaries when they had a semimajor axis of 0.01 pc.

constraints above. Furthermore, in reality the history of a binary’s eccentricity will see phases of growth and circularization depending upon the interplay of environmental factors and GW emission (e.g., Sesana 2010; Kocsis & Sesana 2011). This should be considered in future analyses.

6. LIMITS ON COSMOLOGICAL MODELS

SMBHB mergers are not the only possible source of a GWB in the pulsar timing band. In this section, we use our power-law-spectrum results, shown in Figure 3, to constrain the fractional energy density of the universe in relic (or primordial) GWs, along with stringent limits on the tension of cosmic strings. Note that many observable cosmological quantities, including those associated with CMB anisotropies—such as the

amplitude and spectrum of primordial perturbations—are dimensionless, see, e.g., Mukhanov et al. (1992), Antoniadis & Patil (2015) for an overview. However, when trying to infer an absolute energy scale for inflation, it becomes necessary to choose units, with the natural choice being Planck units.

6.1. Relic GWs

Quantum fluctuations of the gravitational field in the early universe, amplified by an inflationary phase, are expected to produce a stochastic relic GWB, see, e.g., Grishchuk (1976, 1977), Starobinsky (1980), Linde (1982). Observations of this background would provide a unique insight into the highly curved spacetime of the early universe at less than 10^{-32} s after the Big Bang and at energy scales of 10^{16} GeV, when quantum mechanics and general relativity should reconcile, BICEP2/Keck et al. (2015), Ade et al. (2014). This background is expected to produce a characteristic signature in the polarization of the CMB radiation, as well as CMB temperature anisotropies (Mukhanov et al. 1992; Grishchuk 2005). In the context of PTAs, the amplitude of the relic GWB is of astrophysical and cosmological interest due to the fact that it intrinsically depends on the equation of state of the early universe, w , and thus the Hubble constant in the inflationary stage H_* , as well as the tensor-to-scalar ratio r , see, e.g., Zhao (2011), Zhao et al. (2013), Lasky et al. (2015).

Specifically, the spectral index of the stochastic GWB, α , is related to the equation of state of the early universe, w , via

$$\alpha = \frac{n_t}{2} - \frac{2}{3w + 1}, \quad (36)$$

where n_t is the spectral index of the primordial power spectrum, usually set to 0. In current hot Big Bang models, $w = 1/3$ and $n_t = 0$, thus $\alpha = -1$. We therefore fix the spectral index to -1 and apply the Bayesian analysis methods described in Section 4.2, to the 9 year NANOGrav data set. We obtain a 95% upper limit of on the amplitude of the relic GWB of

$$A_{\text{relic}}^{95\%} = 8.1 \times 10^{-16}, \quad (37)$$

assuming a power spectrum for the characteristic strain with $\alpha = -1$ at a reference frequency of $f = \text{yr}^{-1}$. This then constrains the GW energy density content per unit logarithmic frequency, divided by the critical energy density, ρ_c , to close the universe:

$$\Omega_{\text{gw}}(f) = \frac{1}{\rho_c} \frac{d\rho_{\text{gw}}}{d \log f} = \frac{2\pi^2}{3H_0^2} f^2 h_c^2(f), \quad (38)$$

where f is the frequency, $\rho_c = 8\pi/(3H_0^2)$, H_0 is the Hubble expansion rate, and ρ_{gw} is the total energy density in GWs (Allen & Romano 1999; Maggiore 2000). The NANOGrav limit is therefore

$$\Omega_{\text{gw}}(f)h^2 < 4.2 \times 10^{-10}, \quad (39)$$

in a radiation-dominated universe with equation of state of $w = 1/3$. This new limit is a factor of 2.9 better than the previous best limit of $\Omega_{\text{gw}}(f)h^2 = 1.2 \times 10^{-9}$ from LTM15.

Although a radiation-dominated era is usually assumed to follow inflation in the hot Big Bang paradigm, there is currently no evidence to show this held before Big Bang Nucleosynthesis (BBN). In fact, the existence of a reheating stage or the

Table 4
Spectral Bend and Turnover Frequency Constraints

	McWilliams et al. (2014)		Sesana (2013a) Prior	
	f_{bend} (nHz)	f_{turn} (nHz)	f_{bend} (nHz)	f_{turn} (nHz)
Marginalized κ	13.6	12.6	5.29	4.98
$\kappa = 10/3$ (stars)	8.75	9.88	4.30	4.85
$\kappa = 7/3$ (gas)	8.03	7.10	4.34	3.83

Note. Limits correspond to the 95% lower limit deduced from posterior probability distributions. Turnover frequencies are computed using the subset of posterior samples which have $\kappa > 4/3$.

Table 5
Values for $\Omega_{\text{gw}}(f)h^2$ Reported at 1 yr^{-1}

EoS, ω	Spectral Index, α	$A_{\text{relic}}^{95\%}$	$\Omega_{\text{gw}}(f)h^2$	H_*/m_{Pl}
0	-2	1.2×10^{-16}	8.6×10^{-12}	5.4
1/3	-1	8.1×10^{-16}	4.2×10^{-10}	1.6×10^{-2}
1	-1/2	2.0×10^{-15}	2.5×10^{-9}	8.2×10^{-4}

Note. Cosmological parameters used for H_* are $h = 0.702$, $T_{\text{CMB}} = 2.76 \text{ K}$, $\Omega_{\Lambda} = 0.725$, $\Omega_m = 0.275$, and $z_{\text{eq}} = 3454$. Values of H_* are in given in multiples of the Planck mass $m_{\text{Pl}} \equiv 1/\sqrt{G}$.

existence of a cosmic phase transition both violate this assumption (Grishchuk 2001, pp. 167–192; Watanabe & Komatsu 2006; Zhao 2011). For completeness, we now allow for other equations of state of the early universe before the BBN stage. The same analysis can be repeated assuming different equations of state for the early universe: a matter-dominated universe would have $w = 0$ and therefore by Equation (36), $\alpha = -2$. If the universe were instead dominated by the kinetic energy of the inflaton, then $w = 1$ and $\alpha = -1/2$.

Finally, we place limits on the Hubble parameter, H_* in the inflationary stage using methods developed in Zhao (2011). The interested reader is referred to Section 2 of Antoniadis & Patil (2015) for a comprehensive introduction. In Zhao (2011), the authors introduce a way of translating the upper limit on a primordial GWB to a constraint on H_* . Using typical cosmological parameters $h = 0.702$, $T_{\text{CMB}} = 2.76 \text{ K}$, $\Omega_{\Lambda} = 0.725$, $\Omega_m = 0.275$, and $z_{\text{eq}} = 3454$, they obtain the following relation:

$$\log_{10} \Omega_{\text{gw}}(k_*) = 1.25 - \frac{13.48}{3w + 1} + 2 \log_{10} \left(\frac{H_*}{m_{\text{Pl}}} \right), \quad (40)$$

where $k_* = 2\pi f_*$ and is reported at a reference frequency of $f_* = 1 \text{ yr}^{-1}$ and $m_{\text{Pl}} \equiv 1/\sqrt{G}$ is the Planck mass. Using Equation (40), we can then limit H_* for a fixed equation of state. For example, using the limit on $\Omega_{\text{gw}}h^2 = 4.2 \times 10^{-10}$ for $w = 1/3$, see Table 4, one can place a limit $H_* = 1.6 \times 10^{-2} m_{\text{Pl}}$. Results for $w = 0, 1$ are in Table 4.

In LTM15, the limit on a primordial GWB with $w = 1/3$ and $n_t = 0$ is $\Omega_{\text{gw}}(f)h^2 = 1.2 \times 10^{-9}$, resulting in $H_* = 2.74 \times 10^{-2} m_{\text{Pl}}$. One can see that the NANOGrav limit on H_* is an improvement of 1.7 over the EPTA, and the most stringent limit to date.




Figure 12. Left: cosmic string constraints in terms of string tension $G\mu$ and reconnection probability p using the results of recent cosmic string simulations described in Blanco-Pillado et al. (2014). Right: cosmic string tension $G\mu$ vs. loop size parameterized by α_{cs} using the model described in LTM15. The shaded area is ruled out by our GW upper limit in both panels.

6.2. GW Background from Cosmic (Super)Strings

Cosmic strings are linear topological defects that can be produced in the early universe via phase transitions (Kibble 1976; Vilenkin 1981; Vilenkin & Shellard 2000). Cosmic superstrings are fundamental strings stretched to cosmological scales and arise in a wide range of string-theory-inspired cosmological scenarios (Sarangi & Tye 2002; Jones et al. 2003; Copeland et al. 2004). Cosmic (super)strings produce a stochastic background of GWs as well as individual bursts (Damour & Vilenkin 2001, 2005; Siemens et al. 2006, 2007; Ölmez et al. 2010).

Our limits on the amplitude of the stochastic background can also be used to constrain the parameter space of cosmic (super) strings. Recent simulations (Blanco-Pillado et al. 2014) have shown that cosmic (super)string loop densities are dominated by loops that formed at scales comparable to the Hubble size at the time of formation, even though only about 10% of loops are formed with such large sizes. We use the loop distributions derived by Blanco-Pillado et al. (2014), specifically Equations (63), (65), and (67) of that reference with loop size $\alpha_{cs} = 0.05$, together with the techniques described in Ölmez et al. (2010) to compute the stochastic background produced by cosmic string cusps. The cosmological parameters we used are taken from the Planck 2015 data (Ade et al. 2015). In this case the relevant parameters are the string tension $G\mu$ and the reconnection probability p . We explore this parameter space and exclude regions where the cosmic (super)string network would have resulted in a stochastic background amplitude larger than that ruled out by our measurements. The left panel of Figure 12 shows the results of our analysis. On the y -axis we show the reconnection probability and on the x -axis the string tension. The gray shaded area shows the region of cosmic string parameter space that is ruled out. Note that for $p = 1$ our data only allow for cosmic (super)strings with tensions $G\mu < 1.3 \times 10^{-10}$.

Recently LTM15 presented a comprehensive and fully general overview of cosmic string limits from the EPTA, and found a conservative limit on the string tension to be $G\mu < 1.3 \times 10^{-7}$. The limit is conservative in the sense that it is found by considering a wide range of loop sizes and taking the upper limit to be the largest possible value of $G\mu$ consistent with the data. The limit was identical to that set by the Planck

Collaboration, combining Planck and high- l cosmic microwave background data with Atacama Cosmology Telescope (ACT) and the South Pole Telescope (SPT), cf. Planck Collaboration et al. (2014). While the calculation in LTM15 was not carried out explicitly for the Blanco-Pillado et al. (2014) simulations we can use their published limit on $\Omega_{gw}(f)h^2 = 1.2 \times 10^{-9}$ for cosmic strings to place a limit of $G\mu < 8.6 \times 10^{-10}$. Our limit for this model is therefore roughly a factor of 6.6 times more constraining than the inferred previous limit. Using the same analysis developed by the EPTA, (Sanidas et al. 2012, 2013), we compute the upper limit on the string tension $G\mu$ as a function of loop size α_{cs} as shown in the right panel of Figure 12. Our conservative limit on cosmic string tension using this range of cosmic string models is $G\mu < 3.3 \times 10^{-8}$, a factor of four better than both the combined Planck, ACT, SPT limit and the EPTA limit.

7. SUMMARY AND CONCLUSIONS

This paper reports on the search for an isotropic stochastic GW background in NANOGrav’s 9 year data set. We do not find positive statistical evidence for the presence of such a signal. Following up on a series of earlier results by the three PTAs, we report new upper limits on the amplitude of backgrounds described by power-law spectra:

1. For an astrophysical background of SMBH binaries (corresponding to a timing-residual spectral density with exponent $\gamma = 13/3$), we find a 95% confidence limit $A_{gw} < 1.5 \times 10^{-15}$, five times more constraining than the analogous limit for NANOGrav’s 5 year data set (DFG13). Under the assumption of purely GW-driven evolution, leading to an unbroken $\gamma = 13/3$ power law, we compute the probability that our constraint is consistent with the MOP14 and S13/RWS14 theoretical predictions for A_{gw} as 0.8% and 20%, respectively, essentially ruling out the MOP14 model and placing the S13/RWS14 model in tension with our data (Section 4.2.1).
2. We check the consistency of our limit with previously reported scaling relations between SMBH mass and galactic bulge mass, adopting fiducial estimates for galaxy merger rates and the stellar mass function. Under the assumption of circular GW-driven binaries, we find

the scaling relations reported in [Korndy & Ho \(2013\)](#) and [McConnell & Ma \(2013\)](#) to be inconsistent with our data at the 95% and 90% level, respectively (Section 5.1).

3. We also perform an optimal-statistic (cross-correlation) analysis, and find limits that are 5.4 and 1.5 times more constraining than the analogous [NG5](#) and [LTM15](#) results. The cross-correlation S/N is 1.5, indicating that there is little evidence for inter-pulsar residual correlations induced by GWs (Section 4.1).
4. We extend the power-law background search to generic spectral indices, and place stringent limits on the energy density of relic GWs, $\Omega_{\text{gw}}(f = 1 \text{ yr}^{-1})h^2 < 4.2 \times 10^{-10}$, for a $w = 1/3$ early-universe equation of state. From this we obtain limits on the Hubble parameter during inflation, $H_* = 1.6 \times 10^{-2} m_{\text{pl}}$ (Sections 4.2.1, 6.1).
5. We place the most stringent limits to date on a GW background generated by a network of cosmic strings, $\Omega_{\text{gw}}(f = 1 \text{ yr}^{-1})h^2 < 2.2 \times 10^{-10}$, which translates into a conservative upper limit on cosmic string tension $G\mu < 3.3 \times 10^{-8}$, using the model presented in [LTM15](#). This is a factor of four better than both the combined Planck, ACT, SPT limit and the EPTA limit. Using the recent models of [Blanco-Pillado et al. \(2014\)](#) we find $G\mu < 1.3 \times 10^{-10}$, a factor of 6.6 times more constraining than an identical analysis using the EPTA limit (Sections 4.2.1, and 6.2).

We further probe the interface between PTA observations and SMBH-binary population estimates by analyzing the 9 year data set in terms of a GW background described by an inflected power law (Equation (24)). We derive joint posteriors for the spectral parameters (the amplitude A_{gw} , the inflection frequency f_{bend} , and the shape exponent γ) assuming A_{gw} priors from [MOP14](#) and [S13/RWS14](#) (see Figure 5). For both priors (but more so for [MOP14](#)), the inflected power-law model is preferred to an unbroken power law. The f_{bend} posterior can be used to infer astrophysical information about the effects that may reduce GW power at lower frequencies, such as the initial eccentricity of SMBH binaries, and the environmental influences of stars and gas in galactic nuclei. To summarize:

1. We find that the data prefer an inflected spectrum over a pure power law, with Bayes factors of ~ 22 and 2.2 for the [MOP14](#) and [S13](#) amplitude priors, respectively. The same analysis, run on the 5 year DFG13 data set, provides little to no information about the shape of the GWB spectrum (Section 4.2.1).
2. Under several simplifying assumptions, we map the posterior distribution of f_{bend} into posterior distributions for the nuclear stellar mass density ρ (which modulates the strength of binary frequency evolution by stellar scattering, cf. Section 5.2), the SMBH mass accretion rate from circumbinary disks \dot{M}_1 (which can be linked to binary frequency evolution by interactions with gas, cf. Section 5.2), and the initial value of SMBH-binary orbital eccentricity e_0 (which distributes GW power to higher frequency harmonics, cf. Section 5.3).
3. We find that the [MOP14](#) model is in tension with observed and/or predicted values of ρ and \dot{M}_1 and requires a large initial eccentricity (i.e., $e_0 > 0.7$) to be consistent with our residual data. Furthermore, we find that while the [S13](#) model marginally prefers an inflected

spectral model, the inferred astrophysical parameters and eccentricities indicate that this model is still consistent with the assumption of GW-driven circular binaries.

For the last decade (and longer), the three PTA consortia have been engaged in an accelerating race toward higher GW sensitivities—the analysis presented in this paper represents the latest stage of the race, but not the last. While our investigation cannot claim the ultimate prize, a positive detection, it is the first to use information about the amplitude *and* shape of GW background to make concrete (if limited) astrophysical statements about the dynamics and environments of SMBH binaries.

Author contributions. An alphabetical-order author list was used for this paper in recognition of the fact that a large, decade-timescale project such as NANOGrav is necessarily the result of the work of many people. All authors contributed to the activities of the NANOGrav collaboration leading to the work presented here, and reviewed the manuscript, text, and figures prior to the paper’s submission. Additional specific contributions to this paper are as follows (also listed alphabetically). Z.A., K.C., P.B.D., T.D., R.D.F., E.F., M.E.G., G.J., M.J., M.T.L., L.L., M.A.M., D.J.N., T.T.P., S.M.R., I.H.S., K.S., J.K.S., and W.W.Z. made observations for this project and developed timing models. J.A.E. coordinated the writing and submission of the paper. J.A.E., Rv.H., and C.M.F.M. led this search by directly running the analysis pipelines. N.J.C., J.A.E., Rv.H., L.S., S.R.T., and M.V. designed and tuned the statistical analysis. N.J.C., J.A.E., S.T.M., C.M.F.M., L.S., S.A.S., A.S., X.S., J.S., S.B.-S., S.R.T., and M.V. developed the interpretation of astrophysical results. C.M.F.M. developed and interpreted the relic G.W. results. C.M.F.M., S.A.S., and X.S. developed and interpreted the cosmic strings results. J.A.E., Rv.H., S.T.M., C.M.F.M., A.S., X.S., J.S., S.B.-S., S.R.T., and M.V. wrote the paper, collected the bibliography, prepared figures and tables.

We would like to thank David Merritt for his useful comments on the manuscript. The NANOGrav project receives support from National Science Foundation (NSF) PIRE program award number 0968296 and NSF Physics Frontier Center award number 1430284. NANOGrav research at UBC is supported by an NSERC Discovery Grant and Discovery Accelerator Supplement and by the Canadian Institute for Advanced Research. D.R.M. acknowledges partial support through the New York Space Grant Consortium. M.V. acknowledges support from the JPL RTD program. Portions of this research were carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. C.M.F.M. was supported by a Marie Curie International Outgoing Fellowship within the 7th European Community Framework Programme. S.R.T. was supported by appointment to the NASA Postdoctoral Program at the Jet Propulsion Laboratory, administered by Oak Ridge Associated Universities through a contract with NASA. J.A.E. and Rv.H. acknowledge support by NASA through Einstein Fellowship grants PF4-150120 and PF3-140116, respectively. S.A.S. acknowledges funding from an NWO Vidi fellowship (PI: JTW Hessels). A.S. is supported by a University Research Fellowship of the Royal Society. Y.W. is supported by the National Science Foundation of China (NSFC) under grant NO. 11503007. This work was supported in part by National

Science Foundation Grant No. PHYS-1066293 and by the hospitality of the Aspen Center for Physics. This research was performed in part using the Zwicky computer cluster at Caltech supported by NSF under MRI-R2 award no. PHY-0960291 and by the Sherman Fairchild Foundation. A majority of the computational work was performed on the Nemo cluster at UWM supported by NSF grant No. 0923409. Parts of the analysis in this work were carried out on the Nimrod cluster made available by S.M.R. Data for this project were collected using the facilities of the National Radio Astronomy Observatory and the Arecibo Observatory. The National Radio Astronomy Observatory is a facility of the NSF operated under cooperative agreement by Associated Universities, Inc. The Arecibo Observatory is operated by SRI International under a cooperative agreement with the NSF (AST-1100968), and in alliance with Ana G. Méndez-Universidad Metropolitana and the Universities Space Research Association.

REFERENCES

- Aarseth, S. J. 2003, *Ap&SS*, **285**, 367
- Ade, P. A. R., Aikin, R. W., Barkats, D., et al. 2014, *PhRvL*, **112**, 241101
- Ade, P. A. R., et al. 2015, arXiv:1502.01589
- Allen, B., & Romano, J. D. 1999, *PhRvD*, **59**, 102001
- Amaro-Seoane, P., Miller, M. C., & Freitag, M. 2009, *ApJL*, **692**, L50
- Anholm, M., Ballmer, S., Creighton, J. D. E., Price, L. R., & Siemens, X. 2009, *PhRvD*, **79**, 084030
- Antoniadis, I., & Patil, S. P. 2015, *EPJC*, **75**, 182
- Armitage, P. J., & Natarajan, P. 2002, *ApJL*, **567**, L9
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2014, *ApJ*, **794**, 141
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2015, *ApJ*, **813**, 65
- Babak, S., Petiteau, A., Sesana, A., et al. 2016, *MNRAS*, **455**, 1665
- Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, *Natur*, **287**, 307
- Berentzen, I., Preto, M., Berczik, P., Merritt, D., & Spurzem, R. 2009, *ApJ*, **695**, 455
- BICEP2/Keck, Planck Collaborations, et al. 2015, arXiv:1502.00612
- Blanco-Pillado, J. J., Olum, K. D., & Shlaer, B. 2014, *PhRv*, **D89**, 023512
- Burke-Spolaor, S. 2015, arXiv:1511.07869
- Chamberlin, S. J., Creighton, J. D. E., Siemens, X., et al. 2015, *PhRvD*, **91**, 044048
- Colpi, M. 2014, *SSRv*, **183**, 189
- Copeland, E. J., Myers, R. C., & Polchinski, J. 2004, *JHEP*, **06**, 013
- Cordes, J. M. 2013, *CQGra*, **30**, 224002
- Cordes, J. M., & Shannon, R. M. 2010, arXiv:1010.3785
- Cornish, N. J., & Littenberg, T. B. 2015, *CQGra*, **32**, 135012
- Damour, T., & Vilenkin, A. 2001, *PhRvD*, **64**, 064008
- Damour, T., & Vilenkin, A. 2005, *PhRvD*, **D71**, 063510
- Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2013, *ApJ*, **762**, 94
- Detweiler, S. 1979, *ApJ*, **234**, 1100
- Di Matteo, T., Carilli, C. L., & Fabian, A. C. 2001, *ApJ*, **547**, 731
- Dotti, M., Colpi, M., Haardt, F., & Mayer, L. 2007, *MNRAS*, **379**, 956
- Dotti, M., Merloni, A., & Montuori, C. 2015, *MNRAS*, **448**, 3603
- Emsellem, E. 2013, *MNRAS*, **433**, 1862
- Enoki, M., Enoki, M., Nagashima, M., & Nagashima, M. 2007, *PTHPh*, **117**, 241
- Foster, R. S., & Backer, D. C. 1990, *ApJ*, **361**, 300
- Gair, J., Romano, J. D., Taylor, S., & Mingarelli, C. M. F. 2014, *PhRvD*, **90**, 082001
- Gebhardt, K., Adams, J., Richstone, D., et al. 2011, *ApJ*, **729**, 119
- Goicovic, F. G., Cuadra, J., Sesana, A., et al. 2015, arXiv:1507.05596
- Graham, A. W., & Scott, N. 2013, *ApJ*, **764**, 151
- Grishchuk, L. 2001, in *Gyros, Clocks, Interferometers...: Testing Relativistic Gravity in Space*, Vol. 562, ed. C. Lämmerzahl, C. Everitt, & F. Hehl (Berlin, Heidelberg: Springer)
- Grishchuk, L. P. 1976, *JETPL*, **23**, 293
- Grishchuk, L. P. 1977, in *Eighth Texas Symp. on Relativistic Astrophysics 302*, Annals of the New York Academy of Sciences, ed. M. D. Papagiannis (New York: New York Academy of Sciences), 439
- Grishchuk, L. P. 2005, *PhyU*, **48**, 1235
- Haiman, Z., Kocsis, B., & Menou, K. 2009, *ApJ*, **700**, 1952
- Häring, N., & Rix, H.-W. 2004, *ApJL*, **604**, L89
- Hellings, R. W., & Downs, G. S. 1983, *ApJL*, **265**, L39
- Hemsendorf, M., Sigurdsson, S., & Spurzem, R. 2002, *ApJ*, **581**, 1256
- Hobbs, G. 2013, *CQGra*, **30**, 224007
- Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, *CQGra*, **27**, 084013
- Huerta, E. A., McWilliams, S. T., Gair, J. R., & Taylor, S. R. 2015, arXiv:1504.00928
- Ilbert, O., McCracken, H. J., Le Fèvre, O., et al. 2013, *A&A*, **556**, A55
- Ivanov, P. B., Papaloizou, J. C. B., & Polnarev, A. G. 1999, *MNRAS*, **307**, 79
- Jaffe, A. H., & Backer, D. C. 2003, *ApJ*, **583**, 616
- Jenet, F. A., Hobbs, G. B., Lee, K. J., & Manchester, R. N. 2005, *ApJL*, **625**, L123
- Jones, N. T., Stoica, H., & Tye, S.-H. H. 2003, *PhLB*, **563**, 6
- Keenan, R. C., Foucaud, S., De Propriis, R., et al. 2014, *ApJ*, **795**, 157
- Kibble, T. W. B. 1976, *JPhA*, **9**, 1387
- Kocsis, B., & Sesana, A. 2011, *MNRAS*, **411**, 1467
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, **51**, 511
- Kramer, M., & Champion, D. J. 2013, *CQGra*, **30**, 224009
- Kramer, M., Stairs, I. H., Manchester, R. N., et al. 2006, *Sci*, **314**, 97
- Lasky, P. D., Mingarelli, C. M. F., Smith, T. L., et al. 2015, arXiv:1511.05994
- Lentati, L., Alexander, P., Hobson, M. P., et al. 2013, *PhRvD*, **87**, 104021
- Lentati, L., Taylor, S. R., Mingarelli, C. M. F., et al. 2015, arXiv:1504.03692
- Lin, Y.-T., Ostriker, J. P., & Miller, C. J. 2010, *ApJ*, **715**, 1486
- Linde, A. D. 1982, *PhLB*, **108**, 389
- Lorimer, D., & Kramer, M. 2005, *Handbook of Pulsar Astronomy*, Vol. 4 (1st ed.; Cambridge: Cambridge Univ. Press)
- Maggiore, M. 2000, *PhR*, **331**, 283
- McConnell, N. J., & Ma, C.-P. 2013, *ApJ*, **764**, 184
- McConnell, N. J., Ma, C.-P., Gebhardt, K., et al. 2011, *Natur*, **480**, 215
- McLaughlin, M. A. 2013, *CQGra*, **30**, 224008
- McWilliams, S. T., Ostriker, J. P., & Pretorius, F. 2014, *ApJ*, **789**, 156
- Merritt, D., & Ferrarese, L. 2001, in *ASP Conf. Ser. 249*, The Central Kiloparsec of Starbursts and AGN: The La Palma Connection, ed. J. H. Knapen et al. (San Francisco, CA: ASP), 335
- Middleton, H., Del Pozzo, W., Farr, W. M., Sesana, A., & Vecchio, A. 2015, arXiv:1507.00992
- Milosavljević, M., & Merritt, D. 2003, *ApJ*, **596**, 860
- Mingarelli, C. M. F., & Sidery, T. 2014, *PhRvD*, **90**, 062011
- Mingarelli, C. M. F., Sidery, T., Mandel, I., & Vecchio, A. 2013, *PhRvD*, **88**, 062005
- Moore, C. J., Taylor, S. R., & Gair, J. R. 2015, *CQGra*, **32**, 055004
- Mukhanov, V. F., Feldman, H. A., & Brandenberger, R. H. 1992, *PhR*, **215**, 203
- Ölmez, S., Mandic, V., & Siemens, X. 2010, *PhRvD*, **81**, 104028
- Peters, P. C., & Mathews, J. 1963, *PhRv*, **131**, 435
- Phinney, E. S. 2001, arXiv:astro-ph/0108028
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, *A&A*, **571**, A25
- Quinlan, G. D. 1996, *NewA*, **1**, 35
- Rajagopal, M., & Romani, R. W. 1995, *ApJ*, **446**, 543
- Ravi, V., Wyithe, J. S. B., Shannon, R. M., & Hobbs, G. 2015, *MNRAS*, **447**, 2772
- Ravi, V., Wyithe, J. S. B., Shannon, R. M., Hobbs, G., & Manchester, R. N. 2014, *MNRAS*, **442**, 56
- Robotham, A. S. G., Driver, S. P., Davies, L. J. M., et al. 2014, *MNRAS*, **444**, 3986
- Rosado, P. A., Sesana, A., & Gair, J. 2015, *MNRAS*, **451**, 2417
- Sampson, L., Cornish, N. J., & McWilliams, S. T. 2015, *PhRvD*, **91**, 084055
- Sanidas, S. A., Battye, R. A., & Stappers, B. W. 2012, *PhRvD*, **85**, 122003
- Sanidas, S. A., Battye, R. A., & Stappers, B. W. 2013, *ApJ*, **764**, 108
- Sarangi, S., & Tye, S.-H. H. 2002, *PhLB*, **536**, 185
- Sazhin, M. V. 1978, *SvA*, **22**, 36
- Sesana, A. 2010, *ApJ*, **719**, 851
- Sesana, A. 2013a, *MNRAS*, **433**, L1
- Sesana, A. 2013b, *CQGra*, **30**, 224014
- Sesana, A., Haardt, F., Madau, P., & Volonteri, M. 2004, *ApJ*, **611**, 623
- Sesana, A., & Khan, F. M. 2015, arXiv:1505.02062
- Sesana, A., Vecchio, A., & Colacino, C. N. 2008, *MNRAS*, **390**, 192
- Shannon, R. M., & Cordes, J. M. 2010, *ApJ*, **725**, 1607
- Shannon, R. M., Ravi, V., Coles, W. A., et al. 2013, *Sci*, **342**, 334
- Shannon, R. M., Ravi, V., Lentati, L. T., et al. 2015, arXiv:1509.07320
- Siemens, X., Creighton, J., Maor, I., et al. 2006, *PhRvD*, **D73**, 105001
- Siemens, X., Mandic, V., & Creighton, J. 2007, *PhRvL*, **98**, 111101
- Starobinsky, A. A. 1980, *PhLB*, **91**, 99
- Stinebring, D. 2013, *CQGra*, **30**, 224006
- Taylor, S. R., & Gair, J. R. 2013, *PhRvD*, **88**, 084001

- Thrane, E., & Romano, J. D. 2013, [PhRvD](#), **88**, 124032
- Tomeczak, A. R., Quadri, R. F., Tran, K.-V. H., et al. 2014, [ApJ](#), **783**, 85
- van den Bosch, R. C. E., Gebhardt, K., Gültekin, K., et al. 2012, [Natur](#), **491**, 729
- van Haasteren, R., & Vallisneri, M. 2014, [PhRvD](#), **90**, 104012
- van Haasteren, R., & Vallisneri, M. 2015, [MNRAS](#), **446**, 1170
- Vasiliev, E., Antonini, F., & Merritt, D. 2015, arXiv:1505.05480
- Vilenkin, A. 1981, [PhLB](#), **107**, 47
- Vilenkin, A., & Shellard, E. P. S. 2000, *Cosmic Strings and Other Topological Defects* (Cambridge: Cambridge Univ. Press)
- Walsh, J. L., Barth, A. J., Ho, L. C., & Sarzi, M. 2013, [ApJ](#), **770**, 86
- Walsh, J. L., van den Bosch, R. C. E., Gebhardt, K., et al. 2015, arXiv:1511.04455
- Watanabe, Y., & Komatsu, E. 2006, [PhRvD](#), **73**, 123515
- Weisberg, J. M., Nice, D. J., & Taylor, J. H. 2010, [ApJ](#), **722**, 1030
- Woodbury, M. A. 1950, Statistical Research Group, Memorandum Rep. 42 (Princeton, NJ: Princeton Univ. Press), 106
- Wyithe, J. S. B., & Loeb, A. 2003, [ApJ](#), **590**, 691
- Yardley, D. R. B., Hobbs, G. B., Jenet, F. A., et al. 2010, [MNRAS](#), **407**, 669
- Zhao, W. 2011, [PhRvD](#), **83**, 104021
- Zhao, W., Zhang, Y., You, X.-P., & Zhu, Z.-H. 2013, [PhRvD](#), **87**, 124012