



The NANOGrav 15 yr Data Set: Piecewise Power-law Reconstruction of the Gravitational-wave Background

Gabriella Agazie¹ , Akash Anumalapudi² , Anne M. Archibald³ , Zaven Arzumianian⁴ , Jeremy G. Baier⁵ , Paul T. Baker⁶ , Bence Bécsey⁷ , Amit Bhoonah⁸ , Laura Blecha⁹ , Adam Brazier^{10,11} , Paul R. Brook⁷ , Sarah Burke-Spolaor^{12,13,67} , Rand Burnette⁵ , Robin Case⁵ , J. Andrew Casey-Clyde¹⁴ , Maria Charisi^{15,16} , Shami Chatterjee¹⁰ , Tyler Cohen¹⁷ , James M. Cordes¹⁰ , Neil J. Cornish¹⁸ , Fronefield Crawford¹⁹ , H. Thankful Cromartie^{68,72} , Kathryn Crowter²⁰ , Megan E. DeCesar^{21,68} , Paul B. Demorest²² , Heling Deng⁵ , Lankeswar Dey^{12,13} , Timothy Dolch^{23,24} , Elizabeth C. Ferrara^{25,26,27} , William Fiore²⁰ , Emmanuel Fonseca^{12,13} , Gabriel E. Freedman¹ , Emiko C. Gardiner²⁸ , Nate Garver-Daniels^{12,13} , Peter A. Gentile^{12,13} , Kyle A. Gersbach²⁹ , Joseph Glaser^{12,13} , Brenda D. Gómez-Cortés⁸ , Deborah C. Good³⁰ , Kayhan Gültekin³¹ , C. J. Harris³¹ , Jeffrey S. Hazboun⁵ , Ross J. Jennings^{12,13,69} , Aaron D. Johnson^{1,32} , Megan L. Jones¹ , David L. Kaplan¹ , Luke Zoltan Kelley²⁸ , Matthew Kerr³³ , Joey S. Key³⁴ , Nima Laal²⁹ , Michael T. Lam^{35,36,37} , William G. Lamb²⁹ , Bjorn Larsen³⁸ , T. Joseph W. Lazio³⁹ , Natalia Lewandowska⁴⁰ , Monica Leys⁸ , Tingting Liu⁴¹ , Duncan R. Lorimer^{12,13} , Jing Luo^{42,70} , Ryan S. Lynch⁴³ , Chung-Pei Ma^{28,44} , Dustin R. Madison⁴⁵ , Cayenne Matt³¹ , Alexander McEwen¹ , James W. McKee⁴⁶ , Maura A. McLaughlin⁴⁷ , Natasha McMann²⁹ , Bradley W. Meyers^{47,48} , Patrick M. Meyers³² , Chiara M. F. Mingarelli³⁸ , Andrea Mitridate⁴⁹ , Cherry Ng⁵⁰ , David J. Nice⁵¹ , Stella Koch Ocker^{32,52} , Ken D. Olum⁵³ , Timothy T. Pennucci⁵⁴ , Benetge B. P. Perera⁵⁵ , Polina Petrov²⁹ , Nihan S. Pol⁵⁶ , Henri A. Radovan⁵⁷ , Scott M. Ransom⁵⁸ , Paul S. Ray³³ , Joseph D. Romano⁵⁶ , Jessie C. Runnøe²⁹ , Alexander Saffer^{58,69} , Shashwat C. Sardesai¹ , Ann Schmiedekamp⁵⁹ , Carl Schmiedekamp⁵⁹ , Kai Schmitz^{60,61} , Brent J. Shapiro-Albert^{12,13,62} , Xavier Siemens^{1,5} , Joseph Simon^{63,71} , Sophia V. Sosa Fiscella^{36,37} , Ingrid H. Stairs²⁰ , Daniel R. Stinebring⁶⁴ , Kevin Stovall²² , Abhimanyu Susobhanan⁶⁵ , Joseph K. Swiggum^{32,39} , Jacob Taylor⁵ , Stephen R. Taylor²⁹ , Mercedes S. Thompson²⁰ , Jacob E. Turner⁴³ , Michele Vallisneri^{32,39} , Rutger van Haasteren⁶⁵ , Sarah J. Vigeland¹ , Haley M. Wahl^{12,13} , Si Wang⁸ , Kevin P. Wilson^{12,13} , Caitlin A. Witt⁶⁶ , David Wright⁵ , and Olivia Young^{36,37}

¹ Center for Gravitation, Cosmology and Astrophysics, Department of Physics and Astronomy, University of Wisconsin-Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA

² Department of Physics and Astronomy, University of North Carolina, Chapel Hill, NC 27599, USA

³ Newcastle University, NE1 7RU, Newcastle upon Tyne, UK

⁴ X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA

⁵ Department of Physics, Oregon State University, Corvallis, OR 97331, USA

⁶ Department of Physics and Astronomy, Widener University, One University Place, Chester, PA 19013, USA

⁷ Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

⁸ PITT PACC, Department of Physics and Astronomy, University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, PA 15260, USA; amit.bhoonah@nanograv.org

⁹ Physics Department, University of Florida, Gainesville, FL 32611, USA

¹⁰ Cornell Center for Astrophysics and Planetary Science and Department of Astronomy, Cornell University, Ithaca, NY 14853, USA

¹¹ Cornell Center for Advanced Computing, Cornell University, Ithaca, NY 14853, USA

¹² Department of Physics and Astronomy, West Virginia University, P.O. Box 6315, Morgantown, WV 26506, USA

¹³ Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

¹⁴ Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269-3046, USA

¹⁵ Department of Physics and Astronomy, Washington State University, Pullman, WA 99163, USA

¹⁶ Institute of Astrophysics, FORTH, GR-71110, Heraklion, Greece

¹⁷ Department of Physics, New Mexico Institute of Mining and Technology, 801 Leroy Place, Socorro, NM 87801, USA

¹⁸ Department of Physics, Montana State University, Bozeman, MT 59717, USA

¹⁹ Department of Physics and Astronomy, Franklin & Marshall College, P.O. Box 3003, Lancaster, PA 17604, USA

²⁰ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada

²¹ Department of Physics and Astronomy, George Mason University, Fairfax, VA 22030, USA

²² National Radio Astronomy Observatory, 1003 Lopezville Road, Socorro, NM 87801, USA

²³ Department of Physics, Hillsdale College, 33 E. College Street, Hillsdale, MI 49242, USA

²⁴ Eureka Scientific, 2452 Delmer Street, Suite 100, Oakland, CA 94602-3017, USA

²⁵ Department of Astronomy, University of Maryland, College Park, MD 20742, USA

²⁶ Center for Research and Exploration in Space Science and Technology, NASA/GSFC, Greenbelt, MD 20771, USA

²⁷ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

²⁸ Department of Astronomy, University of California, Berkeley, 501 Campbell Hall #3411, Berkeley, CA 94720, USA

²⁹ Department of Physics and Astronomy, Vanderbilt University, 2301 Vanderbilt Place, Nashville, TN 37235, USA

³⁰ Department of Physics and Astronomy, University of Montana, 32 Campus Drive, Missoula, MT 59812, USA

³¹ Department of Astronomy and Astrophysics, University of Michigan, Ann Arbor, MI 48109, USA

³² Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

³³ Space Science Division, U.S. Naval Research Laboratory, Washington, DC 20375-5352, USA

³⁴ University of Washington Bothell, 18115, Campus Way NE, Bothell, WA 98011, USA

³⁵ SETI Institute, 339 N Bernardo Avenue, Suite 200, Mountain View, CA 94043, USA

³⁶ School of Physics and Astronomy, Rochester Institute of Technology, Rochester, NY 14623, USA

³⁷ Laboratory for Multiwavelength Astrophysics, Rochester Institute of Technology, Rochester, NY 14623, USA

³⁸ Department of Physics, Yale University, New Haven, CT 06520, USA

³⁹ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

- ⁴⁰ Department of Physics and Astronomy, State University of New York at Oswego, Oswego, NY 13126, USA
⁴¹ Department of Physics and Astronomy, Georgia State University, 25 Park Place, Suite 605, Atlanta, GA 30303, USA
⁴² Department of Astronomy & Astrophysics, University of Toronto, 50 Saint George Street, Toronto, ON M5S 3H4, Canada
⁴³ Green Bank Observatory, P.O. Box 2, Green Bank, WV 24944, USA
⁴⁴ Department of Physics, University of California, Berkeley, CA 94720, USA
⁴⁵ Department of Physics, Occidental College, 1600 Campus Road, Los Angeles, CA 90041, USA
⁴⁶ Department of Physics and Astronomy, Union College, Schenectady, NY 12308, USA
⁴⁷ Australian SKA Regional Centre (AusSRC), Curtin University, Bentley, WA 6102, Australia
⁴⁸ International Centre for Radio Astronomy Research (ICRAR), Curtin University, Bentley, WA 6102, Australia
⁴⁹ Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany
⁵⁰ Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St George Street, Toronto, ON M5S 3H4, Canada
⁵¹ Department of Physics, Lafayette College, Easton, PA 18042, USA
⁵² The Observatories of the Carnegie Institution for Science, Pasadena, CA 91101, USA
⁵³ Institute of Cosmology, Department of Physics and Astronomy, Tufts University, Medford, MA 02155, USA
⁵⁴ Institute of Physics and Astronomy, Eötvös Loránd University, Pázmány P. s. 1/A, 1117 Budapest, Hungary
⁵⁵ Arecibo Observatory, HC3 Box 53995, Arecibo, PR 00612, USA
⁵⁶ Department of Physics, Texas Tech University, Box 41051, Lubbock, TX 79409, USA
⁵⁷ Department of Physics, University of Puerto Rico, Mayagüez, PR 00681, USA
⁵⁸ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA
⁵⁹ Department of Physics, Penn State Abington, Abington, PA 19001, USA
⁶⁰ Institute for Theoretical Physics, University of Münster, 48149 Münster, Germany; kai.schmitz@nanograv.org
⁶¹ Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan
⁶² Giant Army, 915A 17th Avenue, Seattle, WA 98122, USA
⁶³ Department of Astrophysical and Planetary Sciences, University of Colorado, Boulder, CO 80309, USA
⁶⁴ Department of Physics and Astronomy, Oberlin College, Oberlin, OH 44074, USA
⁶⁵ Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Callinstraße 38, D-30167 Hannover, Germany
⁶⁶ Department of Physics, Wake Forest University, 1834 Wake Forest Road, Winston-Salem, NC 27109, USA

Received 2026 January 27; revised 2026 May 7; accepted 2026 May 19; published 2026 June 16

Abstract

The NANOGrav 15 yr (NG15) data set provides evidence for a gravitational-wave background (GWB) signal at nHz frequencies, which is expected to originate either from a cosmic population of inspiraling supermassive black hole binaries or new particle physics in the early Universe. A firm identification of the source of the NG15 signal requires an accurate reconstruction of its frequency spectrum. In this Letter, we provide such a spectral characterization of the NG15 signal based on a piecewise power-law (PPL) ansatz that strikes a balance between existing alternatives in the literature. Our PPL reconstruction is more flexible than the standard constant power-law model, which describes the GWB spectrum in terms of only two parameters: an amplitude A and a spectral index γ . Concurrently, it better approximates physically realistic GWB spectra—especially those of cosmological origin—than the free spectral model, since the latter allows for arbitrary variations in the GWB amplitude from one frequency bin to the next. Our PPL reconstruction of the NG15 signal relies on individual PPL models with a fixed number of internal nodes (i.e., constant power law, broken power law, doubly broken power law, etc.), which are ultimately combined in a Bayesian model average. The data products resulting from our analysis provide the basis for fast refits of spectral GWB models.

Unified Astronomy Thesaurus concepts: [Gravitational waves \(678\)](#)

1. Introduction

Pulsar timing arrays (PTAs) search for a stochastic gravitational-wave background (GWB) at nHz frequencies by monitoring the times of arrival (TOAs) of pulses from galactic millisecond pulsars over time spans of years and decades (S. R. Taylor 2021). The hallmark signature of a GWB signal in PTA data consists of a characteristic cross correlation between the timing residuals of pairs of pulsars, where the

timing residuals of each individual pulsar are obtained by comparing its observed TOAs with the predictions of a sophisticated, pulsar-specific timing model. The expected cross correlation is described by the Hellings–Downs (HD) curve (R. W. Hellings & G. S. Downs 1983), an overlap reduction function that only depends on the angular separation of pulsars in the sky and which is dominated by its quadrupole contribution when decomposed into Legendre polynomials. The latest generation of PTA datasets—specifically, the latest datasets from CPTA (H. Xu et al. 2023), EPTA and InPTA (J. Antoniadis et al. 2023), MPTA (M. T. Miles et al. 2024), NANOGrav (G. Agazie et al. 2023a), and PPTA (D. J. Reardon et al. 2023)—contain an excess timing delay signal that features, at varying levels of statistical significance, the expected HD correlation. It is thus reasonable to assume that PTAs are currently on the brink of a 5σ detection of a GWB.

In this Letter, we focus on the NANOGrav 15 yr (NG15) dataset (G. Agazie et al. 2023b), which contains evidence for an HD-correlated signal with a statistical significance of $3\sigma \cdots 4\sigma$, and assume that the signal seen in this dataset does

⁶⁷ Sloan Fellow.

⁶⁸ Resident at the U.S. Naval Research Laboratory.

⁶⁹ NANOGrav Physics Frontiers Center Postdoctoral Fellow.

⁷⁰ Deceased.

⁷¹ NSF Astronomy and Astrophysics Postdoctoral Fellow.

⁷² National Research Council Research Associate, National Academy of Sciences.



indeed correspond to a genuine GWB. A crucial question then is that of the origin of the signal. The most plausible source of the signal would arguably be a cosmic population of supermassive black hole binaries (SMBHBs) at the centers of galaxies (G. Agazie et al. 2023c; D. J. D’Orazio & M. Charisi 2023; Y. Chen et al. 2026; C. Matt et al. 2026). While no individual SMBHB system close to merger has been detected thus far, it is known that most galaxies host a supermassive black hole at their core, and galaxies merge during structure formation. Galaxy mergers across cosmic history should thus give rise to an abundance of SMBHBs and hence a viable source of a GWB signal in the PTA frequency band. Alternatively, the signal may originate from exotic sources in the early Universe, necessarily involving new physics beyond the Standard Model of particle physics, e.g., cosmic inflation, scalar-induced gravitational waves (SIGWs), cosmic defects, or a cosmological first-order phase transition (C. Caprini & D. G. Figueroa 2018; A. Afzal et al. 2023; J. Antoniadis et al. 2024; J. Ellis et al. 2024). This cosmological interpretation is more speculative than the astrophysical interpretation in terms of SMBHBs. Meanwhile, the detection of a gravitational-wave (GW) echo of the Big Bang would be particularly exciting, as it would open an observational window into the history of our Universe at very early times and hence particle physics at very high energies.

Making progress toward the firm identification of the source of the NG15 signal requires a multipronged approach. A PTA observable of particular interest in this context is the GWB angular power spectrum (G. Agazie et al. 2023d; T. Konstandin et al. 2025): the GWB signal from SMBHBs is expected to feature much more pronounced anisotropies across the sky than a smooth GWB signal from the early Universe (E. C. Gardiner et al. 2024). Similarly, searches for continuous wave (CW) signals from individual SMBHB systems could, in case of a positive signal, confirm that SMBHBs exist and succeed in acting as sources of nHz GWs (G. Agazie et al. 2023e; N. Agarwal et al. 2026). However, neither GWB anisotropies nor CW signals have been detected thus far, which has two important implications. First, the cosmological interpretation of the NG15 signal remains a viable and relevant option for now. Second, current attempts at model discrimination must focus on yet another PTA observable: the GWB frequency spectrum.

Spectral GWB models are often formulated in terms of unique model parameters that do not exist in other models, e.g., the cosmic string tension $G\mu$ in models of GWs from a cosmic string network, the phase transition temperature T_* in models of GWs from a cosmological phase transition, etc. To unify the discussion of these different GWB models and facilitate a global model comparison, it is therefore helpful to develop a more general language based on physics-agnostic reference models. The most prominent reference model in this regard is the constant power-law (CPL) model, which describes the GW energy density power spectrum $h^2\Omega_{\text{GW}}$ by a simple power law parameterized in terms of two parameters, an amplitude A and a spectral index γ (see below for more details). Furthermore, the CPL model can be easily generalized to a running power-law (RPL) model by allowing for a logarithmic running of the spectral index, $\gamma \rightarrow \gamma_{\text{run}}(f)$ (G. Agazie et al. 2025a). As recently discussed

in D. Esmyol et al. (2025), it is straightforward to map beyond the Standard Model predictions for the GWB spectrum onto the CPL and RPL models and then use these maps for refitting analyses.

A second prominent physics-agnostic reference model is the free spectral model. While the CPL model can be regarded as a minimal model that fixes the GWB amplitude in each frequency bin in terms of just two parameters, the free spectral model is maximally flexible. Indeed, by definition, the free spectral model treats the GWB amplitude in each frequency bin as a free parameter that can be chosen independently of all other GWB amplitudes. The posterior distributions for these amplitude parameters then yield the well-known NG15 violins that are often used as a graphical illustration of the NG15 signal (G. Agazie et al. 2023a).

Many physically realistic GWB models are expected to be quite different from these simple reference models. The GWB spectrum from SMBHBs, e.g., can be modeled as a CPL only to first approximation. More realistic SMBHB models predict a spectral turnover at low frequencies due to environmental interactions (B. Kocsis & A. Sesana 2011) and a spectral break at high frequencies caused by the discrete nature of the SMBHB population (A. Sesana et al. 2008; B. Bécsy et al. 2022; G. Agazie et al. 2025b; W. G. Lamb et al. 2025). Another example is GWs from cosmic inflation (M. C. Guzzetti et al. 2016), which do not result in an exact CPL spectrum, even if the primordial tensor power spectrum produced during inflation should have a perfect CPL shape (which by itself is a strong assumption). In this case, the primordial tensor power spectrum will be modulated by the transfer function that relates the primordial spectrum to the present-day GWB spectrum and which is sensitive to changes in the effective number of relativistic degrees of freedom across cosmic history. Together, the primordial tensor spectrum and the transfer function then result in a GWB spectrum that deviates from a CPL (Y. Watanabe & E. Komatsu 2006; K. Saikawa & S. Shirai 2018).

Meanwhile, the free spectral model represents a maximally conservative reconstruction of the GWB signal. Irrespective of physical expectations, it accounts for any conceivable shape of the GWB spectrum, no matter how strongly the GWB amplitude may vary up and down from one frequency bin to the next. It is therefore legitimate to use the free spectral model as a basis for conservative refitting analyses, which is the underlying idea of the `ceffyl` refitting framework (W. G. Lamb et al. 2023). At the same time, it is clear that the free spectral model encapsulates a multitude of spectral shapes that are unlikely to be realized in any physically realistic model.

This observation motivates us to seek alternative reference models. We are interested in models that (i) better reflect our physical expectation that the GWB spectrum should be described by a smooth and well-behaved function of frequency, in particular if it originates from a cosmological source, and that (ii) are more flexible than a rigid CPL. The RPL model discussed in G. Agazie et al. (2025a) can be regarded as a first step in this direction. In the present Letter, we continue our construction of physics-agnostic reference models and investigate yet another promising approach: the reconstruction of the GWB signal in terms of a piecewise power law (PPL), i.e., a set of linear spline interpolations between movable nodes in the $\log_{10}f$ – $\log_{10}h^2\Omega_{\text{GW}}$ plane.

Our analysis is primarily inspired by literature on the cosmic microwave background (CMB) that is concerned with the reconstruction of the primordial scalar power spectrum (J. A. Vázquez et al. 2012a; G. Aslanyan et al. 2014; F. Finelli et al. 2018; W. J. Handley et al. 2019). Our reconstruction of the GWB spectrum based on the NG15 data can notably be regarded as the equivalent of the Bayesian reconstruction of the primordial scalar power spectrum by the Planck Collaboration based on the Planck 2015 (Planck Collaboration et al. 2016) and Planck 2018 data (Planck Collaboration et al. 2020). Apart from this, the PPL ansatz (sometimes also referred to as the “nodal” or “flexknot” reconstruction) is routinely used in a variety of analyses in cosmology, ranging from reconstructions of the equation of state of dark energy (J. A. Vázquez et al. 2012b; S. Hee et al. 2016; A. N. Ormondroyd et al. 2025) over reconstructions of the cosmic reionization history (M. Millea & F. Bouchet 2018; S. Heimersheim et al. 2022) to reconstructions of galaxy cluster profiles (M. Olamaie et al. 2018) and the 21 cm signal (S. Heimersheim et al. 2023; E. Shen et al. 2024). In the context of the SIGW interpretation of the NG15 signal, the PPL ansatz was recently also used by A. Ghaleb et al. (2025), in order to constraint the primordial scalar power spectrum in the NANOGrav frequency band.

A decisive advantage of our PPL ansatz is that it properly accounts for the level of complexity supported by the NG15 data. As explained in more detail below, our method relies on a Bayesian model average (BMA; R. van Haasteren 2024) that weights PPL models with different numbers of frequency nodes according to their relative Bayesian evidence. This approach guarantees that our reconstruction of the GWB signal ends up being sufficiently complex, but not more complex than necessary. In addition, we will eventually allow all internal frequency nodes to freely move, such that our final PPL reconstruction of the GWB signal is sensitive to localized features in the GWB spectrum. Conversely, in the absence of such features, our ansatz allows us to constrain local deviations from a smooth CPL-like spectrum; in practice, this is what we will find when applying our methodology to the NG15 data. In summary, the two characteristic properties of our approach—a data-driven level of complexity and sensitivity to local features (i.e., local constraining power)—distinguish our ansatz in this work from other physics-agnostic reference models in the literature, such as the CPL or RPL model and the free spectral model.

The rest of the Letter is organized as follows. In Section 2, we will introduce our framework for the reconstruction of the GWB, specifically, our set of PPL models. In particular, we will distinguish between BINNED and FREE versions of our models, in which the locations of the internal node frequencies are either kept fixed or allowed to vary. In Section 3, we subsequently describe our Bayesian fit to the NG15 data and discuss in turn our main results: first, our results for the individual PPL models and then our results for the BMA over these models. Finally, in Section 4, we give a brief outlook on how the output of our analysis could be used for refitting analyses in future work, before we conclude in Section 5. The busy reader may directly skip to Figure 1, which shows the main result of our analysis: our final PPL reconstruction of the GWB signal in the NG15 data, alongside the standard NG15 violins.

2. Framework and Models

2.1. GW Power Spectrum and Spectral Index

A GWB induces extra timing delays in the TOAs of the observed pulsars whose covariance can be expressed in terms of the timing-residual cross-power spectrum. For a stochastic, Gaussian, stationary, isotropic, and unpolarized signal in standard Einstein gravity, one finds

$$S_{ab}^{\text{GWB}}(f) = \Gamma_{ab} \frac{S_h(f)}{6\pi^2 f^2}. \quad (1)$$

Here, Γ_{ab} is the cross-correlation coefficient for the pair of pulsars a and b , which follows from evaluating the HD overlap reduction function at the angle ξ_{ab} , i.e., the angular separation of a and b in the sky,

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

where x_{ab} is shorthand for $x_{ab} = \frac{1}{2}(1 - \cos \xi_{ab})$. The spectral composition of the signal is accounted for by the GW strain power spectrum S_h ,⁷³ which can be related to the characteristic GW strain amplitude

$$h_c(f) = \sqrt{2f S_h(f)}, \quad (3)$$

or, alternatively, to the GW energy density power spectrum $h^2\Omega_{\text{GW}}$, which measures the fractional GW energy density per logarithmic frequency interval in units of the critical energy density of the present Universe, ρ_{crit} ,

$$h^2\Omega_{\text{GW}}(f) = \frac{1}{\rho_{\text{crit}}} \frac{d\rho_{\text{GW}}(f)}{d \ln f} = \frac{4\pi^2 f^3 S_h(f)}{3 (H_0/h)^2}. \quad (4)$$

Here, H_0 is the Hubble constant and h its dimensionless version, $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$. Consequently, H_0/h is just a constant factor that has the units of a frequency, and $h^2\Omega_{\text{GW}}$ is independent of the observed value of H_0 .

Our goal in this Letter is to work out a spectral characterization of the GWB signal in the NG15 data, which can be done equally well in terms of S_h , h_c , and $h^2\Omega_{\text{GW}}$. For definiteness, we will present all of our results in the following, in terms of $h^2\Omega_{\text{GW}}$, which represents the standard choice in large parts of the cosmology and early Universe literature. In addition, we will be interested in a reconstruction of the frequency-dependent index of the GW energy density power spectrum $h^2\Omega_{\text{GW}}$,

$$n(f) = \frac{d \ln h^2\Omega_{\text{GW}}(f)}{d \ln f}, \quad (5)$$

which is closely related to the spectral index of the timing-residual cross-power spectrum in Equation (1),

$$\gamma(f) = -\frac{d \ln S_{ab}^{\text{GWB}}(f)}{d \ln f} = 5 - n(f), \quad (6)$$

where the negative sign in the first identity is a standard convention. For definiteness and in order to facilitate the comparison with other parts of the PTA literature, we will

⁷³ In our convention, S_h is normalized such that the variance of the tensor metric perturbations h_{ij} , summed over the two GW polarization states, can be written as $\langle h_{ij} h^{ij} \rangle = 4 \int_0^\infty df S_h(f)$.




Figure 1. Left panel: PPL reconstruction of the energy density power spectrum $h^2 \Omega_{\text{GW}}$ for the GWB signal in the NG15 data, based on the BMA of six PPL models with a fixed number of freely moving internal nodes (PPL1-FREE,... PPL6-FREE). Right panel: same, but for the frequency-dependent spectral index of the energy density power spectrum, $\gamma = 5 - d \ln h^2 \Omega_{\text{GW}} / d \ln f$. Both reconstructions cover the frequency range from $F_1 = \frac{1}{T_{\text{obs}}}$ to $F_{14} = \frac{14}{T_{\text{obs}}}$, where $T_{\text{obs}} = 16.03$ yr (see the text for details). In the left panel, we also show the NG15 violins in orange, which represent the amplitude posterior densities for the free spectral model. The lower end of the violins at higher frequencies follows from the prior density chosen in G. Agazie et al. (2023a).

report our results for the spectral index in terms of γ rather than n in the remainder of this Letter.

2.2. PPL Models

Let us now define the PPL models that our spectral reconstruction of the GWB is going to be based on. The simplest PPL model is equivalent to the CPL model,

$$h^2 \Omega_{\text{GW}}(f) \stackrel{\text{CPL}}{=} h^2 \Omega_{\text{GW}}(f_0) \left(\frac{f}{f_0} \right)^{5-\gamma_1}, \quad (7)$$

for some reference frequency f_0 . This model has two free parameters: the amplitude of $h^2 \Omega_{\text{GW}}$ at the reference frequency, $h^2 \Omega_{\text{GW}}(f_0)$, and the spectral index γ_1 , which takes the same constant value at all frequencies. In our Bayesian fit analyses below, we will account for the presence of a GWB signal in the frequency range from $f_{\min} = \frac{1}{T_{\text{obs}}}$ to $f_{\max} = \frac{14}{T_{\text{obs}}}$, where T_{obs} denotes the total observing time span, $T_{\text{obs}} = 16.03$ yr. This frequency range corresponds to the first 14 frequency bins in the discrete Fourier series that we shall employ to model the NG15 data. In the simplest PPL model introduced in Equation (7), but also in all other PPL models, we will fix the reference frequency f_0 that determines the overall amplitude factor $h^2 \Omega_{\text{GW}}(f_0)$ at $f_0 = f_{\min}$.

The next simplest PPL model describes the GW power spectrum in terms of a broken power law (BPL),

$$h^2 \Omega_{\text{GW}}(f) \stackrel{\text{BPL}}{=} h^2 \Omega_{\text{GW}}(f_0) \times \begin{cases} \left(\frac{f}{f_0} \right)^{5-\gamma_1}, & f_{\min} \leq f \leq f_1 \\ \left(\frac{f_1}{f_0} \right)^{5-\gamma_1} \left(\frac{f}{f_1} \right)^{5-\gamma_2}, & f_1 \leq f \leq f_{\max} \end{cases}, \quad (8)$$

which now features one amplitude, $h^2 \Omega_{\text{GW}}(f_0)$, and two spectral indices, γ_1 and γ_2 , which describe the power-law

behavior of the GW spectrum at $f < f_1$ and $f > f_1$, respectively. The frequency f_1 marks the position of the spectral break in the BPL model, i.e., the internal node in this PPL model.

For all PPL models with at least two spectral indices, we shall distinguish between two model variants in the following: a BINNED and a FREE version. In the BINNED models, we do not treat the location of the node frequencies as free parameters. Instead, we keep them fixed at equidistant separations on a logarithmic frequency axis in order to set up equally sized frequency bins spanning the range from $f_{\min}$ to $f_{\max}$. In the case of the BPL model introduced above, this results in what we shall refer to as the PPL2-BINNED model, in which $f_0 = f_{\min} = \frac{1}{T_{\text{obs}}}$, $f_1 = \sqrt{f_0 f_2} = \frac{\sqrt{14}}{T_{\text{obs}}}$, and $f_2 = f_{\max} = \frac{14}{T_{\text{obs}}}$. In this model, we thus work with frequency bins for the spectral index: the frequency-dependent index γ is set to γ_1 in the bin from f_0 to f_1 and to γ_2 in the bin from f_1 to f_2 . By contrast, in the FREE versions of our PPL models, we allow the position of the internal node frequencies to vary freely. For the BPL model introduced above, this results in what we shall refer to as the PPL2-FREE model, in which only f_0 and f_2 are kept fixed, but f_1 is allowed to vary between f_0 and f_2 . In summary, this means that the PPL2-BINNED model features three model parameters, $\{h^2 \Omega_{\text{GW}}(f_0), \gamma_1, \gamma_2\}$, while the PPL2-FREE model features four model parameters, $\{h^2 \Omega_{\text{GW}}(f_0), \gamma_1, \gamma_2, f_1\}$. Note that a similar distinction is not possible for the CPL model, i.e., the PPL1 model, which does not feature any internal node. Therefore, when we use model names such as PPL1-FREE (see the caption of Figure 1) and PPL1-BINNED (see Section 3.4), it should be kept in mind that all three names (PPL1-FREE, PPL1-BINNED, PPL1) denote the same model.

Below, our analysis of the BINNED PPL models will serve as a warm-up for our analysis of the FREE PPL models. The BINNED PPL models are slightly easier to handle because of the lower dimensionality of their parameter space. In particular, this will allow us to avoid parameter degeneracies, i.e., flat directions in two-dimensional (2D) marginalized posterior densities. In addition, the BINNED models are inspired by related approaches in the literature; see, e.g.,

C. Caprini et al. (2019), who discuss a method for the spectral reconstruction of the GWB signal that will be measured by the Laser Interferometer Space Antenna. The method described by C. Caprini et al. (2019) also relies on a bin-by-bin power-law reconstruction of the GWB signal (i.e., a binned PPL reconstruction) and provides the basis for the dedicated software package `SGWBinner`. At the same time, it is clear that the BINNED PPL models are more restricted than the FREE PPL models. Our most general and flexible reconstruction of the GWB signal will therefore be based on the FREE PPL models, ensuring that the data itself decides where breaks or features need to be placed on the frequency axis.

Before we move on to the next PPL model, we would like to comment on an alternative possible parameterization of the PPL2-BINNED and PPL2-FREE models. In our analysis, we deliberately construct our parameterization around the spectral indices γ_1 and γ_2 , as this will allow us to directly reconstruct the GW power spectrum and its frequency-dependent spectral index simultaneously. Alternatively, both models can also be parameterized in terms of the amplitudes at the boundaries of the frequency band, $h^2\Omega_{\text{GW}}(f_0)$ and $h^2\Omega_{\text{GW}}(f_2)$, and the location of the internal node, $\{f_1, h^2\Omega_{\text{GW}}(f_1)\}$, where f_1 is either treated as a free parameter (PPL2-FREE) or not (PPL2-BINNED). The relation between these two possible parameterizations is straightforward and directly follows from Equation (8),

$$h^2\Omega_{\text{GW}}(f_1) = h^2\Omega_{\text{GW}}(f_0) \left(\frac{f_1}{f_0} \right)^{5-\gamma_1}, \quad (9)$$

$$h^2\Omega_{\text{GW}}(f_2) = h^2\Omega_{\text{GW}}(f_1) \left(\frac{f_2}{f_1} \right)^{5-\gamma_2}. \quad (10)$$

In their reconstruction of the primordial scalar power spectrum (Planck Collaboration et al. 2016; Planck Collaboration et al. 2020), the Planck Collaboration works with a parameterization of the second type, i.e., in terms of amplitudes and frequencies rather than spectral indices and frequencies. From the perspective of Bayesian inference, both approaches are equivalent as long as the prior densities of all relevant parameters are chosen accordingly. In our analysis, we will work, e.g., with uniform priors on all spectral indices and log-uniform priors on all amplitudes and frequencies. For appropriately chosen prior ranges, our analysis can thus be mapped onto an analysis that parameterizes the GW power spectrum exclusively in terms of amplitudes and frequencies, and that uses log-uniform priors for all these parameters.

The next PPL model describes the GW power spectrum in terms of a doubly broken power law (DBPL),

$$h^2\Omega_{\text{GW}}(f) \stackrel{\text{DBPL}}{=} h^2\Omega_{\text{GW}}(f_0) \times \begin{cases} \left(\frac{f}{f_0} \right)^{5-\gamma_1}, & f_{\min} \leq f \leq f_1 \\ \left(\frac{f_1}{f_0} \right)^{5-\gamma_1} \left(\frac{f}{f_1} \right)^{5-\gamma_2}, & f_1 \leq f \leq f_2 \\ \left(\frac{f_1}{f_0} \right)^{5-\gamma_1} \left(\frac{f_2}{f_1} \right)^{5-\gamma_2} \left(\frac{f}{f_2} \right)^{5-\gamma_3}, & f_2 \leq f \leq f_{\max} \end{cases}. \quad (11)$$

Again, we distinguish between two versions of the model: PPL3-BINNED and PPL3-FREE, where the former features four parameters, $\{h^2\Omega_{\text{GW}}(f_0), \gamma_1, \gamma_2, \gamma_3\}$, and the latter six parameters, $\{h^2\Omega_{\text{GW}}(f_0), \gamma_1, \gamma_2, \gamma_3, f_1, f_2\}$. In the PPL3-BINNED model, the two internal node frequencies are fixed at $f_1 = \frac{14^{1/3}}{T_{\text{obs}}}$

and $f_2 = \frac{14^{2/3}}{T_{\text{obs}}}$. Similarly, as in the case of a BPL, we could switch again to an alternative parameterization purely in terms of node amplitudes $h^2\Omega_{\text{GW}}(f_i)$ and node frequencies f_i .

At this point, it is now obvious how to generalize the PPL models introduced thus far to PPL models with an even larger number of internal nodes, i.e., to a triply BPL, a quadruply BPL, etc. For a general PPL n model, featuring in total n different spectral indices γ_i with $i = 1, \dots, n$, we can write

$$h^2\Omega_{\text{GW}}(f) \stackrel{\text{PPL}n}{=} h^2\Omega_{\text{GW}}(f_0) \quad (12)$$

$$\times \left[\prod_{i=1}^{m-1} \left(\frac{f_i}{f_{i-1}} \right)^{5-\gamma_i} \right] \left(\frac{f}{f_{m-1}} \right)^{5-\gamma_m}, \quad (13)$$

when f falls into the interval $[f_{m-1}, f_m]$. The PPL n model features n such intervals, $m = 1, \dots, n$. In the FREE version of the model, the internal node frequencies f_i with $i = 1, \dots, n-1$ are treated as free parameters that are only subject to the requirement that they be correctly ordered, $f_{\min}=f_0 \leq f_1 \leq f_2 \leq \dots \leq f_n = f_{\max}$. In the BINNED version of the model, on the other hand, all internal frequencies are fixed at $f_i = \frac{14^{i/n}}{T_{\text{obs}}}$. The two nodes at the boundary of the frequency range are always kept constant in both models, $f_0 = \frac{1}{T_{\text{obs}}}$ and $f_n = \frac{14}{T_{\text{obs}}}$. The PPL n -BINNED model thus features $n+1$ parameters, $\{h^2\Omega_{\text{GW}}(f_0), \gamma_1, \dots, \gamma_n\}$, while the PPL n -FREE model features $(n+1) + (n-1) = 2n$ parameters, $\{h^2\Omega_{\text{GW}}(f_0), \gamma_1, \dots, \gamma_n, f_1, \dots, f_{n-1}\}$.

3. Bayesian Fit to the NG15 Data

3.1. PTA Likelihood and Prior Choices

We are now ready to fit the PPL n -BINNED and PPL n -FREE models to the NG15 data. This dataset spans a total observing time of $T_{\text{obs}} \simeq 16.03$ yr and features pulse TOAs from 68 millisecond pulsars (G. Agazie et al. 2023b). One of these pulsars was observed for less than 3 yr and is not included in our analysis. The Bayesian fit framework employed in this Letter is closely related to the methods used in our earlier work; see our search for signals from new physics (A. Afzal et al. 2023) and our fit of the RPL model to the NG15 data (G. Agazie et al. 2025a). In the following, we will therefore only briefly summarize the key ingredients of our analysis framework and provide more details only when certain aspects of our analysis differ from our approach in earlier work.

The central observables at the heart of our fit analysis are the NG15 timing residuals, which follow from the observed TOAs after subtracting the expected TOAs predicted by the respective timing models. In our Bayesian fit framework, we model these timing residuals in terms of three different contributions: (i) linear offsets from the best-fit values of the timing model parameters, (ii) frequency-independent white (jitter) noise, and (iii) frequency-dependent red noise. On top, we model pulse dispersion, again following A. Afzal et al.

Table 1
Prior Probability Density Distributions for the FREE Parameters in Our Fit of the PPL Models to the NG15 Data.

Parameter	Description	Prior	Comments
Pulsar-intrinsic Red Noise			
A_a^{red}	amplitude	log-uniform $[-20, -11]$	one parameter per pulsar a
γ_a^{red}	spectral index	uniform $[0, 7]$	one parameter per pulsar a
GWB in the PPL n -binned Models			
$h^2\Omega_{\text{GW}}\left(\frac{1}{f_{\text{obs}}}\right)$	amplitude	log-uniform $[-12, -8]$	one parameter per PTA
γ_i with $i = 1, \dots, n$	spectral indices	uniform $[-3, 12]$	n parameters per PTA
GWB in the PPL n -free Models			
$h^2\Omega_{\text{GW}}\left(\frac{1}{f_{\text{obs}}}\right)$	amplitude	log-uniform $[-12, -8]$	one parameter per PTA
γ_i with $i = 1, \dots, n$	spectral indices	uniform $[-3, 12]$	n parameters per PTA
f_i with $i = 1, \dots, n - 1$	node frequencies	log-uniform $[\log_{10} \frac{1}{f_{\text{obs}}}, \log_{10} \frac{14}{f_{\text{obs}}}]^{\text{a}}$	$n - 1$ parameters per PTA

Note.

^a After drawing the node frequencies from this prior density, we needed to sort them according to their size (see the text for more details).

(2023), as a piecewise constant with the inclusion of DMX parameters (Z. Arzoumanian et al. 2015; M. L. Jones et al. 2017).

Similar to our earlier work, we marginalize over the timing model uncertainties and keep the white-noise parameters fixed at the maximum posterior values that were previously determined in single-pulsar analyses (G. Agazie et al. 2023f). The physically interesting contributions to the timing residuals reside in the red-noise terms, which encompass pulsar-intrinsic red noise and the GWB signal. More specifically, the GWB signal is represented by a common-spectrum red-noise term with HD cross-correlation coefficients Γ_{ab} . In the limit of trivial cross-correlation coefficients, $\Gamma_{ab} \rightarrow \delta_{ab}$, the GWB signal reduces to a common-spectrum uncorrelated red-noise (CURN) signal. We will come back to the relation between these two types of signals (GWB and CURN) in our discussion on Bayes factors below.

We model both types of red noise, pulsar-intrinsic red noise and the GWB signal, in terms of discrete Fourier series with harmonic frequencies $F_k = \frac{k}{T_{\text{obs}}}$.⁷⁴ The intrinsic red noise of each pulsar is described by a power law, which amounts to two red-noise parameters per pulsar a : an amplitude A_a^{red} and a spectral index γ_a^{red} . The GWB signal, on the other hand, is modeled by the timing-residual cross-power spectrum in Equation (1), which, together with the relation in Equation (4), can be written as

$$S_{ab}^{\text{GWB}}(f) = \Gamma_{ab} \frac{(H_0/h)^2}{8\pi^4 f^5} h^2\Omega_{\text{GW}}(f). \quad (14)$$

Here, $h^2\Omega_{\text{GW}}$ needs to be chosen accordingly to the PPL model under consideration. Moreover, the GWB contribution to the discrete Fourier series is only sensitive to the values of S_{ab}^{GWB} at $f = F_k$. We include pulsar-intrinsic red noise at all harmonic frequencies across the whole NANOGrav frequency band from F_1 to F_{30} , but restrict the GWB signal to the first 14 harmonic frequencies from F_1 to F_{14} . This is consistent with the

⁷⁴ Here, we choose a notation that clearly distinguishes between the node frequencies in our PPL models, $f_i = \frac{14/n}{T_{\text{obs}}}$, and the harmonic frequencies in the Fourier series, $F_k = \frac{k}{T_{\text{obs}}}$.

observation that the GWB signal in the NG15 data mostly appears at lower frequencies (G. Agazie et al. 2023a).

We caution that the reconstruction of the GWB signal is sensitive to the assumptions that go into the modeling of the pulsar-intrinsic noise. This limitation is, however, a general issue in PTA data analysis: signal and noise power are partially degenerate, and an unambiguous separation of the two is challenging given the strength of the signal in current PTA datasets. In view of this situation, our best option is to stick to the modeling assumptions in earlier analyses of the NG15 dataset, in order to make sure that our results can be compared to the CPL and free spectral reconstructions in G. Agazie et al. (2023a). In ongoing and future work, we will relax these assumptions and generalize our reconstruction of the GWB signal to a broader class of pulsar noise models; see also J. S. Hazboun et al. (2026) for related work on the NANOGrav 12.5 yr dataset.

The power spectrum in Equation (14), alongside the power spectrum for pulsar-intrinsic red noise, determines the covariance matrix of the coefficients in our discrete Fourier series. We are not interested in the explicit realization of the red noise in the NG15 dataset, only its stochastic properties that are described by the parameters (or more precisely, hyperparameters) entering the covariance matrix. We, therefore, marginalize the likelihood for the timing residuals not only over the timing model uncertainties, but also over the Fourier series coefficients. This procedure results in the standard marginalized PTA likelihood (R. van Haasteren & Y. Levin 2013; L. Lentati et al. 2013), which forms the basis of our numerical data analysis.

The marginalized PTA likelihood quantifies the probability $\mathcal{L}(D|\theta)$ of the NG15 data D as a function of the hyperparameters θ , which encompass 134 parameters for pulsar-intrinsic red noise (67 amplitudes A_a^{red} and 67 spectral indices γ_a^{red}) as well as $n + 1$ GWB parameters in the case of the PPL n -binned models or $2n$ GWB parameters in the case of the PPL n -free models. In our Bayesian fit analysis, we need to combine the likelihood $\mathcal{L}(D|\theta)$ with prior densities for all model parameters θ . Our choices for these prior densities are summarized in Table 1. As discussed before, we work with uniform priors for all spectral indices as well as log-uniform priors for all amplitude and frequency parameters. In particular, our prior ranges are chosen large enough such that we are typically able

to capture the entire relevant posterior volume in the high-dimensional parameter space. At the same time, our prior choice for the spectral indices γ_i is supposed to reflect that we consider even smaller or larger indices, i.e., a GW spectrum scaling like $h^2\Omega_{\text{GW}} \propto f^n$ with $n > 8$ or $n < -7$, highly improbable from a physical perspective. The boundaries of this prior range are somewhat arbitrary, which represents a typical issue in Bayesian statistics. On the other hand, our choice for the γ_i prior densities is certainly a conservative one, which allows us to explore the bulk of the posterior density, even in the presence of parameter degeneracies in the PPLn-FREE models. Moreover, a second important factor to consider is that our prior choices are universal across all PPL models under investigation, which is essential for a fair model comparison and the determination of meaningful Bayes factors.

The node frequencies f_i are first drawn from log-uniform prior ranges as indicated in Table 1 and then sorted according to their size, $f_1 \leq f_2 \leq \dots \leq f_{n-1}$, before they are used to evaluate the GW power spectrum $h^2\Omega_{\text{GW}}$. At the level of the posterior density, this algorithm effectively amounts to an analysis in which we simultaneously explore several identical copies of the posterior density that are eventually mapped onto one common posterior density. Consider, e.g., the PPL3 model, which features two break frequencies, f_1 and f_2 . The posterior density that we would like to explore lives in the region of the f_1 - f_2 parameter space where $f_1 \leq f_2$. We can, however, continue this posterior density to the region where $f_1 > f_2$ by exchanging f_1 and f_2 in the likelihood. From this, we obtain a mirrored copy of the original posterior density in the second half of parameter space that we can combine with the original posterior density to obtain one final posterior density. In short, first drawing the node frequencies from independent prior ranges as in Table 1 and then sorting them allows us to avoid a more complicated construction of a multidimensional prior density that accounts for the correct order of the node frequencies from the start.

In passing, we also stress that our treatment of the node frequencies leaves the marginal likelihoods of the PPL models unaffected, which means that it does not distort the Bayes factors that we will need later on in our analysis. Indeed, for PPLn and $n \geq 3$, our analysis operates on $(n-1)!$ copies of the relevant parameter space volume. But at the same time, we renormalize the prior density for the internal node frequencies f_i , such that it integrates to unity, not on the initial parameter space where all node frequencies are sorted, but on the whole enlarged parameter space. This introduces a factor of $\frac{1}{(n-1)!}$ at the level of the prior density, which cancels the $(n-1)!$ factor in the parameter space volume. In practice, this cancellation takes place automatically. We simply have to make sure that the prior density for the internal node frequencies f_i is normalized to unity on the full $\left[\log_{10} \frac{1}{T_{\text{obs}}}, \log_{10} \frac{14}{T_{\text{obs}}}\right]^{n-1}$ hypercube, which is exactly what we do in our analysis (see Table 1).

With the PTA likelihood and prior choices in Table 1 in place, we are now ready to sample from the posterior densities of our PPL models. We do so using standard Markov chain Monte Carlo (MCMC) techniques (J. Ellis & R. van Haasteren 2017) provided by the software package PTArcade (A. Mitridate et al. 2023), a wrapper for ENTERPRISE (J. A. Ellis et al. 2019) and ENTERPRISE_EXTENSIONS (S. R. Taylor et al. 2021). Our main results presented below are based on MCMC runs in

PTArcade’s *enterprise* mode, i.e., Bayesian fits based on the PTA likelihood for the timing residuals. However, for testing and validation purposes, we also make use of PTArcade’s *ceffyl* mode that allows one to perform Bayesian fits based on a factorized likelihood that is induced by the posterior densities of the free spectral model (i.e., by the NG15 violins) (W. G. Lamb et al. 2023).

3.2. Results and Discussion: Individual Models

The results of our Bayesian fits at the level of individual PPL models are shown in Figures 2, 3, and 4. In Figure 2, we present what we refer to as the Ω -PPL and γ -PPL bands for all of our BINNED and FREE models. Slicing these bands at a given frequency, one obtains the distributions of $h^2\Omega_{\text{GW}}$ and γ values at this frequency that are induced by the posterior density in the GWB parameter space. In practice, we construct all Ω -PPL and γ -PPL bands from our MCMC chains as follows: For a given model, we first pick a large number of MCMC samples. Each sample corresponds to a choice of GWB parameters, which in turn fix the precise form of the GWB spectrum. From this spectrum, we then read off the $h^2\Omega_{\text{GW}}$ and γ values at any given frequency. Repeating this procedure for all selected MCMC samples thus results in distributions of $h^2\Omega_{\text{GW}}$ and γ values at all frequencies of interest. The 68% and 95% bands in Figure 2 specifically mark the boundaries of the median credible intervals of 68% and 95% probability in these distributions at each frequency.

In our description of the NG15 timing residuals, the GWB component is only present at frequencies between $f_{\text{min}} = \frac{1}{T_{\text{obs}}}$ and $f_{\text{max}} = \frac{14}{T_{\text{obs}}}$. Correspondingly, the Ω -PPL and γ -PPL bands are only well defined in this frequency range, which we indicate by darker colors in Figure 2. Meanwhile, at smaller and larger frequencies, we extrapolate our results, which we indicate by lighter colors in Figure 2. In practice, we construct these parts of our Ω -PPL and γ -PPL bands by a slight extension of our spectral models, assuming that the GWB spectrum maintains its power-law shape with spectral index γ_1 even below f_{min} and its power-law shape with spectral index γ_n even above f_{max} . In the case of the Ω -PPL bands, the resulting extrapolation highlights the behavior of the bands around f_{min} and f_{max} and the trend in the bands beyond these frequencies. In the case of the γ -PPL bands, on the other hand, the resulting extrapolation simply reflects our underlying assumption, namely, that the spectral index remains fixed at γ_1 for $f < f_{\text{min}}$ as well as fixed at γ_n for $f > f_{\text{max}}$. In all cases, the extrapolated parts of the Ω -PPL and γ -PPL bands serve as a useful visual aid. However, beyond that, they contain no new information, which is why we include them in Figure 2 for purely illustrative purposes.

Next, let us compare our results for the BINNED and FREE models. In the case of our simplest model, PPL1, we find identical results by construction. As mentioned before, PPL1-BINNED and PPL1-FREE denote one and the same model, simply because PPL1 features no node frequency that could either remain fixed or vary freely. In fact, our results for PPL1 are in one-to-one correspondence with the results for the CPL fit to the NG15 data presented in G. Agazie et al. (2023a), where this model is referred to as the HD $^\gamma$ model. The Ω -PPL band for PPL1 in Figure 2 contains the same information as the blue band labeled “power-law posterior” in Figure 1(a) of G. Agazie et al. (2023a), and the γ -PPL band for PPL1 in




Figure 2. Ω -PPL and γ -PPL bands for all individual BINNED and FREE models considered in this work (see the text for details).

Figure 2 represents the same posterior density as the γ_{GWB} posterior density in Figure 1(b) of G. Agazie et al. (2023a).

For all other PPL models, the results of the BINNED models differ from those of the FREE models. In the case of the BINNED models, the Ω -PPL and γ -PPL bands in Figure 2 clearly reflect the fixed frequency bins entering the construction of the models, which always remain the same, no matter the choice of GWB model parameters. In the case of the FREE models, on the other hand, all boundaries between frequency bins are washed out, as the internal node frequencies are allowed to move around. Despite these differences, it is

apparent from Figure 2 that the Ω -PPL and γ -PPL bands for the BINNED models represent reasonable (literally, binned) approximations of their respective counterparts in the FREE models.

All PPL reconstructions of the GWB spectrum are also consistent with the NG15 violins, which we show on top of each Ω -PPL band in Figure 2. In particular, one can ignore the fact that some Ω -PPL bands extend to lower amplitudes than the NG15 violins around f_{max} ; the NG15 violins are simply cut off from below by a prior bound that rises like f^5 in units of $h^2\Omega_{\text{GWB}}$.




Figure 3. Corner plots for the BINNED models. Parameter regions shaded in darker and lighter colors correspond to 68% and 95% credible regions, respectively. At the top of each column, we state median values and 95% equal-tailed credible intervals.

The γ -PPL bands for the BINNED models allow us to construct what may be called γ violins; see the gray violins in Figure 2. These γ violins represent the marginalized posterior densities for the spectral indices γ_i in the BINNED models. In a sense, these violins can thus be regarded as the equivalent—at

the level of the index rather than the amplitude of the GWB spectrum—of the NG15 violins in the free spectral model. Furthermore, we observe that, at frequencies close to $f_{\max}$, the γ -PPL bands in both the BINNED and FREE models often saturate the upper boundary of the γ_i prior range (see Table 1).




Figure 4. Corner plots for the FREE models. Parameter regions shaded in darker and lighter colors correspond to 68% and 95% credible regions, respectively. At the top of each column, we state median values and 95% equal-tailed credible intervals.

This observation reflects the lack of constraining power of the NG15 data at larger frequencies, which is consistent with the behavior of the Ω -PPL bands and the NG15 violins at larger frequencies.

The posterior densities for the actual GWB model parameters, i.e., $\{h^2\Omega_{\text{GWB}}(f_0), \gamma_1, \dots, \gamma_n\}$ in the BINNED models and $\{h^2\Omega_{\text{GWB}}(f_0), \gamma_1, \dots, \gamma_n, f_1, \dots, f_{n-1}\}$ in the FREE models, can be found in Figures 3 and 4. These figures present corner plots,

featuring one-dimensional (1D) marginalized posterior densities in the plots on the diagonal as well as 2D marginalized posterior densities in the off-diagonal plots, for all PPL models of interest. Again, the corner plots for PPL1-FREE and PPL1-BINNED are identical by definition. For all other models, we notice several differences.

The posterior densities in the BINNED models are all well converged and mostly unimodal. In particular, no 2D marginalized posterior density exhibits a flat direction that could in principle be continued far beyond the boundaries of our prior ranges. The posterior densities of the BINNED models are therefore well behaved, which underlines the utility of these models, despite the fact that we merely regard them as an intermediate step in the construction of the FREE models. The situation in the FREE models is more complicated. For these models, multimodal distributions are more common, and parameter degeneracies often result in flat directions that extend all the way up to the boundaries of the prior ranges. Nonetheless, the 1D marginalized posterior densities in Figure 4 indicate that our MCMC chains cover the most relevant regions of the GWB parameter space, which justifies our prior choices in Table 1. Finally, we note that the 2D marginalized posterior densities for pairs of node frequencies (in all FREE models with at least two node frequencies) reflect the correct ordering of the internal node frequencies, $f_1 \leq f_2 \leq \dots \leq f_{n-1}$. In the PPL3-FREE model, e.g., the 2D marginalized posterior density for $\log_{10} f_1$ and $\log_{10} f_2$ is only nonzero when the two frequencies are correctly ordered, $f_1 \leq f_2$.

3.3. Bayesian Model Comparison and Average

Our central goal in this Letter is to perform a Bayesian average over the results for the individual PPL models. To this end, we need to carry out a Bayesian model comparison that allows us to determine appropriate weights (probabilities) $P_n = P(\text{PPL}n|D)$ for each PPL model. According to Bayes' theorem, these model probabilities can be expressed in terms of marginal likelihoods (model evidences) $\mathcal{L}(D|\text{PPL}n)$ and model priors $\pi(\text{PPL}n)$,

$$P_n = \frac{\mathcal{L}(D|\text{PPL}n) \pi(\text{PPL}n)}{\sum_m \mathcal{L}(D|\text{PPL}m) \pi(\text{PPL}m)}. \quad (15)$$

Then, if we assume equal prior probability for all models, P_n can be written as a normalized marginal likelihood,

$$P_n = \frac{\mathcal{L}(D|\text{PPL}n)}{\sum_m \mathcal{L}(D|\text{PPL}m)}, \quad (16)$$

which is sometimes referred to as relative evidence. Next, we can divide by the marginal likelihood of the PPL1 model and express P_n in terms of Bayes factors,

$$P_n = \frac{\mathcal{B}_{n1}}{\sum_m \mathcal{B}_{m1}}, \quad \mathcal{B}_{n1} = \frac{\mathcal{L}(D|\text{PPL}n)}{\mathcal{L}(D|\text{PPL}1)}. \quad (17)$$

In the following, we shall use product-space sampling methods (B. P. Carlin & S. Chib 1995; S. J. Godsill 2001; S. Hee et al. 2016) in order to compute Bayes factors. Product-space sampling, however, works best for similar models that provide fits of the data of comparable quality. Therefore, instead of evaluating the Bayes factors $\mathcal{B}_{n1}$ directly, we first rewrite them in two steps that prove advantageous from the

perspective of our numerical analysis. First, we make use of the fact that Bayes factors are transitive, which means that $\mathcal{B}_{n1}$ can be written as a product of Bayes factors $\mathcal{B}_{k \ k-1}$ for pairs of successive PPL models,

$$\mathcal{B}_{n1} = \prod_{k=2}^n \mathcal{B}_{k \ k-1}. \quad (18)$$

The BINNED and FREE PPL k and PPL $k - 1$ models differ by only one and two parameters, respectively, which facilitates the evaluation of the Bayes factors $\mathcal{B}_{k \ k-1}$. However, we still do not compute $\mathcal{B}_{k \ k-1}$ directly. In a second step, we distinguish between HD and CURN versions of the PPL k and PPL $k - 1$ models and write

$$\mathcal{B}_{k \ k-1} = \frac{\mathcal{B}_k^{\text{HD/CURN}}}{\mathcal{B}_{k-1}^{\text{HD/CURN}}} \mathcal{B}_{k \ k-1}^{\text{CURN}}. \quad (19)$$

Here, $\mathcal{B}_{k \ k-1}$ assumes HD cross correlations in Equation (14) for both models, PPL k and PPL $k - 1$, whereas $\mathcal{B}_{k \ k-1}^{\text{CURN}}$ assumes no cross correlations ($\Gamma_{ab} \rightarrow \delta_{ab}$) in Equation (14),

$$\mathcal{B}_{k \ k-1}^{\text{CURN}} = \frac{\mathcal{L}(D|\text{PPL}k, \text{CURN})}{\mathcal{L}(D|\text{PPL}k - 1, \text{CURN})}. \quad (20)$$

In that sense, $\mathcal{B}_{k \ k-1}^{\text{CURN}}$ is the CURN analog of the Bayes factor that we are actually interested in, $\mathcal{B}_{k \ k-1}$. Meanwhile, $\mathcal{B}_k^{\text{HD/CURN}}$ compares the marginal likelihoods of PPL k in its HD and CURN versions, respectively,

$$\mathcal{B}_k^{\text{HD/CURN}} = \frac{\mathcal{L}(D|\text{PPL}k, \text{HD})}{\mathcal{L}(D|\text{PPL}k, \text{CURN})}. \quad (21)$$

With these ingredients at our disposal, we are now ready to apply product-space sampling, which effectively turns model selection into a parameter estimation problem. Specifically, in order to compute $\mathcal{B}_{k \ k-1}^{\text{CURN}}$, we sample the union of unique parameters of PPL k and PPL $k - 1$ in their CURN versions, alongside an indexing parameter that activates the selected parameter space and likelihood model. $\mathcal{B}_{k \ k-1}^{\text{CURN}}$ is then given by the ratio between the number of samples collected for each model. Furthermore, we make use of likelihood reweighting (E. Payne et al. 2019; S. Hourihane et al. 2023), which applies the idea of importance sampling (J. S. Liu 2008; C. Robert & G. Casella 2010) to the problem of likelihood evaluation, in order to evaluate the $\mathcal{B}_k^{\text{HD/CURN}}$ Bayes factors. In this way, we are able to avoid product-space sampling in this step of our analysis, which would be prohibitively costly because of the HD likelihood entering Equation (21).

In principle, there are two use cases for likelihood reweighting: this technique allows one to model a target posterior by reweighting an approximate posterior, and it yields an estimate of the corresponding Bayes factor. In our case, we shall use likelihood reweighting solely for the latter purpose, i.e., for computing $\mathcal{B}_k^{\text{HD/CURN}}$. To do so, we evaluate the likelihood of PPL k in its HD version on the accepted samples for PPL k in its CURN version that we obtain from our product-space sampling runs. For each sample, we thus obtain a weight factor,

$$w_i = \frac{\mathcal{L}(D|\theta_i, \text{PPL}k, \text{HD})}{\mathcal{L}(D|\theta_i, \text{PPL}k, \text{CURN})}. \quad (22)$$




Figure 5. Bayes factors $\mathcal{B}_{n1}$ for the model comparisons PPLn vs. PPL1, both for the BINNED and FREE models. Normalizing the $\mathcal{B}_{n1}$ by their sum yields the model probabilities P_n [see Equation (17), shown here for the FREE models].

In the CURN versions of the PPL models, we assume the same prior densities as in their HD versions, which allows us to compute $\mathcal{B}_k^{\text{HD/CURN}}$ as the sample mean of the weight factors w_i (S. Hourihane et al. 2023). We stress once more that this is the only part of our analysis in which we resort to likelihood reweighting. All other results for the PPL models in their HD version, in particular all results in Figures 2, 3, and 4, are based on Bayesian MCMC fits employing the full HD likelihoods.

3.4. Results and Discussion: Averaged Models

Our results for the Bayes factors $\mathcal{B}_{n1}$ and model probabilities P_n are shown in Figure 5 and listed in Table 2. The central values and error bars correspond to mean values and standard deviations that we compute using statistical bootstrapping methods (B. Efron & R. Tibshirani 1986). As evident from our results, the NG15 data is best fitted by the PPL1 model: adding more complexity by introducing further power-law segments results in a decreasing Bayes factor $\mathcal{B}_{n1}$. This observation reflects our expectation and justifies the construction of BMAs over a finite number of individual PPL models.

At face value, the preference for the PPL1 model may not appear very significant—Bayes factors between $\frac{1}{2}$ and $\frac{1}{3}$ or so are “barely worth mentioning” on the Jeffreys scale (H. Jeffreys 1939). We, however, emphasize that our models are nested, which means that a model with $n + 1$ spectral indices is always able to reach at least the same likelihood values as a model with only n spectral indices. At the likelihood level, we therefore actually expect an *increase* in the quality of the fit to the data from model to model. A *decrease* in the Bayes factor can only result from an “Occam penalty,” i.e., from the increased dimensionality of the parameter space and the prior volume therein, which is what we see in Figure 5. We therefore conclude that PPLn models with $n > 1$ fail to significantly improve over the PPL1 model at the level of the likelihood, but at the same time receive a penalty due to their larger number of parameters.

This conclusion is supported by our findings in Figure 2: for frequencies roughly between $F_1 = \frac{1}{T_{\text{obs}}}$ and $F_{10} = \frac{10}{T_{\text{obs}}}$, the Ω -PPL and γ -PPL bands of our models are similar to what we obtain for PPL1. Consider, in particular, the 68% γ -PPL

Table 2
Bayes Factors $\mathcal{B}_{n1}$ and Model Probabilities P_n .

Model	Bayes Factor $\mathcal{B}_{n1}$	Probability P_n
PPL1-BINNED	1	0.348 ± 0.017
PPL2-BINNED	0.305 ± 0.031	0.106 ± 0.012
PPL3-BINNED	0.208 ± 0.022	0.072 ± 0.008
PPL4-BINNED	0.612 ± 0.102	0.213 ± 0.037
PPL5-BINNED	0.367 ± 0.062	0.128 ± 0.022
PPL6-BINNED	0.384 ± 0.065	0.133 ± 0.023
PPL1-FREE	1	0.236 ± 0.002
PPL2-FREE	0.887 ± 0.018	0.210 ± 0.005
PPL3-FREE	0.709 ± 0.019	0.167 ± 0.005
PPL4-FREE	0.619 ± 0.019	0.146 ± 0.005
PPL5-FREE	0.452 ± 0.016	0.107 ± 0.004
PPL6-FREE	0.568 ± 0.020	0.134 ± 0.005

bands for the FREE models, which remain rather flat at $f \lesssim F_{10}$: slicing these bands in the horizontal direction by setting the frequency-dependent spectral index γ to a constant value between, say, 2 and 4 is always guaranteed to result in a good fit of the data. Consequently, models with several spectral indices (γ_1, γ_2 , etc.) will always yield a good fit of the data in the special case where these indices are set to the same value, $\gamma_1 = \dots = \gamma_n$, i.e., in the limit of a PPL1-like spectrum.

At the same time, the Bayes factors $\mathcal{B}_{n1}$ in Table 2 provide no statistical grounds to reject any PPLn model with $n > 1$. In summary, this means that the NG15 data is well fitted by the PPL1 model, but more complex models with two or more PPLs are not excluded, justifying once more our BMA approach. Indeed, in our analysis, we consider all viable PPLn models up to $n = 6$. For larger values of n , the PPLn-free models would feature at least as many free parameters as the free spectral model for frequencies up to F_{14} .

In the case of the BINNED models, we observe that the data disfavors power-law breaks at the fixed node frequencies f_1 and f_2 . If these frequencies are allowed to vary, the corresponding Bayes factors $\mathcal{B}_{21}$ and $\mathcal{B}_{31}$ are less suppressed. The gap between the Bayes factors for the BINNED and FREE models only partially closes when a third internal node frequency, f_3 , is added. That is, the BINNED models can only approximately match the flexibility of their FREE counterparts starting from PPL4. At the same time, we observe that the Bayes factors for the BINNED models are always bounded from above by the Bayes factors for the FREE models.

The model probabilities P_n in Table 2 allow us to construct the BMA of the Ω -PPL and γ -PPL bands in Figure 2. To this end, we combine at each frequency f the $h^2\Omega_{\text{GW}}$ and γ distributions for the individual PPL models, weighting each model by its respective P_n value,

$$P(h^2\Omega_{\text{GW}}|D) = \sum_{n=1}^6 P_n P(h^2\Omega_{\text{GW}}|\text{PPLn}, D), \quad (23)$$

and similarly for the spectral index,

$$P(\gamma|D) = \sum_{n=1}^6 P_n P(\gamma|\text{PPLn}, D). \quad (24)$$

In both cases, we evaluate the model weights P_n at their central values listed in Table 2. The outcome of our analysis is shown in Figures 1 and 6. The BMA bands in these figures constitute the main result of this work. As in Figure 2, all bands in




Figure 6. Left panel: PPL reconstruction of the energy density power spectrum $h^2\Omega_{\text{GW}}$ for the GWB signal in the NG15 data, based on the BMA of six PPL models with a fixed number of equidistant internal nodes (PPL1-BINNED,... PPL6-BINNED). Right panel: same, but for the frequency-dependent spectral index of the energy density power spectrum, $\gamma = 5 - d \ln h^2\Omega_{\text{GW}}/d \ln f$. Both reconstructions cover the frequency range from $F_1 = \frac{1}{T_{\text{obs}}}$ to $F_{14} = \frac{14}{T_{\text{obs}}}$, where $T_{\text{obs}} = 16.03$ yr (see the text for details). In the left panel, we also show the NG15 violins in orange, which represent the amplitude posterior densities for the free spectral model. The lower end of the violins at higher frequencies follows from the prior density chosen in G. Agazie et al. (2023a).

Figures 1 and 6 again indicate the 68% and 95% median credible intervals of the corresponding posterior distributions at each frequency f .

In view of our results, several comments are in order. First, it is apparent that the Ω -PPL-BMA bands are dominated by their PPL1 contributions. In other words, the overall shape of the bands is rather straight on log-log axes to first approximation. There is no pronounced peak structure in the sense that the bands first rise with a positive spectral index up to a certain peak frequency and then decline with a negative spectral index. Similarly, there is no double-peak structure, etc. This observation is in agreement with the fact that (i) the model weight for PPL1 is the largest and that (ii) the posterior densities of the PPL n models with $n > 1$ receive large contributions from regions in parameter space where the GWB spectrum is PPL1-like (i.e., where all spectral indices are similar). Second, the BMA bands for the BINNED models reflect again the rigid frequency bins entering the construction of these models. At the same time, the results for the BINNED models serve as an approximation of the results for the FREE models.

Finally, we conclude that the Ω -PPL-BMA bands indeed provide us with a physics-agnostic reconstruction of the signal in the NG15 data that assumes a middle role in between the simple CPL reconstruction and the more conservative free spectral reconstruction. First of all, our Ω -PPL-BMA bands allow for the amount of complexity that is supported by the data. In particular, it is interesting to observe a slight upward trend around $f \simeq F_8 = \frac{8}{T_{\text{obs}}}$, which is also visible in the NG15 violins (violin #8 is very sharp), but which is not captured by the CPL reconstruction. While this feature is statistically not significant, it illustrates the local constraining power of our Ω -PPL-BMA bands, which enable us to constrain local-in-frequency deviations from a globally smooth GWB spectrum. Second of all, our reconstruction of the NG15 signal is more continuous than the free spectral reconstruction. That is, it does not allow for large variations in GWB power from frequency F_k to frequency F_{k+1} , which better reflects our physical expectation for a large class of spectral GWB models.




Figure 7. Bayesian periodograms for the free spectral model and the BMA over the PPL-FREE models. The band in light blue in the background corresponds to the Ω -PPL-BMA band for the PPL-FREE models; see the left panel of Figure 1.

The fact that our PPL reconstructions lead to tighter constraints on the GWB spectrum than the free spectral model can also be illustrated as follows. At each frequency $F_k = \frac{k}{T_{\text{obs}}}$, we can extract the distribution of $h^2\Omega_{\text{GW}}$ values from the Ω -PPL-BMA band for the FREE models and use it to construct what we may refer to as a PPL violin. Together, these PPL violins result in the Bayesian periodogram for the BMA over our PPL models. We show this periodogram in Figure 7 and compare it to the periodogram for the free spectral model, i.e., the standard NG15 violins. Clearly, the PPL violins follow the same trend as the free-spectrum violins. But at the same time, their extent along the $h^2\Omega_{\text{GW}}$ axis is, in most cases, much shorter, reflecting their larger constraining power. We therefore conclude that, as expected, the NG15 data, in combination with the additional assumption of a PPL-type GWB spectrum, lead to stronger constraints on the GWB amplitudes than the more conservative free spectral model.

In particular, our Ω -PPL-BMA band for the FREE models allows us to construct an even larger number of PPL violins. Thanks to the assumption of a PPL-type GWB spectrum, our reconstruction of the NG15 signal provides us with $h^2\Omega_{\text{GW}}$ posterior distributions at every f between f_{min} and f_{max} . To highlight this point, we show 13 additional PPL violins at half-integer multiples of $\frac{1}{T_{\text{obs}}}$ in Figure 7, i.e., in between the violins at $F_k = \frac{k}{T_{\text{obs}}}$. We emphasize that a similar construction is not possible in the free spectral model, which, by definition, only knows about the GWB amplitude at each F_k . In summary, this means that our PPL construction of the NG15 signal leads to a more constrained Bayesian periodogram with a finer (in principle, arbitrarily fine) frequency resolution than the free spectral model.

4. Outlook: PPL Refits

In the literature, numerous spectral GWB models have been proposed as possible interpretations of the NG15 signal. Standard parameter inference on these models requires a Bayesian MCMC fit to the NG15 timing residuals for each individual model, which is computationally expensive. Therefore, in order to facilitate faster model fits to the NG15 data, approximate refitting techniques have been developed in recent years.

The `ceffyl` framework (W. G. Lamb et al. 2023), e.g., builds upon the posterior densities for the amplitude parameters in the free spectral model (i.e., the NG15 violins). The basic idea behind the `ceffyl` framework is that the NG15 violins can be used to construct an induced likelihood on the parameter space of spectral GWB models. In a given spectral model, one simply evaluates the 14 GWB amplitudes $h^2\Omega_{\text{GW}}(F_k)$ for each set of values of the underlying model parameters and reads off the posterior probabilities of these amplitudes from the 14 NG15 violins. The product of these 14 probabilities then defines, up to a constant normalization factor, the factorized `ceffyl` likelihood. The `ceffyl` refitting framework thus consists of two steps: first, one fits, once and for all, the free spectral model to the NG15 data; then, in a second step, one *refits* other spectral GWB models to the amplitude posteriors obtained in the first step.

A similar idea, but for a simpler physics-agnostic reference model, was recently proposed in D. Esmiol et al. (2025), which discusses spectral refits to the RPL model. Again, this refitting framework consists of two steps: first, one fits, once and for all, the RPL model to the NG15 data; then, in a second step, one *refits* other spectral GWB models to the posterior density obtained in the first step. Here, an advantage of the RPL model is the low dimensionality of its parameter space. While the free spectral model introduces a new amplitude parameter for each frequency bin, the RPL model features only three parameters: the amplitude A , the spectral index γ , and the running of the spectral index β . It is therefore easy to construct a numerical approximation of the full three-dimensional posterior density for A , γ , β , e.g., via kernel density estimation (KDE) applied to the MCMC chain that results from fitting the RPL model to the NG15 data. On the other hand, it is now less trivial to map arbitrary spectral GWB models onto the RPL parameter space, which is where the RPL posterior density lives. The analysis in D. Esmiol et al. (2025) solves this problem by introducing a χ^2 -based matched-filtering approach that maps every input GWB spectrum to a triple of best-fit A ,

γ , and β parameter values, i.e., a best-fit RPL proxy for the given input spectrum. This map notably allows one to pull back the RPL posterior density onto the parameter space of the spectral GWB model, where it can then be used as an induced likelihood.

The purpose of this short outlook section now is to point out that, next to refitting spectral models to the free spectral or RPL model, it would be interesting to refit spectral models to the PPL models discussed in the present work. As in the other two cases, the procedure would consist again of two steps: first, one fits, once and for all, all relevant PPL models to the NG15 data; then, in a second step, one *refits* other spectral GWB models to the PPL posterior densities obtained in the first step. The present Letter provides a comprehensive completion of the first of these two tasks. In particular, the corner plots in Figures 3 and 4 show (marginalized versions of) the PPL posterior densities that one would need for a PPL refitting analysis. As in D. Esmiol et al. (2025), one could, e.g., attempt to construct KDE approximations of the respective PPL posterior densities, which live in $2n$ dimensions in the case of the FREE models.

In a second step, one would then need to construct a map from spectral GWB models to PPL-type spectra. We anticipate that χ^2 -based matched filtering might again provide a viable option. Furthermore, one would need to select not only the best-fit points in the PPL parameter space, but also the most appropriate PPL model (CPL, BPL, etc.) to begin with. This additional model selection may also be accomplished by comparing the values of a χ^2 statistic to each other. For now, however, we content ourselves with the fact that the posterior densities obtained in the present work provide the numerical basis for any PPL refit. We leave a discussion of all remaining details and the numerical implementation of a complete PPL refitting framework for future work.

Finally, we conclude by stressing an important conceptual difference between `ceffyl`-style refits to the NG15 violins on the one hand and RPL or PPL refits on the other hand. In the case of refits to the free spectral model, the input GWB spectrum is always already contained in the class of spectra covered by the reference model (i.e., the free spectral model) by construction—at least under the assumptions of the current standard PTA analysis pipeline. The point is that, in the approach followed by G. Agazie et al. (2023a) and A. Afzal et al. (2023), the map from the spectral model to the reference model is trivial: the GWB spectrum is simply modeled by a discrete Fourier series; information in between Fourier frequencies does not enter the analysis. The GWB spectrum is therefore exclusively represented by the set of amplitudes $h^2\Omega_{\text{GW}}(F_k)$, which precisely correspond to the parameters of the free spectral model. The model-specific amplitudes $h^2\Omega_{\text{GW}}(F_k)$ thus play the role of the coordinates of the input spectrum in the parameter space of the reference model. In the case of RPL or PPL refits, the situation is, by contrast, more complex. Most input spectra will not correspond precisely to a specific RPL or PPL spectrum. Hence, in these cases, it is necessary to first establish a map between models that can then be used to pull back the posterior density of the reference model onto the parameter space of the spectral model of interest. An important implication of this observation is that we cannot simply refit spectral models to the PPL violins. Instead, we first have to find the best-fit PPL proxy for a given

input spectrum, before we can then evaluate the PPL posterior density at the location of this proxy spectrum.

5. Conclusions

In this Letter, we presented a new physics-agnostic spectral reconstruction of the signal in the NG15 data that is complementary to the well-known reconstructions based on the CPL model, RPL model, or free spectral model. In our approach, we model the GWB spectrum in terms of a set of PPLs and then marginalize over the number of PPLs in a BMA. The final result of this reconstruction is shown in Figure 1 and can be regarded as the NG15 equivalent of other similar results in the literature, e.g., the PPL reconstruction of the primordial scalar power spectrum by the Planck Collaboration based on the Planck 2015 (Planck Collaboration et al. 2016) and Planck 2018 data (Planck Collaboration et al. 2020). Unlike the CPL reconstruction, our PPL reconstruction of the GWB spectrum accounts for the amount of complexity supported by the NG15 data. In particular, it is sensitive to local features in the GWB spectrum. Conversely, in the absence of local features, which is what we find for the NG15 data, our reconstruction allows us to constrain local-in-frequency deviations from a globally smooth GWB spectrum. At the same time, it offers a more continuous reconstruction of the GWB spectrum than the free spectrum, which better reflects our physical expectation, especially from the perspective of cosmological sources.

Apart from our specific ansatz for the GWB spectrum, our analysis in this Letter relied on the same modeling assumptions as the CPL and free spectral reconstructions of the NG15 signal in G. Agazie et al. (2023a), in particular, the same treatment of pulsar-intrinsic noise. This approach guarantees that our Ω -PPL-BMA band in Figure 1 can be compared to the results in Figure 1(a) in G. Agazie et al. (2023a) in a straightforward apples-to-apples fashion. In ongoing work, we are currently investigating customized chromatic noise models for the individual pulsars in the NG15 dataset and the consequences of this improved noise modeling for the free spectral reconstruction of the GWB signal. Beyond this current analysis, it would be interesting to combine our improved pulsar noise models with a PPL reconstruction of the GWB signal in future work.

We also compared the PPL and free spectral reconstruction of the NG15 signal at the level of their Bayesian periodograms (see Figure 7), i.e., at the level of their amplitude violins, and commented on the possibility to use the PPL posterior densities shown in Figures 3 and 4 in future PPL refit analyses. In particular, we discussed the ingredients that would be necessary to set up a complete PPL refitting framework—an interesting task that we leave for future work. Finally, we complemented every reconstruction of the GWB energy density power spectrum $h^2\Omega_{\text{GW}}$ by an analogous reconstruction of the frequency-dependent spectral index γ . Our final result for the reconstruction of γ is shown in Figure 1. When comparing the two panels in this figure, it is important to keep in mind that the panel for γ does *not* simply follow from taking the logarithmic frequency derivative of the Ω -PPL-BMA band in the left panel. Instead, both the Ω -PPL-BMA band in the left panel and the γ -PPL-BMA band in the right panel represent BMAs over distributions of $h^2\Omega_{\text{GW}}$ and γ values, at each frequency f , that follow from individual PPL spectra.

In summary, we conclude that the PPL reconstruction of the NG15 signal represents a useful addition to the list of commonly employed reference models for the GWB, which motivates us to apply the methodology developed in this work also to future NANOGrav datasets.

Acknowledgments

The authors would like to thank the Mainz Institute for Theoretical Physics (MITP) of the Cluster of Excellence PRISMA+ (Project ID 39083149) for its hospitality and support. The authors would also like to thank Joseph Lazio, Chiara Mingarelli, and Ken Olum for helpful feedback and comments on the manuscript, as well as all members of the NANOGrav New Physics Working Group who provided feedback and comments.

Author Contributions

This Letter uses over a decade’s worth of pulsar timing observations and is the product of the work of many people. K. Sc. initiated and led the project, designed the analysis, and wrote the manuscript. S.W. prepared the first exploratory results. A.B. and W.G.L. carried out the final numerical analysis. W.G.L. prepared numerical data for cross-checks in the `ceffyl` framework. B.D.G.C. and M.L. cross-checked parts of the numerical analysis based on `ceffyl`-type runs. W.G.L. and D.W. contributed to statistical and software development. A.B. and K.Sc. prepared all figures.

Dataset

G.A., A.A., A.M.A., Z.A., P.T.B., P.R.B., H.T.C., K.C., M. E.D., P.B.D., T.D., E.C.F., W.F., E.F., G.E.F., N.G.D., D.C.G., P.A.G., J.G., R.J.J., M.L.J., D.L.K., M.K., M.T.L., D.R.L., J. L., R.S.L., A.M., M.A.M., N.M., B.W.M., C.N., D.J.N., T.T. N., B.B.P.P., N.S.P., H.A.R., S.M.R., P.S.R., A.S., C.S., B.J.S. A., I.H.S., K.St., A.S., J.K.S., and H.M.W. developed timing models and ran observations for the NG15 dataset.

Computing Resources

This work was conducted in part using the resources of the Advanced Computing Center for Research and Education (ACCRE) at Vanderbilt University, Nashville, TN.

Funding Information

The NANOGrav Collaboration receives support from the National Science Foundation (NSF) Physics Frontiers Center under award Nos. 1430284 and 2020265, the Gordon and Betty Moore Foundation, NSF AccelNet award No. 2114721, an NSERC Discovery Grant, and CIFAR. The Arecibo Observatory is a facility of the NSF operated under cooperative agreement (AST-1744119) by the University of Central Florida (UCF) in alliance with Universidad Ana G. Méndez (UAGM) and Yang Enterprises (YEI), Inc. The Green Bank Observatory is a facility of the NSF operated under cooperative agreement by Associated Universities, Inc. The National Radio Astronomy Observatory is a facility of the NSF operated under cooperative agreement by Associated Universities, Inc.

The work of A.B., B.D.G.C., M.L., and S.W. was supported in part by the US Department of Energy under grant No. DE-SC0007914 and in part by Pitt PACC. L.B. acknowledges

support from the NSF under award No. AST-2307171 and from the National Aeronautics and Space Administration under award No. 80NSSC22K0808. P.R.B. is supported by the Science and Technology Facilities Council under grant No. ST/W000946/1. S.B. gratefully acknowledges the support of a Sloan Fellowship and the support of NSF under award No. 1815664. The work of R.B., R.C., X.S., J.T., and D.W. is partly supported by the George and Hannah Bolinger Memorial Fund in the College of Science at Oregon State University. M.C. acknowledges support by the European Union (ERC, MMonsters, 101117624). Support for this work was provided by the NSF through the Grote Reber Fellowship Program administered by Associated Universities, Inc./National Radio Astronomy Observatory. H.T.C. acknowledges funding from the U.S. Naval Research Laboratory. Pulsar research at UBC is supported by an NSERC Discovery Grant and by CIFAR. K.C. is supported by a UBC Four Year Fellowship (6456). M.E.D. acknowledges support from the Naval Research Laboratory by NASA under contract No. S-15633Y. T.D. and M.T.L. received support from an NSF Astronomy and Astrophysics Grant (AAG) award No. 2009468 during this work. E.C.F. is supported by NASA under award number 80GSFC24M0006. G.E.F., S.C.S., and S. J.V. are supported by NSF award No. PHY-2011772. K.A.G. and S.R.T. acknowledge support from NSF CAREER award No. 2146016. D.C.G. is supported by NSF Astronomy and Astrophysics Grant (AAG) award No. 2406919. A.D.J. and M.V. acknowledge support from the Caltech and Jet Propulsion Laboratory President's and Director's Research and Development Fund. A.D.J. acknowledges support from the Sloan Foundation. N.La. was supported by the Vanderbilt Initiative in Data Intensive Astrophysics (VIDA) Fellowship. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). W.G.L. was supported by a Royal Astronomical Society travel grant, which inspired this work. D.R.L. and M.A.M. are supported by NSF grant No. 1458952. M.A.M. is supported by NSF grant No. 2009425. C.M.F.M. was supported in part by the NSF under grant Nos. NSF PHY-1748958 and AST-2106552. A.Mi. is supported by the Deutsche Forschungsgemeinschaft under Germany's Excellence Strategy - EXC 2121 Quantum Universe - 390833306. The Dunlap Institute is funded by an endowment established by the David Dunlap family and the University of Toronto. K.D.O. was supported in part by NSF grant No. 2207267. T.T.P. acknowledges support from the Extragalactic Astrophysics Research Group at Eötvös Loránd University, funded by the Eötvös Loránd Research Network (ELKH), which was used during the development of this research. P.P. and S.R.T. acknowledge support from NSF AST-2007993. H.A.R. is supported by NSF Partnerships for Research and Education in Physics (PREP) award No. 2216793. S.M.R. and I.H.S. are CIFAR Fellows. Portions of this work performed at NRL were supported by ONR 6.1 basic research funding. J.D.R. also acknowledges support from start-up funds from Texas Tech University. K.Sc. is an affiliate member of the Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU) at the University of Tokyo and, as such, is supported by the World Premier International Research Center Initiative (WPI), MEXT, Japan (Kavli IPMU). J.S. is supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship

under award No. AST-2202388, and acknowledges previous support by the NSF under award No. 1847938. J.P.W.V. acknowledges support from NSF AccelNet award No. 2114721. O.Y. is supported by the National Science Foundation Graduate Research Fellowship under grant No. DGE-2139292.

ORCID iDs

Gabriella Agazie  <https://orcid.org/0000-0001-5134-3925>
 Akash Anumalapudi  <https://orcid.org/0000-0002-8935-9882>
 Anne M. Archibald  <https://orcid.org/0000-0003-0638-3340>
 Zaven Arzumianian  <https://orcid.org/0009-0008-6187-8753>
 Jeremy G. Baier  <https://orcid.org/0000-0002-4972-1525>
 Paul T. Baker  <https://orcid.org/0000-0003-2745-753X>
 Bence Bécsy  <https://orcid.org/0000-0003-0909-5563>
 Amit Bhoonah  <https://orcid.org/0009-0005-8851-7286>
 Laura Blecha  <https://orcid.org/0000-0002-2183-1087>
 Adam Brazier  <https://orcid.org/0000-0001-6341-7178>
 Paul R. Brook  <https://orcid.org/0000-0003-3053-6538>
 Sarah Burke-Spolaor  <https://orcid.org/0000-0003-4052-7838>
 Rand Burnette  <https://orcid.org/0009-0008-3649-0618>
 J. Andrew Casey-Clyde  <https://orcid.org/0000-0002-5557-4007>
 Maria Charisi  <https://orcid.org/0000-0003-3579-2522>
 Shami Chatterjee  <https://orcid.org/0000-0002-2878-1502>
 Tyler Cohen  <https://orcid.org/0000-0001-7587-5483>
 James M. Cordes  <https://orcid.org/0000-0002-4049-1882>
 Neil J. Cornish  <https://orcid.org/0000-0002-7435-0869>
 Fronefield Crawford  <https://orcid.org/0000-0002-2578-0360>
 H. Thankful Cromartie  <https://orcid.org/0000-0002-6039-692X>
 Kathryn Crowter  <https://orcid.org/0000-0002-1529-5169>
 Megan E. DeCesar  <https://orcid.org/0000-0002-2185-1790>
 Paul B. Demorest  <https://orcid.org/0000-0002-6664-965X>
 Heling Deng  <https://orcid.org/0000-0002-1918-5477>
 Lankeswar Dey  <https://orcid.org/0000-0002-2554-0674>
 Timothy Dolch  <https://orcid.org/0000-0001-8885-6388>
 Elizabeth C. Ferrara  <https://orcid.org/0000-0001-7828-7708>
 William Fiore  <https://orcid.org/0000-0001-5645-5336>
 Emmanuel Fonseca  <https://orcid.org/0000-0001-8384-5049>
 Gabriel E. Freedman  <https://orcid.org/0000-0001-7624-4616>
 Emiko C. Gardiner  <https://orcid.org/0000-0002-8857-613X>
 Nate Garver-Daniels  <https://orcid.org/0000-0001-6166-9646>
 Peter A. Gentile  <https://orcid.org/0000-0001-8158-683X>
 Kyle A. Gersbach  <https://orcid.org/0009-0009-5393-0141>
 Joseph Glaser  <https://orcid.org/0000-0003-4090-9780>
 Brenda D. Gómez-Cortés  <https://orcid.org/0000-0003-4350-6565>
 Deborah C. Good  <https://orcid.org/0000-0003-1884-348X>
 Kayhan Gültekin  <https://orcid.org/0000-0002-1146-0198>
 C. J. Harris  <https://orcid.org/0000-0002-4231-7802>
 Jeffrey S. Hazboun  <https://orcid.org/0000-0003-2742-3321>
 Ross J. Jennings  <https://orcid.org/0000-0003-1082-2342>
 Aaron D. Johnson  <https://orcid.org/0000-0002-7445-8423>
 Megan L. Jones  <https://orcid.org/0000-0001-6607-3710>

David L. Kaplan  <https://orcid.org/0000-0001-6295-2881>
 Luke Zoltan Kelley  <https://orcid.org/0000-0002-6625-6450>
 Matthew Kerr  <https://orcid.org/0000-0002-0893-4073>
 Joey S. Key  <https://orcid.org/0000-0003-0123-7600>
 Nima Laal  <https://orcid.org/0000-0002-9197-7604>
 Michael T. Lam  <https://orcid.org/0000-0003-0721-651X>
 William G. Lamb  <https://orcid.org/0000-0003-1096-4156>
 Bjorn Larsen  <https://orcid.org/0000-0001-6436-8216>
 Natalia Lewandowska  <https://orcid.org/0000-0003-0771-6581>
 Monica Leys  <https://orcid.org/0009-0000-8694-2872>
 Tingting Liu  <https://orcid.org/0000-0001-5766-4287>
 Duncan R. Lorimer  <https://orcid.org/0000-0003-1301-966X>
 Jing Luo  <https://orcid.org/0000-0001-5373-5914>
 Ryan S. Lynch  <https://orcid.org/0000-0001-5229-7430>
 Chung-Pei Ma  <https://orcid.org/0000-0002-4430-102X>
 Dustin R. Madison  <https://orcid.org/0000-0003-2285-0404>
 Cayenne Matt  <https://orcid.org/0000-0002-9710-6527>
 Alexander McEwen  <https://orcid.org/0000-0001-5481-7559>
 James W. McKee  <https://orcid.org/0000-0002-2885-8485>
 Maura A. McLaughlin  <https://orcid.org/0000-0001-7697-7422>
 Natasha McMann  <https://orcid.org/0000-0002-4642-1260>
 Bradley W. Meyers  <https://orcid.org/0000-0001-8845-1225>
 Patrick M. Meyers  <https://orcid.org/0000-0002-2689-0190>
 Chiara M. F. Mingarelli  <https://orcid.org/0000-0002-4307-1322>
 Andrea Mitridate  <https://orcid.org/0000-0003-2898-5844>
 Cherry Ng  <https://orcid.org/0000-0002-3616-5160>
 David J. Nice  <https://orcid.org/0000-0002-6709-2566>
 Stella Koch Ocker  <https://orcid.org/0000-0002-4941-5333>
 Ken D. Olum  <https://orcid.org/0000-0002-2027-3714>
 Timothy T. Pennucci  <https://orcid.org/0000-0001-5465-2889>
 Benetge B. P. Perera  <https://orcid.org/0000-0002-8509-5947>
 Polina Petrov  <https://orcid.org/0000-0001-5681-4319>
 Nihan S. Pol  <https://orcid.org/0000-0002-8826-1285>
 Henri A. Radovan  <https://orcid.org/0000-0002-2074-4360>
 Scott M. Ransom  <https://orcid.org/0000-0001-5799-9714>
 Paul S. Ray  <https://orcid.org/0000-0002-5297-5278>
 Joseph D. Romano  <https://orcid.org/0000-0003-4915-3246>
 Jessie C. Runnoe  <https://orcid.org/0000-0001-8557-2822>
 Alexander Saffer  <https://orcid.org/0000-0001-7832-9066>
 Shashwat C. Sardesai  <https://orcid.org/0009-0006-5476-3603>
 Ann Schmiedekamp  <https://orcid.org/0000-0003-4391-936X>
 Carl Schmiedekamp  <https://orcid.org/0000-0002-1283-2184>
 Kai Schmitz  <https://orcid.org/0000-0003-2807-6472>
 Brent J. Shapiro-Albert  <https://orcid.org/0000-0002-7283-1124>
 Xavier Siemens  <https://orcid.org/0000-0002-7778-2990>
 Joseph Simon  <https://orcid.org/0000-0003-1407-6607>
 Sophia V. Sosa Fiscella  <https://orcid.org/0000-0002-5176-2924>
 Ingrid H. Stairs  <https://orcid.org/0000-0001-9784-8670>
 Daniel R. Stinebring  <https://orcid.org/0000-0002-1797-3277>

Kevin Stovall  <https://orcid.org/0000-0002-7261-594X>
 Abhimanyu Susobhanan  <https://orcid.org/0000-0002-2820-0931>
 Joseph K. Swiggum  <https://orcid.org/0000-0002-1075-3837>
 Jacob Taylor  <https://orcid.org/0000-0001-9118-5589>
 Stephen R. Taylor  <https://orcid.org/0000-0003-0264-1453>
 Mercedes S. Thompson  <https://orcid.org/0009-0001-5938-5000>
 Jacob E. Turner  <https://orcid.org/0000-0002-2451-7288>
 Michele Vallisneri  <https://orcid.org/0000-0002-4162-0033>
 Rutger van Haasteren  <https://orcid.org/0000-0002-6428-2620>
 Sarah J. Vigeland  <https://orcid.org/0000-0003-4700-9072>
 Haley M. Wahl  <https://orcid.org/0000-0001-9678-0299>
 Si Wang  <https://orcid.org/0000-0002-6019-5511>
 Kevin P. Wilson  <https://orcid.org/0000-0003-4231-2822>
 Caitlin A. Witt  <https://orcid.org/0000-0002-6020-9274>
 David Wright  <https://orcid.org/0000-0003-1562-4679>
 Olivia Young  <https://orcid.org/0000-0002-0883-0688>

References

- Agazie, G., & Anumalapudi, A. 2023, *ApJL*, **951**, L11
 Agarwal, N., Agazie, G., Anumalapudi, A., et al. 2026, *ApJL*, **998**, L11
 Agazie, G., Anumalapudi, A., Archibald, A. M., et al. 2023a, *ApJL*, **951**, L8
 Agazie, G., Alam, M. F., Anumalapudi, A., et al. 2023b, *ApJL*, **951**, L9
 Agazie, G., Anumalapudi, A., Archibald, A. M., et al. 2023c, *ApJL*, **952**, L37
 Agazie, G., Anumalapudi, A., Archibald, A. M., et al. 2023d, *ApJL*, **956**, L3
 Agazie, G., Anumalapudi, A., Archibald, A. M., et al. 2023e, *ApJL*, **951**, L50
 Agazie, G., Anumalapudi, A., & Archibald, A. M. 2023f, *ApJL*, **951**, L10
 Agazie, G., Anumalapudi, A., Archibald, A. M., et al. 2025a, *ApJL*, **978**, L29
 Agazie, G., Anumalapudi, A., Archibald, A. M., et al. 2025b, *ApJ*, **978**, 31
 Antoniadis, J., Arumugam, P., Arumugam, S., et al. 2023, *A&A*, **678**, A50
 Antoniadis, J., Arumugam, P., Arumugam, S., et al. 2024, *A&A*, **685**, A94
 Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2015, *ApJ*, **813**, 65
 Aslanyan, G., Price, L. C., Abazajian, K. N., & Easter, R. 2014, *JCAP*, **08**, 052
 Bécsy, B., Cornish, N. J., & Kelley, L. Z. 2022, *ApJ*, **941**, 119
 Caprini, C., & Figueroa, D. G. 2018, *CQGrA*, **35**, 163001
 Caprini, C., Figueroa, D. G., Flauger, R., et al. 2019, *JCAP*, **2019**, 017
 Carlin, B. P., & Chib, S. 1995, *J. R. Stat. Soc. B*, **57**, 473
 Chen, Y., Daniel, M., D'Orazio, D. J., et al. 2026, *NatAs*, **10**, 554
 D'Orazio, D. J., & Charisi, M. 2023, arXiv:2310.16896
 Efron, B., & Tibshirani, R. 1986, *StatSc*, **57**, 54
 Ellis, J., & van Haasteren, R. 2017, jellis18/PTMCMCSampler: Official Release, v1.0.0, Zenodo, doi:10.5281/zenodo.1037579
 Ellis, J., Fairbairn, M., Franciolini, G., et al. 2024, *PhRvD*, **109**, 023522
 Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T. 2019, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, Astrophysics Source Code Library, ascl:1912.015
 Esmyol, D., Iovino, A. J., & Schmitz, K. 2025, *PhRvD*, **113**, 063016
 Finelli, F., Bucher, M., Achúcarro, A., et al. 2018, *JCAP*, **2018**, 016
 Gardiner, E. C., Kelley, L. Z., Lemke, A.-M., & Mitridate, A. 2024, *ApJ*, **965**, 164
 Ghaleb, A., Malhotra, A., Tasinato, G., & Zavala, I. 2025, *PhRvD*, **112**, 123538
 Godsill, S. J. 2001, *J. Comput. Graphical Stat.*, **10**, 230
 Guzzetti, M. C., Bartolo, N., Liguori, M., & Matarrese, S. 2016, *NCimR*, **39**, 399
 Handley, W. J., Lasenby, A. N., Peiris, H. V., & Hobson, M. P. 2019, *PhRvD*, **100**, 103511
 Hazboun, J. S., Simon, J., Baier, J., et al. 2026, *ApJ*, **1003**, 34
 Hee, S., Handley, W., Hobson, M. P., & Lasenby, A. N. 2016, *MNRAS*, **455**, 2461
 Heimersheim, S., Rønneberg, L., Linton, H., Pagani, F., & Fialkov, A. 2023, *MNRAS*, **527**, 11404
 Heimersheim, S., Sartorio, N. S., Fialkov, A., & Lorimer, D. R. 2022, *ApJ*, **933**, 57
 Hellings, R. W., & Downs, G. S. 1983, *ApJL*, **265**, L39

- Hourihane, S., Meyers, P., Johnson, A., Chatziioannou, K., & Vallisneri, M. 2023, [PhRvD](#), **107**, 084045
- Jeffreys, H. 1939, *The Theory of Probability*, Oxford Classic Texts in the Physical Sciences (Oxford Univ. Press)
- Jones, M. L., et al. 2017, [ApJ](#), **841**, 125
- Kocsis, B., & Sesana, A. 2011, [MNRAS](#), **411**, 1467
- Konstandin, T., Lemke, A.-M., Mitridate, A., & Perboni, E. 2025, [JCAP](#), **2026**, 084
- Lamb, W. G., Taylor, S. R., & van Haasteren, R. 2023, [PhRvD](#), **108**, 103019
- Lamb, W. G., Wachter, J. M., Mitridate, A., et al. 2025, [arXiv:2511.09659](#)
- Lentati, L., Alexander, P., Hobson, M. P., et al. 2013, [PhRvD](#), **87**, 104021
- Liu, J. S. 2008, *Monte Carlo Strategies in Scientific Computing* (Springer Series in Statistics)
- Matt, C., Gültekin, K., Kelly, L. Z., et al. 2026, [ApJ](#), **997**, 188
- Miles, M. T., et al. 2024, [MNRAS](#), **536**, 1489
- Millea, M., & Bouchet, F. 2018, [A&A](#), **617**, A96
- Mitridate, A., Wright, D., von Eckardstein, R., et al. 2023, [arXiv:2306.16377](#)
- Olamaie, M., Hobson, M. P., Feroz, F., et al. 2018, [MNRAS](#), **481**, 3853
- Ormondroyd, A. N., Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2026, [MNRAS](#), **547**, staf2207
- Payne, E., Talbot, C., & Thrane, E. 2019, [PhRvD](#), **100**, 123017
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, [A&A](#), **594**, A20
- Planck Collaboration, Akrami, Y., Arroja, F., et al. 2020, [A&A](#), **641**, A10
- Reardon, D. J., et al. 2023, [ApJL](#), **951**, L6
- Robert, C., & Casella, G. 2010, *Introducing Monte Carlo Methods with R* (Springer)
- Saikawa, K., & Shirai, S. 2018, [JCAP](#), **2018**, 035
- Sesana, A., Vecchio, A., & Colacino, C. N. 2008, [MNRAS](#), **390**, 192
- Shen, E., Anstey, D., de Lera Acedo, E., & Fialkov, A. 2024, [MNRAS](#), **529**, 1642
- Taylor, S. R. 2021, [arXiv:2105.13270](#)
- Taylor, S. R., Baker, P. T., Hazboun, J. S., Simon, J., & Vigeland, S. J. 2021, *enterprise_extensions*, https://github.com/nanograv/enterprise_extensions
- van Haasteren, R. 2024, [MNRAS](#), **537**, L1
- van Haasteren, R., & Levin, Y. 2013, [MNRAS](#), **428**, 1147
- Vázquez, J. A., Bridges, M., Hobson, M. P., & Lasenby, A. N. 2012a, [JCAP](#), **2012**, 006
- Vázquez, J. A., Bridges, M., Hobson, M. P., & Lasenby, A. N. 2012b, [JCAP](#), **2012**, 020
- Watanabe, Y., & Komatsu, E. 2006, [PhRvD](#), **73**, 123515
- Xu, H., Chen, S., Guo, Y., et al. 2023, [RAA](#), **23**, 075024